Pedro Miguel Amaral Melo da Cruz

Semantic Figurative Metaphors in Information Visualization

Doctoral thesis submitted to the Doctoral Program in Information Science and Technology, supervised by Assistant Professor Fernando Jorge Penousal Martins Machado and presented to the Department of Informatics Engineering of the Faculty of Sciences and Technology of the University of Coimbra

May 2016

· U C ·

Universidade de Coimbra

Pedro Miguel Amaral Melo da Cruz

Semantic figurative metaphors in information visualization

Doctoral thesis submitted to the Doctoral Program in Information Science and Technology, supervised by Assistant Professor Fernando Jorge Penousal Martins Machado and presented to the Department of Informatics Engineering of the Faculty of Sciences and Technology of the University of Coimbra

May 2016

· U · C ·

Universidade de Coimbra

This document was composed using Minion Pro, standard web technologies and PrinceXML.

PhD thesis funded under grant SFRH/BD/77133/2011 — FCT (Fundação para a Ciência e a Tecnologia).

No man can live for another. He cannot share his spirit just as he cannot share his body. But the second-hander has used altruism as a weapon of exploitation and reversed the base of mankind's moral principles. Men have been taught every precept that destroys the creator. Men have been taught dependence as a virtue.

—In Howard Roark's courtroom speech
The Fountainhead (1943)
Ayn Rand

A word of gratitude

(in Portuguese)

Quero agradecer antes de mais, ao meu orientador, o Professor Penousal Machado, por todas horas em que aprendi com ele como produzidos, estruturamos e mostramos conhecimento. Foi uma verdadeira jornada em que sintonizámos ideias que resultam não só no trabalho aqui descrito, mas na forma em que é apresentado.

Ao meu amigo João Bicker, que me deixa chamá-lo assim. Foram o companheirismo, a sinceridade e a troca em que aprendi a ver melhor a grandeza do subtil e do elegante, e a tentar perceber o que funciona do que "não funciona". Foram o apoio e o acreditar agudo no meu trabalho. Nesta sequência quero também agradecer ao Alexandre Matos, que também acreditou em mim e abriu-me as portas para uma convivência de saberes que até então não tinha tido.

Dentro do meu grupo de investigação, quero agradecer à Cátia Costa e ao António Cruz pelas contribuições diretas que deram ao meu trabalho, e ao Pedro Martins por poder partilhar os trâmites da escrita desta tese e os seus contextos.

Aos meus grandes amigos de Coimbra, a Olga Filipova, o Luís Pureza e o Alexandre Martins porque mesmo estando longe estão sempre perto, e a minha tese é um período de vida que não pode ser dissociado da convivência, partilha, apoio e gargalhadas que tive com eles.

Aos meus amigos de Boston e Singapura, Prudence Robinson, Anthony Vanky e Inês Dias que me acolheram nas minhas intempéries emocionais e tornaram qualquer lugar ou situação muito mais tranquilos.

À Licia Matos por tanto termos conversado sobre a língua portuguesa, e por todos os caminhos que explorámos juntos. À Ana Cláudia pelo apoio e por todas as fotos que tirámos nos mais pequenos cantos. À Marta Pinto, pelo alegre entusiasmo com o meu trabalho, pelo apoio, e pela a bonita amizade que passámos.

À Isabel Meirelles também por acreditar no meu trabalho, em mim, e me ter aberto tantas portas para novos saltos na vida.

Nenhuma palavra passa mais sem antes passar pelo André Cunha, como tantas vezes acontece. Por ser meu amigo de longa data, meu conselheiro e primeiro confidente. Muitos dos aspetos da minha vida partilhei-os com ele, e sobre eles construímos abstrações, intelectualizações, ou cânones, do que é, do que tem de ser, incluindo a minha tese. Com o André, tem de vir a Alexandra Gouveia, se bem que ela veio antes dele e isso será sempre assim, sem brumas de memórias ou distâncias que o ocluam.

Ao Felipe Shibuya, pelo companheirismo, pela partilha, pelo apoio e por todos os trilhos que percorremos juntos. A história da minha tese, por isso, já teve um final feliz.

À minha mãe, ao meu pai e à minha irmã, com quem cresci e me formei, e que estão sempre prontos para me ajudarem no que for preciso.

Aos meus pais.

Abstract

Visual metaphors are an intrinsic part of visualization. They structure information in order to provide an interpretation framework. *Semantic figurative metaphors* are an approach to information visualization that carries the weight of authorial intent. This approach is suitable for application in the context of casual visualizations, aimed at contemplative and momentary usages. Implementing semantic figurative metaphors encompasses two steps: the adaption of the structural model chosen for a visualization, and the introduction of figurative visual cues. These steps add non-data aspects to the visualizations, introducing inaccuracies that are used for figurative representations of data. These non-data aspects are not gratuitous – they have a semantic meaning for the designer's interpretation of the data. Here, this approach is applied across three visualizations, ranging from topics such as the decline of major empires, traffic in cities, and political-corporate relations in Portugal. The objective is to show the exploration of visual metaphors with the distortion of visualization models to the point of deliberately inserting inaccuracies in data representation, which can be used to emotionally connect with an audience while being clear and faithful to data. The Empires' Decline visualization uses a system of colliding soft bodies to display the growth and disintegration of four empires that vary in size through time. The system was designed to carry metaphors of disintegration, competition, ephemerality and fragility. These visualizations received wide adoption and engagement. The system portrays the empires' areas with an error that is negligible. The city traffic work contributes with a new model for implementing edge-based cartograms. These cartograms are used to depict traffic conditions based on the simulation of a mesh of interconnected springs, using a time-distance metaphor. The edges representing the road network of Lisbon and London expand when velocities are low and contract when velocities are high. The model was tested for several parameterizations in order to obtain results with the best errors. Two additional visualizations were developed and applied on top of the deformed traffic map: the visualization of trajectories that shows vehicles' trails, colored according to their velocities, and a figurative visualization that turns the city into a living organism made of pulsing blood vessels. The Ecosystem of Corporate Politicians shows the relations among Portuguese politicians and companies in an interactive way and adopts a satirical tone. Here, politicians are living organisms that frenetically jump between companies. Data for this visualization was collected and made publicly available. Users can explore companies, catch politicians, and build their own discourses from the data. The visualization reached a large audience and users were engaged by the added metaphorical cues. Finally, these works contribute to a general reflection on the usage of semantic figurative metaphors: designers should make it what reflects their own interpretations, clear; non-data metaphorical aspects should be specifically tailored for a data-message and thus are not usually generalizable; the approach should be part of the design process since its inception – it is not an add-on; figurative representations should be subtle and not obfuscate data.

Keywords
Visual metaphors, information visualization, figurative visualization, visualization design, casual visualization

Resumo

As metáforas visuais são uma parte intrínseca da visualização, estruturando informação para a interpretação. *Metáforas semânticas e figurativas* são uma abordagem à visualização de informação que contemplam intenções autorais. Esta abordagem é apropriada para aplicações no contexto de visualizações casuais que se dirigem a usos momentâneos e contemplativos. Implementar esta abordagem, passa por dois passos: a adaptação do modelo estrutural escolhido para uma visualização, e a introdução de pistas visuais e figurativas. Estes passos adicionam aspetos não relacionados com os dados às visualizações, inserindo imprecisões que são usadas para a representação figurativa dos dados. Estes aspetos não são gratuitos e têm um significado semântico em relação à interpretação que o designer faz dos dados. Aqui, a abordagem é aplicada transversalmente a três trabalhos que englobam tópicos como o declínio de grandes impérios, o tráfego em cidades, e as relações político-empresariais em Portugal. O objetivo é mostrar que a exploração de metáforas visuais juntamente com a distorção de modelos visualização, até ao ponto em que se introduzem inexatidões na representação, pode ser usada de forma a conectar emotivamente a visualização com uma audiência, mas mantendo-se fiel aos dados. A visualização do Declínio dos Impérios usa um sistema de corpos moles que colidem entre si de forma a mostrar o crescimento e a desintegração de quatro impérios, que variam em tamanho ao longo do tempo. O sistema foi desenhado de forma a mostrar metáforas de desintegração, competição, efemeridade e fragilidade. Estas visualizações tiveram uma adoção alargada e mostraram um envolvimento por parte do público. O sistema mostra as áreas dos impérios com um erro que é negligenciável. O trabalho sobre o tráfego em cidades contribui com um modelo novo que implementa cartogramas baseados em arestas. Estes cartogramas são usados para mostrar estados de tráfego, usando para isso uma rede de molas interconectadas, e recorrendo a uma metáfora de distância-tempo. As arestas representam a rede de estradas de Lisboa e Londres, expandindo quando as velocidades são baixas, e comprimindo quando as velocidades são altas. O modelo foi testado para várias parametrizações de forma a obter os resultados com os melhores erros. Mais duas visualizações foram desenvolvidas e aplicadas sobre o mapa deformado das cidades: a visualização de trajetórias que mostra rastos de veículos coloridos de acordo com as suas velocidades, e uma visualização figurativa que transforma a cidade num organismo vivo feito de artérias pulsantes. O Ecossistema Político-Empresarial: mostra as relações entre políticos e empresas portuguesas de forma interativa e adotando um tom satírico. Os políticos são organismos vivos que saltam freneticamente entre empresas. Os dados foram coletados e tornados públicos. Os utilizadores podem explorar empresas, apanhar políticos, e construir os seus próprios discursos sobre os dados. A visualização chegou a uma audiência alargada que se mostrou envolvida pelas pistas metafóricas que foram inseridas. Finalmente, estes trabalhos contribuem para uma reflexão geral sobre o uso de metáforas semânticas e figurativas: o designer deve tornar claro o que é que reflete as suas próprias interpretações; aspetos não relacionados com os dados devem ser adaptados especificamente para uma mensagem autoral, e por isso não são geralmente generalizáveis; a abordagem não é posterior ao processo de design, mas deve ser incorporado nele desde o início; as representações figurativas devem ser subtis e não devem ofuscar os dados.

Palavras-chave

Metáforas visuais, visualização de informação, visualização figurativa, design de visualização, visualização casual

Number formatting: lining and tabular figures

Numbers from zero to nine and when rounding hundreds, thousands, millions, hundred millions and hundred billions, take their cardinal form.

- zero, … , nine.
- one hundred, … , nine hundred.
- one thousand, … , nine thousand.
- one million, … , nine million.

Numbers used to count use lining figures articulated with the above mentioned suffixes.

- 10, 11, … , 99.
- 101, … , 999.
- 1,001; … ; 9,999.
- 10,001; … ; 99,999.
- 100,001; … ; 999,999.
- 9.51 million, 99 million.

Ordinal numbers are adjectives from the first to the ninth, and take their numeral form using lining figures, from the 10th and beyond.

- first, second, …, ninth
- 10th, 11th, …, 21st, …, 343rd

Big numbers are expressed in scientific notation using tabular figures.

- 1.34×10^{9}
- 5.74×10^{10}
- 6.31×10^{11}

Numbers with decimal places, units, or that have an usage other than counting, use tabular figures.

- 63 km/h
- 16 mm
- 87°
- an average of 6.66
- a median of 4

Exceptions

- Percentages (89%, 5%) use lining figures unless the percentages have decimal points in which case tabular figures should be used instead (98.0%, 5.00%).
- Time is expressed in the hh:mm 24-hour format, always using lining figures (e.g. 00:34, 16:38).
- The numbers that identify images, tables and graphs throughout the document are lining figures.

Tables and graphs

In detriment of all the above mentioned, every number inside a table or a graph should be a tabular figure and always be typed as a number without any spell-outs.

Table of contents

1

Introduction

Information visualization is a broad field with practitioners and researchers that incorporate approaches from science, art, and design. The bulk of research in information visualization is scientific. This type of scientific research aims to test and develop computational techniques, tools, and models for visualization. Validating the results of a new technique can involve either comparing it in terms of performance with previous techniques, or it may be grounded in the presentation and visual discussion of its results. Often, certain ways of visualizing data are also discussed through users' evaluation in terms of efficiency and efficacy for perceptual and cognitive tasks (Bertin 2010; Mackinlay 1986; Cleveland and McGill 1984). There is common ground in terms of approach to develop new techniques in visualization. Visualization is often seen as neutral, relinquishing discussions on design choices for communication (Munzner 2008) and assuming that those who produce a visualization must have no authorial intent. This practice of visualization research embodies a reductionist view of visualization. In this view, data are seen as "sacred", where representations should be stripped of any non-data or redundant elements. Embellishments in visualization are viewed with suspicion and as something to completely avoid (Tufte 1983).

Outside research laboratories and corporate settings, there is a diverse and vibrant set of communities made of artists, designers, and coders that produce expressive visualizations for broad audiences (Viegas and Wattenberg 2007b; Lima 2011; Vande Moere and Purchase 2011; Cawthon and Vande Moere 2007; Holmes 1984). Their aim is not to visually solve domain-specific and information-driven problems, or build tools that enable the extensive exploration of a dataset for decision-making. Their aim is to show the data, tell a story that targets a broad audience. They do so in a rich and expressive way in order to raise awareness for a certain subject. They assume that visualization designers can have an authorial role by delivering their personal view on the data. Their main objective is not problem solving outside of communication, but problem solving inside and for communication. The research that comes from these vibrant communities is often lacking, limiting itself to broad and plain descriptions of solutions (Munzner 2008). Visual design solutions can be rich, expressive, complex, and compelling, but are often presented as a set of mystical options, solely grounded in intuitions without constructing a rationale. The construction of a design rationale that logically explains options leading to clear and memorable communication is important to set a common body of practice for visualization design.

This dissertation assumes that the designer of a visualization can have a strong authorial role, and that the resulting visualizations can be described by a rationale that takes into account the data story, the author's message and the visual fundaments for information representation. In order to attain such authorial intent and communicate messages from data in a memorable way, metaphorical cues are incorporated into the visualizations that are, in a strict sense, non-data related, but are connected to the interpretation of data. These metaphorical cues are specific to each visualization application, have the semantic weight of the authorial message and are figurative towards the message and the data, being herein called "semantic figurative metaphors". Semantic figurative metaphors are an approach to designing visualizations that bring into discussion how the author's message acts with the resulting visualization as well as the advantages and pitfalls of incorporating non-data aspects into visualizations.

This approach explores visual metaphors together with the distortion of visualization models, to the point of deliberately inserting inaccuracies in data representation, as a way to portray authorial messages and approximate the audience to the visualization, while being faithful to the data. The applications resulting from this approach are scattered through different problems, but incorporate specific computational systems that dynamically adapt to the dataset, portraying the story and enabling variations in representation.

The visualizations presented here, concern diverse topics, such as the decline of Empires, the traffic in Lisbon and London, and political-corporate relations in Portugal. They present provocative perspectives on common topics, and try to produce an emotional response in the audience while communicating a meaningful data story in a clear way. From a purist's point of view, some design decisions come at the expense of what is regarded as good practice in information visualization – the addition of visual embellishments, the unorthodox use of color, or the deliberate use of effects that can induce errors in the accurate portrayal of information. The counterpoint is made by justifying that these options can be adequate for each problem in question, on a case-by-case basis, given their semantic content in the data and the author's message. Moreover, if errors are inserted in the visualization, they are accounted for in order to ensure that they do not significantly distort data portrayal.

With this, the main objective of the dissertation is to show that the exploration of visual metaphors together with the distortion of visualization models, to the point of deliberately inserting inaccuracies in data representation, can be used to approximate the audience to the visualization while being faithful to the data. Exploring semantic figurative metaphors, poses the following questions:

— Starting with a familiar visualization model for a dataset, which characteristics can be modified in order to introduce a more expressive form of communication?
— If inaccuracies are inserted into the visualization, how can they be used to enhance the communication of certain aspects of the data narrative?
— To which extent inaccuracies are inserted, and what are the risks in terms of data communication?
— If figurative portrayals are sought, when does a certain degree of figurativeness is considered extreme to the point that it detracts from the data itself?

The first visualization application is called "Visualizing Empires' Decline", a dynamic animation that, through subtle variations, is able to slightly vary how the same story is graphically expressed. This is tied with a conceptual framework aimed at "generative storytelling", where actors and events in a data narrative are translated into agents with a set of behaviors in a dynamic system. The story portrayed shows the decline of the four major overseas empires, i.e. the British, the Spanish, the French, and the Portuguese, during the 19th and 20th centuries. The visualization implements specific semantic and figurative metaphors of disintegration and competition. The design approach is described, both on a graphical and a technical level. This system inserts errors in the portrayal of information, which are measured and analyzed. Finally, adoption and reaction to the visualization are discussed.

The second application explores several metaphorical mappings of city traffic. A new type of cartograms called the "edge-based cartogram" shows city traffic from a dramatic perspective. These cartograms are formulated as a new model to depict traffic information by distorting the city's road network. In the model, roads that have higher velocities compress as if distances were shorter, and roads with slow velocities expand as if traffic congestion made perceived distances larger. The model adapts dynamically to data, animating between different data states, and showing a continuous narrative of the evolution of traffic in a city. The cartogram model is tested with traffic data for the cities of Lisbon and London. As with any contiguous cartogram, this model brings errors to data representation. These errors are reported and analyzed in order to obtain the parameterizations that present the best compromise among several types of errors, while also observing how it impacts the overall dramatic effect of expansion and contraction of the city. The cartogram is used as a skeleton map for two additional visualizations. The first is a classical technique where the trajectories of vehicles are drawn and colored according to their speed. The second is more figurative and aims to portray the city as "a living organism", showing a series of intertwined blood vessels that pulse according to traffic conditions.

The third application deals with data pertaining to major Portuguese politicians and their positions in companies through time. The visualization was developed to raise awareness on these relations, but doing so using a satirical tone. In this context, the application was designed with the intent to make this invisible matter more visible, tailoring a provocative visual form that engages viewers through a satirical portrayal of an important matter. The application is an interactive visualization that displays the relations between Portuguese politicians and government positions in a figurative way. The "Ecosystem of corporate politicians" shows living organisms that frenetically jump between different companies and going about their business. The "living organism" has a graphical expression that provides cues on what type of organism they would be, but leaves the representation abstract enough so that the viewers can have their own interpretation. The adoption and reaction to these metaphors is described.

Each of these three works addresses the implications of exploring semantic figurative metaphors. Modifying visualization models in order to make them less orthodox, brings a sense of unfamiliarity that can be explored in order to connect with the audience in an emotional way. Moreover, each application substantiates a design solution for the communication of specific datasets with an authorial intent, but in doing so, they also contribute to a general reflection on the usage of semantic figurative metaphors in visualization.

2

Visualization and its many faces

Employing semantic figurative metaphors targets broad audiences and aims to emotionally connect with them. Their implementations have the functional purpose of clearly communicating data, but doing so by embodying an author's strong point of view. For this, they make use of figurative metaphors, adding non-data aspects to the visualization, which are often seen as mere "embellishments". Having non-data aspects to support the data and an author's message, without obscuring it or inserting visual difficulties to a point that impairs communication, is a design problem. In this context, it is important to discuss how design, metaphors and embellishments are regarded in visualization.

2.1 Design in visualization

> The Chart, No. 1, representing the rise and fall of all nations or countries, that have been particularly distinguished for wealth or power, is the first of the sort that ever was engraved, and has, therefore, not yet met with public approbation.
>
> —(Playfair 1805 p. xv)

> It seems that 300 years after William Playfair's amazement at the cognitive power of information visualisation, others are finally acknowledging the projective power of his method.
>
> —(Manovich 2011, p.37)

Most definitions of information visualization by computer science researchers encompass the usage of computational and interactive visual representations and interfaces.

> The use of computer-supported, interactive, visual representations of abstract data to amplify cognition.
>
> —(Card et al. 1999, p.7)

> Information visualization is the communication of abstract data through the use of interactive visual interfaces.
>
> —(Keim et al. 2006)

> Information visualization utilizes computer graphics and interaction to assist humans in solving problems.
>
> —(Purchase et al. 2008)

These definitions fail to accommodate all kinds of information visualization projects being created today, as they not always make intensive usage of interaction or do not solve specific problems. Information visualization nowadays is scattered between scientific and communicationally driven approaches. For some authors, the classic definitions of information visualization refer instead to "scientific visualization" and it should be distinguished from information visualization of today. "In contrast to scientific visualization, information visualization typically deals with nonnumeric, nonspatial, and high-dimensional data." (C. Chen 2005). According to Manovich (2011), scientific visualization and information visualization are distinct and came from different cultures: science and design. Scientific visualization developed in the

1980s along with the emergence of 3D computer graphics (Friendly 2008, p.40). Contemporary information visualization started its development in the 1990s with the availability of 2D graphics software in personal computers (Manovich 2011). The development of information visualization accelerated in the 2000s, with designers adopting high-level programming languages and APIs (e.g. Processing[1], D3[2]).

For Vande Moere and Purchase (2011), visualization has traditionally focused on supporting expert users in exploring and analyzing complex data as efficiently and effectively as possible, portraying data in a scientific and neutral way. In recent years, information visualization has moved from an advanced research topic to mainstream adoption, encompassing lay users as opposed to a traditional audience of scientists and analysts. For example, Heer et al. (2008) classify a few of the many information visualization projects across this spectrum. There are new users utilizing visualization tools, from expert developers to savvy designers and novice consumers. Additionally, there are new goals in information visualization detached from traditional interactive exploratory analysis. These new goals are communication-driven and do not demand as much effort from its users.

Considering a more general characterization of information visualization, superseding the duality between science and communication, Vande Moere and Purchase (2011) place visualization on three major axes: "visualization studies", "visualization practice", and "visualization exploration". Visualization studies have their ground in academic research and computer science, using systemic inquiry and searching for empirically proven insights. Visualization practice uses the body of knowledge produced in visualization studies to produce visualizations. Visualization practice is seen by Vande Moere and Purchase (2011) as a design practice in the sense that it intends to create workable solutions rather than accumulate knowledge. Visualization as a practice is found in data analysis firms, design studios and individual freelancers. Visualization exploration has more idealistic and visionary goals, often advancing not the techniques themselves, but how these techniques are contextually used and intertwined with their communicational purposes. From this point of view, visualization exploration also covers more artistic approaches to visualization. For example, Viegas and Wattenberg (2007a) have used the term "artistic data visualization" to refer to several data visualizations used for mass communication that employ vigorous perspectives of their authors. This contrasts with scientific visualization where the importance of an authorial role is disregarded, and visualizations are for the most part seen as neutral. In scientific visualization, "developers" build visualizations, and in exploratory visualizations "authors" build their visualizations.

> The value of the artworks rests on the fact that their creators recognize the power of visualization to express a point of view. By contrast, traditional analytic visualization tools have sought to minimize distortions, since these may interfere with dispassionate analysis. Is it possible that this focus on minimizing "point of view" is misguided? For one thing, it is generally impossible to create a vi-

> sualization that is truly neutral, just as it is impossible to create a flat map of the Earth's surface without distorting distances. (…) Perhaps instead of seeking simply to minimize the intrusion of point-of-view, a more realistic attitude for a designer of a visualization should be, as with traditional maps, to choose which perspective is the right one for a given analytic task.
>
> —(Viegas and Wattenberg 2007a, p.190)

This exploratory side of visualization is on par with "casual information visualization", defined as an umbrella term by Pousman et al. (2007) in order to extend the traditional view on information visualization. This complementary view, unifies ambient, social, and artistic visualization domains. According to the authors, there are four main differences that characterize casual information visualization, that range from the characterization of the target audience, to the objectives of authors and the contexts of the visualization:

- The target audience is broad and includes a wide spectrum of users, from experts to novices. The target audience might not be trained in analytical thinking or interpreting visualizations.
- The usage patterns of the users can include momentary, repeatable, or contemplative usages. A contemplative usage pattern would be, for example, a long observation of an intricate visualization in an art gallery.
- The users are more tightly coupled with the data, since they can connect with the messages that the visualization delivers. The data and the message, instead of being solely work-related, have personal relevance, such as providing a provocative view on society.
- The authors of casual information visualization are not strictly interested in providing analytical insights, but instead focus on providing awareness, social, and reflective insights. Such may range from raising consciousness to a particular situation, to providing a different perspective on one's role in society.

Casual information visualization has an increased focus on activities that are less task-driven, datasets that are personally meaningful, and visualizations that are built for wide audiences. According to the authors, this view points to new questions in visualization, such as: "How can we design systems whose highest aims are not focused on productivity, but instead on notions of usefulness, enjoyment, and reflection?" (Pousman et al. 2007).

Vande Moere and Purchase's triadic view on visualization is not discrete, but is instead continuously defined on a triangle with vertices on practice, studies, and exploration (see figure 2.1). In the middle ground of these approaches to visualization is *design*. All of these three approaches contribute with their own unique perspectives on information visualization. The view of Vande Moere and Purchase is that most innovation can occur in the design domain, when the three domains overlap. They then characterize the role of design in information visualization:

[1] https://www.processing.org
[2] http://d3js.org

- Design in visualization is not an "add on" reasoning that can be added after the decisions involved in making the visualization. That is to say that designing a visualization does not only refer to the "visual mapping" stage (Card et al. 1999, p.17), or the "representation" stage (Fry 2004, p.65) in the information visualization process; instead, it should be present in every step, from data acquisition, data processing, up to refinement and interaction. Furthermore, the design process does not have to be streamlined – it may be iterative or non-linear.
- Design in visualization is not a "black box" and can be disseminated. Design is a well-reasoned activity that involves making conscious decisions that are consistent with each other, which have a conceptual ground and can be objectively reproduced. In fact, as pointed to by Muzner (2008), specific solutions for visualizations can be validated by making a case, whereby the devised visual representation is suitable for the domain problem. This rationale should be explicit and not hidden by appropriately justifying visual encoding choices, reflecting on aesthetic considerations, and logically defining the metaphors being used. One example of this is Byron and Wattenberg's (2008) paper on "streamgraphs", where they put several arguments forward that ground their decisions in terms of aesthetics and functionality, such as shape, order, labelling, and color. Their considerations are subjective, but strong and logical. These rationales contribute to a body of generalizable knowledge, and to other cases beyond the one that is described. Byron and Wattenberg's paper is an example of how design can be integrated into research, since it allows peers to assess the visualization process at its core. According to Fallman (2007), this type of integration is described as "research-oriented design" and should be distinguished from "design-oriented research", where design is only present in the dissemination part of a research project. As noted by Vande Moere and Purchase (2011), researchers in design-oriented approaches to visualization, have the scientific responsibility to be as articulate as possible about their rationale in their design processes.
- Design in visualization is not art since it still has a functional purpose that is the clear communication of a data-driven message. As Fallman (2003) notes, design should not be thought as positioned on a continuum between science and art. He emphasizes that design builds on its own body of knowledge, being a "tradition guiding action and thought, which spans across many disciplines" (Fallman 2003, p.231).

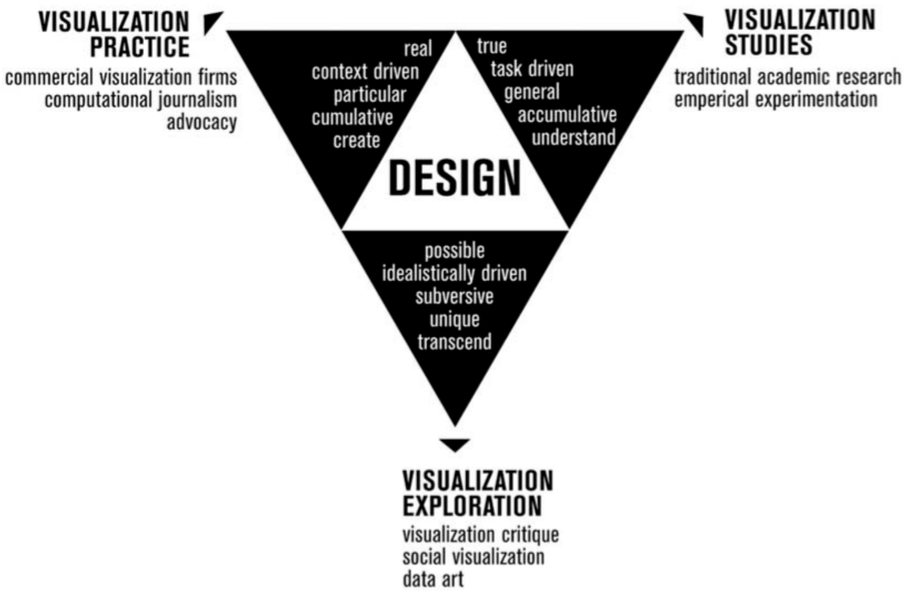

Figure 2.1 The middle ground of design in information visualization, in on a continuum triad among "visualization practice", "visualization studies", and "visualization exploration", as seen in (Vande Moere and Purchase 2011, p.367).

Several visualization applications enjoy wide success and visibility, perhaps due to their design, from the message and the data to the visual representations. There are differences in approach from the traditional visualization community and visualization designers. The later apply design principles that equate both form and function to achieve subjective qualities such as clarity, appeal, and elegant aesthetics. The former applies visual solutions that are empirically proven to be task-effective, hence placing an emphasis on usability, regardless of their intrinsic unappeal. There are clear advantages in crossing those approaches and merging the perspectives and experiences of both communities. Interestingly, merging these different approaches is still unorthodox in some research communities, contributing and agreeing with Vande Moere and Purchase (2011) on the rift between visualization researchers and designers.

2.2 Metaphors

> To draw attention to a philosopher's metaphors is to belittle him – like praising a logician for his beautiful hand-writing. Addiction to metaphor is held to be illicit, on the principle that whereof one can speak only metaphorically, thereof one ought not to speak at all. Yet the nature of the offence is unclear. I should like to do something to dispel the mystery that invests the topic; but since philosophers (…) have so neglected the subject, I must get what help I can from the literary critics. They, at least, do not accept the commandment, "Thou shalt not commit metaphor," or assume that metaphor is incompatible with serious thought.
>
> —(Black 1962, p.25)

A metaphor is a thing regarded as representative or symbolic of something else. In Greek it means “to transfer” and nowadays it still embeds this meaning: the transference of concepts between different domains[3]. Visualization is the formation of a mental image of something[4]. Data visualization, as a practice of providing a visual aid to transform data into mental models (Card et al. 1999, pp.1-3), is intrinsically related with metaphors. These mental models should translate into cognitive amplifications as Card et al. (1999, p.7) note in their classical definition of information visualization: “The use of computer-supported, interactive, visual representations of abstract data to amplify cognition”.

Metaphors are well studied in their verbal form. Using a simple formulation, metaphors can be seen as “two thoughts, of different things active together and supported by a single word, or phrase, whose meaning is a resultant of their interaction” (Richards 1936, p.93). They are usually positioned between two poles: the linguistic and the cognitive (Muller 2008, p.54). In its more linguistic view, a verbal metaphor is only a specific form of language, but from a cognitive perspective, metaphors are forms of thought. Muller (2008, p.57) argues that, independent from the view, metaphors figuratively structure common language and are a basic form of human creativity.

Black (1962, p.35) suggests that figures of speech that involve semantic changes, beyond syntactic changes alone, consist of a transformation of a literal meaning. The author of a metaphor does not provide his intended meaning, but instead the result of a set of transformations. The readers’ task is to employ a convenient set of inverse transformations so that they can decode the original meaning or get close to it. The reader usually decodes the metaphor by either analogy or similarity from the author’s transformations, in order to extract the intended meaning. Naturally, when there is a mismatch between encoding and interpretation, the author can say the opposite of what he meant. A reader’s interpretation is also dependent on a system of cultural commonplaces, and because of this, a metaphor that works in one culture may seem preposterous in another (Black 1962, p.40). Much of the interpretation of certain metaphors are dependent on cultural contexts. For example, Chamfort's[5] (1796, p.26) epigram "The poor are the Negroes of Europe"[6] forces the comparison between the poor in Europe and African slaves in the 18th century. As Black (1962, pp.39-40) notes, the interpretation of this epigram is not a mere substitution of the poor as an oppressed class, or that the poor are discriminated amidst the ideals of society, even that poverty is inherited and ineradicable. The word "negroes" obtains a new meaning that extends its literal sense which, depending on cultural contexts, results in the interaction of two active thoughts that produce a new meaning. This new meaning is the "mystery" of a metaphor, resulting from the dynamic aspects of a reader's interpretation of non-trivial metaphors.

Lakoff and Johnson (1980) pioneered the view of metaphors from a cognitive perspective, framing a theory on "conceptual metaphors". In their view, metaphors are beyond unusual expressive devices and are pervasive in everyday language. Conceptual metaphors map structural properties between a source domain and a target domain and are, a cross-domain mapping process. With this, one can understand one domain in terms of another. The metaphorical expression to convey such process is just a linguistic expression, a surface realization of such cross-domain mapping (Lakoff 1993, p.203). As an example of such mapping, Lakoff (1993, p.207) refers to the "love as a journey" metaphor as a set of correspondences that characterize mapping, where lovers are travelers, the love relationship is a vehicle, and the lovers' goals are their destinations on the journey. It is important to note that rather than a fixed set of correspondences, a metaphor is a dynamic mapping characterized by a set of correspondences, and not the correspondences themselves. The cognitive mapping is not unique to a metaphor, nor is arbitrary. For example, in the metaphor "man is a wolf", some characteristics are mapped from "wolf" onto "man", such as he preys upon others, is fierce, is hungry, etc. (Black 1962, p.41). Other characteristics are ignored from the mapping, such as "animal with tail" or "animal that howls". As Muller (2008, p.57) synthesizes, Lakoff and Johnson's (1980) cognitive theory on metaphors adheres to the following assumptions: language is fundamentally figurative; metaphors are pervasive in ordinary language; metaphors are fundamentally figurative; and metaphors are a general principle of conceptual organization and are fundamental to human understanding. In the cognitive theory, metaphors are not seen as a simple projection of linguistic metaphors or systems of linguistic metaphors for the cognitive system. They are abstract structures of thought that frame how words are produced and comprehended and are thus cognitively real (2008, p.63).

Metaphors can express creative thinking, and that is measured in the novelty of a metaphor. Metaphors operate in a continuum from the most familiar, literal or conventional, to their more novel and figurative forms (Indurkhya 1992, pp.19-20). For example, a metaphor such as "love is a fresh fruit" is more novel and figurative than "Their marriage is a rocky road", which is widely

---

[3] Abate, Frank R, and Elizabeth Jewell. *The New Oxford American Dictionary*. New York: Oxford University Press, 2001. "metaphor".
[4] *Ibid.* "visualization".

used in language but still manages to maintain its interpretative nature. On the conventional and familiar side of the spectrum, are metaphors that are nowadays rendered invisible, but that were once novel, such as "savings will be plowed back in investment" or a "table has legs". As Richards notes regarding the "table has legs" expression, its underlying metaphor can be considered dead because of its pervasiveness in everyday language to the point that it has been absorbed as a literal expression. Nonetheless, it is a metaphor that can instantly come to life just by framing it as a metaphor: "We call it dead, but it comes to life very readily" (Richards 1936, pp.117-118). Most theorists agree that novel verbal metaphors can evolve into conventional language as a culture absorbs the metaphor and transforms the novelty to literality (Cox 2006). This difference between literal and non-literal interpretation of verbal metaphors is gradual and not categorical (2008, p.57).

Indurkhya (1992) sees metaphors as a new meaning created from a mapping from a source domain to a target domain. Each domain is a system of beliefs that Indurkhya calls "concept networks", and "realms". Realms are less abstract interpretations of the concept networks. According to Indurkhya, a metaphor is as metaphorical relation made from a set of correspondences between the target realm and the source concept network. The metaphorical mapping is formed by making correspondences between the source and target concept networks, and then extending these relations to the target realm via its conventional cognitive relations with the target. This process is illustrated in figure 2.2. In this mapping, only parts of the target realm are viewed by the metaphorical relation.

Ortony (1993, p.53) argues that metaphors are a pedagogic device with figurative uses of language, as a way to transfer learning and understanding. In communication, they avoid the problem of enumerating all the characteristics of a domain space, enabling the transfer of a set of often innumerable properties on a specific topic (as he calls them "chunks" of information). He presents a "vividness" thesis on metaphors, arguing that metaphors communicate in a vivid manner, since they are a way of reconstructing experience or experience-like representations. "Metaphors are closer to emotional reality for the same reasons they are closer to perceptual experience" (Ortony 1993, pp.51-52). This brings metaphors closer to perceived experiences since they make emotive, sensory, and cognitive aspects available in the communication of "chunks" of properties of a specific domain.

---

[5] Nicolas Chamfort (1741-1794) was a French writer best known for his epigrams and aphorisms.

[6] Chamfort's original epigram in French on slavery reads "Les Pauvres sont les Negres de l'Europe. Semblable aux animaux qui ne peuvent respirer l'air à une certaine hauteur sans périr, l'esclave meurt dans l'atmosphère de la Liberté." (Chamfort 1796, p.26), roughly translating to "The Poor are the Negroes of Europe. Like the animals who cannot breathe air at certain heights without perishing, slaves die in the atmosphere of Freedom." The usage of the word "negroes" for translating the epigram is also used by Black in 1962 (Black 1962, pp.39-40) in order to give his stance on the creation of new meanings brought by metaphors. The usage of a less culturally charged word would dissociate the metaphor from slavery, which would limit the richness of its interpretation.

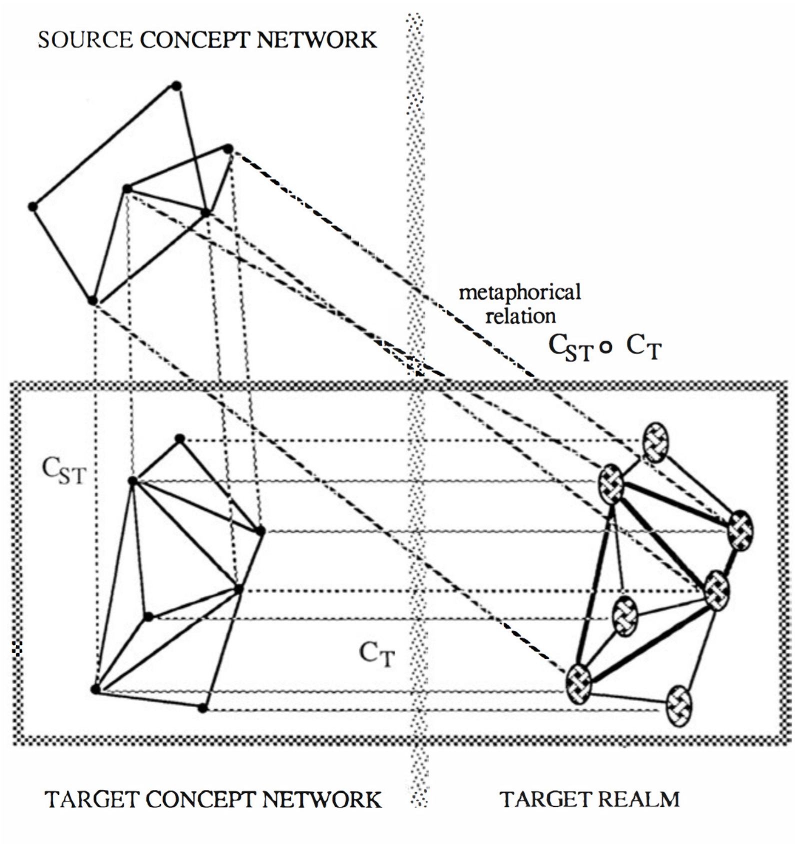

Figure 2.2 A metaphorical relation is formed by first finding the relations between the source and the target concept networks, and then extending these relations to the target realm, which renders visible, in the metaphorical context, only parts of it, as depicted through ticker the lines (Indurkhya 1992, p.259).

Visual metaphors in information visualization

> Creativity is the awareness that alternatives might work, but need to be tested. "Thinking outside the box" requires new perspectives that are practical and visual and make things work. Thinking out of the box is really "thinking out of the metaphor."
>
> —(Cox 2006, p.110)

Designing a visualization is mapping abstract data to a visual space with its own cognitive percepts. Because of this, visualization inherently makes use of metaphors. These metaphors range from simple and seamless to novel and intricate. Structural elements in visualization often use spatial concepts to depict information. For example, an amount can be shown through spatial positioning, size, and length: "down is less", "bigger is larger", "short is low". When working with spatial arrangement in a network, the most important nodes can be positioned near the center in order to depict that those are the most connected nodes. Furthermore, maps and diagrams of a cartographic style, which are a common approach to visualization, also embed spatial metaphors as they translate the spatial domain into an informational one.

Many designs in visualization draw their inspiration from nature, as they metaphorically bring abstract data to the physicality of reality. The natural environment is an excellent source of metaphors for synthesizing complexity in information. During the course of millions of years, nature has adapted its elements into robust complex structures, learning which designs are functional and which are not, and which structures are sustainable in their own context. Many ideas in the representation of the complex are drawn from the visual characteristics of nature and its ecological principles. The subjacent hierarchy in trees, the organic adaption of rivers, the layers of mountains, the motion of the stars, and the arrangement of flowers, exist in a subtle but complex balance and provide an enormous source of inspiration for visualization design.

A traditional discussion of visual metaphors and how they relate with information visualization comes from Cox, who coined the term "visaphor". She argues that there is a direct relationship between data visualization and the cognitive mapping in metaphor theory. These visual metaphors that map and structure data into visual objects are entrenched in culture to an extent that sometimes we no longer recognize their metaphorical nature and thus interpret them as literal or conventional. One example of this is the pie chart, which is about two hundred years old[7] and uses the visual metaphor of sliced portions of a circle as "parts of a whole". This metaphorical interpretation was lost in its cultural absorption, but this conventional way to show data was once novel. Cox presents a framework, that adapts typical concepts from the conceptual metaphor theory to visualization, namely:

- Metaphors in visualization have a target domain and source domain, providing the understanding of the target domain in terms of the source domain.
- The target and source domains represent concept networks that include collections of assumptions, beliefs, symbols, technologies, cultural, and personal biases.
- The mapping of characteristics from the source to the target domain is not a one-to-one mapping since characteristics are selectively included. Nonetheless, this process is not arbitrary.
- New meaning arises with these mappings through novel associations. Some of these metaphors have become embedded in culture so that their metaphorical nature is no longer recognized. Therefore, the metaphors can be seen in a continuum, from conventional everyday metaphors to novel and more figurative metaphors.
- Metaphors in visualization are approximations to data models, and not the data itself, and therefore there are alternative ways to display the same data as much as there are varying interpretations of the same metaphor.
- Finally, the interpretation of metaphors in visualization depends on the audience and the context of the communication.

The classical view of information visualization sees it as a process of encoding numerical values to visual variables, such as position, size, orientation, value, hue, or texture. Bertin (1967; 2010) makes several recommendations on which visual variables are more recommendable to be decoded as one of four informational perceptions: associate, select, sort, and quantify. These variable-encoding models for information graphics focus on how low-level details are interpreted, leaving out more high-level interpretation of features, such as patterns or the interpretation of the overall structure of the visualization. Cleveland and McGill's work (1984) on graphical perception implicitly assumes this view on visualization, and, building on Bertin's theory, they point to certain graphic properties such as angle, direction, area, and curvature as they are more or less suitable for specific elementary perceptual tasks. These two works have since then laid out the theoretical foundation of information visualization, and have been frequently used to devise approaches to visualization design (Mackinlay 1986), or to describe and classify visualization methods in taxonomies (Card and Mackinlay 1997).

Ziemkiewicz (2010, pp.26-29) highlights that high-level percepts from visualizations do not fall under the "reductionist" theory of variable encoding. She proposes that a broader theory is necessary for discussing the effects of visual structure in the communication of information, and suggests that a promising source for such theory are visual metaphors. In her view, the process of making sense of a visualization yields the following steps:

- A viewer first sees a collection of shapes or marks, and can read object-level data, but probably only for one mark or a small number of marks at a time.
- The viewer will be constantly trying to organize the information into a structure by grouping marks based on proximity, connectedness, and similarity.
- Higher-level interpretations of relationships among individual marks and groups are first seen as purely visual patterns, but are then mapped to conceptual relationships based on visual metaphors.
- These metaphors are either directly suggested by the visualization, taught through instructions, or in the absence of instructions, invented by the viewer.

Ziemkiewicz and Kosara (2008) see metaphors as "a set of structural properties that provide a framework for meaning". They point out that the nature of information visualization can lie on the visual metaphors that it uses to structure information. Subtle differences in visual or verbal metaphors can suggest different interpretations of the same information. Their study shows that compatibility of visual metaphors with the verbal metaphors in the mind of the viewers influences both correctness and response time when analyzing certain visualization models. Inducing viewers to think of a metaphor that is incompatible with the source information can slow their response time for information decoding tasks. For example, a treemap uses a metaphor of hierarchy as containment. If a viewer conceives such representation as a series of boxes nested inside each other, then the correspondence to containment should be natural. On the other hand, if the viewer conceives the treemap as a hierarchy of higher and lower levels, the correspondence to containment will be more difficult. In visualizations that naturally make use of the hierarchical metaphor, such as node-link diagrams, interpreting them with a containment metaphor such as "elements inside elements", instead of "elements parenting

[7] One of the earliest pie charts found in information graphics is from William Playfair (1801, pp.59-60)

other elements", will slow the process of decoding information. This is why the compatibility between the inherent visual metaphor of a visualization model, and how it is framed in terms of verbal metaphors, is fundamental. Other authors detected subtler metaphors when viewers decode a graph. Zacks and Tversky (1999) found that when presenting the same two-point data either in the form of a bar chart or a line graph, the users associated the bar chart as depicting separate groups whilst the line graph was associated with the data showing a trend. This suggests that a line graph is preferred over a bar chart when the points should be decoded as connected and representing an overall trend for the same variable.

The idea that visual metaphors are used to structure information is closely related to how people build a mental model of a system. In human-computer interaction, the way that people derive a mental model from an interface is viewed in terms of perceived affordances (Norman 1999). Affordances refer to visual aspects of an interface that suggest actions to a user, such as an interface element that is portrayed as a physical button, which suggests that it can pressed. The interpretation of an interface is therefore metaphorical, suggesting interactions from the physical domain.

Visual metaphors are used in several visualization techniques. For example, edge bundling (Holten and Wijk 2009; Holten 2006) manipulates the edges of a graph in order to show connection flows within a network. The nodes on a graph are drawn in a way to cluster them together as if they were flows of connections. Similarly, the ThemeRiver (Havre et al. 2000) and the Streamgraph (Byron and Wattenberg 2008) layouts also embody the metaphor of a sequence of events as a flow of liquid. They manipulate the layout on a stacked area chart in order to increase its vertical evenness, by using smooth curves that resemble organic flows of liquids.

Another metaphor for the creation of visualizations is related to the fractal architecture of natural systems (Judelman 2004). In nature, complexity spreads across several scales, from the microscopic to the galactic. On a figurative level, as Lima (2011, p.15) points out, several abstract network structures such as the blogosphere, food webs, airline routes or protein chains may resemble biological neural networks. On a conceptual and functional level, the infinite detail of a fractal structure may be present to a certain extent in zoomable maps. Zooming reveals the magnitude and complexity of the universe being portrayed. In fact, current digital maps usually provide a great level of detail close up, and adjust to different information resolutions depending on scale. This property is often referred to as adaptive zoom, since elements are removed or merged when zooming out and added or portioned when zooming in, adding greater detail. On a par with zooms and maps, other techniques exist as an alternative to the adaptive zoom strategy. For example, zooming in on a map can preserve all the elements in a map only if the portion to be zoomed in on is magnified through a seamless spatial transition between the zoomed and the non-zoomed spaces (Keahey and Robertson 1996). This transition is a metaphor of the map through a physical zoom lens. One of such examples is the fisheye view (Furnas 1986) that delivers both a zoomed out portion of the map together with a magnified area with a seamless spatial transition between both.

Visual metaphors can be used for more than structuring information. They can be figurative to a point where they force the association of the source domain with very specific concepts that can attribute new meanings to data. One example of this is Chernoff faces (Chernoff 1973). This model displays multivariate data-points by associating facial features such as the shape of the head, the size of the nose, and the curvature of the mouth, with data features. For example, a face can be used to portray socio-demographic properties of a certain population. If the curvature of the mouth is associated with GDP per capita, where high incomes are transformed into smiles, new meanings to data arise, such as a population being happier or sadder than another (see figure 2.3). Chernoff faces are not frequently used since its inherent metaphor of human emotions is often incompatible with the nature of the dataset. Because of this, they are not a conventional way to depict data and with this they carry a sense of novelty that positions them in the figurative side of the metaphoric continuum.

> …with cartoon faces and even numbers becoming data measures, we would appear to have reached the limit of graphical economy of presentation, imagination, and let it be admitted, eccentricity.
>
> —(Tufte 1983, p.142)

Comparatively with Chernoff faces, the "We Feel Fine" (Kamvar and Harris 2011) project embodies a less figurative take on visual metaphors. "We Feel Fine", is an application that harvests and displays emotions on the blogosphere in a series of visualizations. The authors' metaphorical intent for emotions is clear: individual data points are represented as a "playful" swarming particle system in order to show a "bird's eye view of humanity", and bars in a bar graph are replaced by large jiggling bulbous mounds. Neither the particles' positions nor the shape and jiggle of the bars are strictly data related, but instead reflect the authors' own view of the data and how they visually interpreted a live stream of actual people's feelings. These are concretizations built from the authors' own interpretation of feelings, as if the feelings and data were alive, providing a more compelling imagery for the audience to connect.

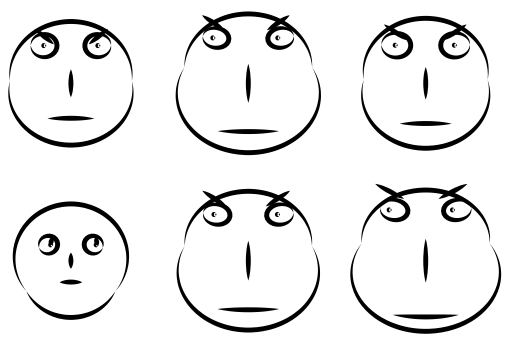

Figure 2.3 An illustration of Chernoff faces that can be used to communicate new meanings within data, such as "happy", "angry", "relaxed", etc.

2.3 Embellishments

Information visualization is a broad field that encompasses approaches from science, design, and art. Since the earliest statistical graphics, graphs are often presented with purely decorative elements, or added elements that do not mirror the data directly, but that instead have the purpose of strengthening a certain high-level perspective from data. The traditional view of information visualization practice is a purist one, and defends neutrality of data-graphics, removing any decorative and even emotional elements. Nevertheless, there are other emerging perspectives on the practice of information visualization that allow the usage of decorative elements, or embellishments, with the intent of strengthening narratives for a broader audience.

Edward Tufte is a scholar well known for introducing a series of principles to graphical excellence and integrity (Tufte 1983). Amongst other rules of thumb, he writes "above all else show the data", defending that within reason, non-data ink or redundant ink should be erased as a way to maximize the data-ink ratio (Tufte 1983 ch. 4). Tufte introduced the term "chartjunk" in his first book, referring to unnecessary graphical elements that add no value to data graphics: "The interior decoration of graphics generates a lot of ink that does not tell the viewer anything new... Regardless of its cause, it is all non-data-ink or redundant data-ink, and it is often chartjunk" (Tufte 1983, p.107). He adds that often, chartjunk does not involve artistic considerations, and it is simply conventional paraphernalia that is added routinely to data graphics. He states that good data graphics are dependent on good datasets, and that overly simple data does not benefit from the addition of decorative elements.

> Chartjunk can turn bores into disasters, but it can never rescue a thin data set. The best designs... are intriguing and curiosity-provoking, drawing the viewer into the wonder of the data, sometimes by narrative power, sometimes by immense detail, and sometimes by elegant presentation of simple but interesting data.
>
> —(Tufte 1983, p.107)

His remarks emerged in a context when many statistic and information graphics, aimed either at the scientific community or the general public, frequently incorporated decorative elements without any semantic value, such as adding an unnecessary third dimension to a chart, recurring to heavy patterns that can cause moiré vibrations (see figure 2.4) and making use of unnecessarily heavy ink. Ellen Lupton, a well known design educator addressed Tufte's vision, presenting a counter-perspective.

Tufte's purist point of view is profound and compelling, but it may be overly restrictive. Information graphics have a role to play in the realm of expressive and editorial graphics... They can be clean and reductive or richly expressive, creating evocative pictures that reveal surprising relationships and impress the eye with the sublime density and grandeur of a body of data.

—(Lupton and Phillips 2008, p.199)

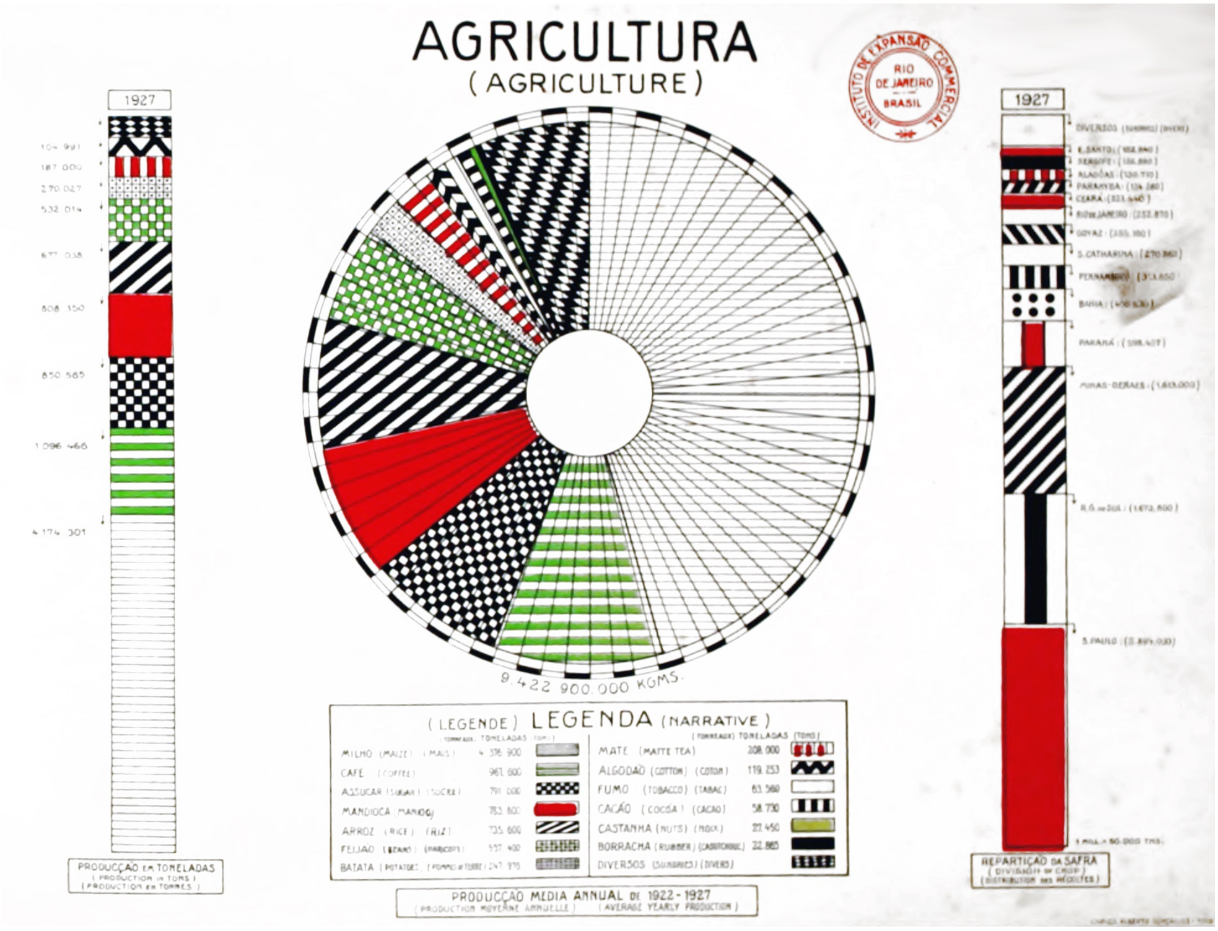

Figure 2.4 Several graphs portioning the agricultural products of Brazil ("Brasil: Graphicos Econômicos - Estatísticas" 1929, p.5). This example employs heavy patterns that, according to Tufte, are non-data ink and should be erased as they also may cause moiré vibrations when observing the graphs.

On this point, the nemesis of Tufte is perhaps Nigel Holmes. Holmes is best known for the graphics he did for TIME magazine from 1978 to 1994. His graphics are often very figurative, merging data depictions into illustrations that try to reinforce a specific message (see figure 2.5). Often, the datasets the he illustrates are fairly small, so a question ponders whether this particular style would be adequate for datasets that demand a more complex visual display. Holmes uses humor to illustrate data that tells a story (Holmes 2014). "A good approach to information graphics includes an appeal to the reader, immediately followed by a true accounting on the story, whether it is statistical, geographic, or diagrammatic. (…) I want to make room for enjoyment, delight, aesthetic appreciation and wit, and a friendly 'you can understand this' approach. I believe all this can be done at the same time as delivering the data." (Heller 2006, p.78). But Tufte, in his book "Envisioning Information" (Tufte 1990), refers to a colored version of Holmes' chart in figure 2.5 in the following way.

> Consider this unsavory exhibit (…) with cliché and stereotype, coarse humor, and a content-empty third dimension. It is the product of a visual sensitivity in which a thigh-graph with a fishnet-stocking grid counts as a Creative Concept. Everything counts, but nothing matters. (…) Lurking behind chartjunk is contempt both for the information and for the audience. (…) Credibility vanishes in clouds of chartjunk; who would trust a chart that looks like a video game?
>
> —(Tufte 1990, p.34)

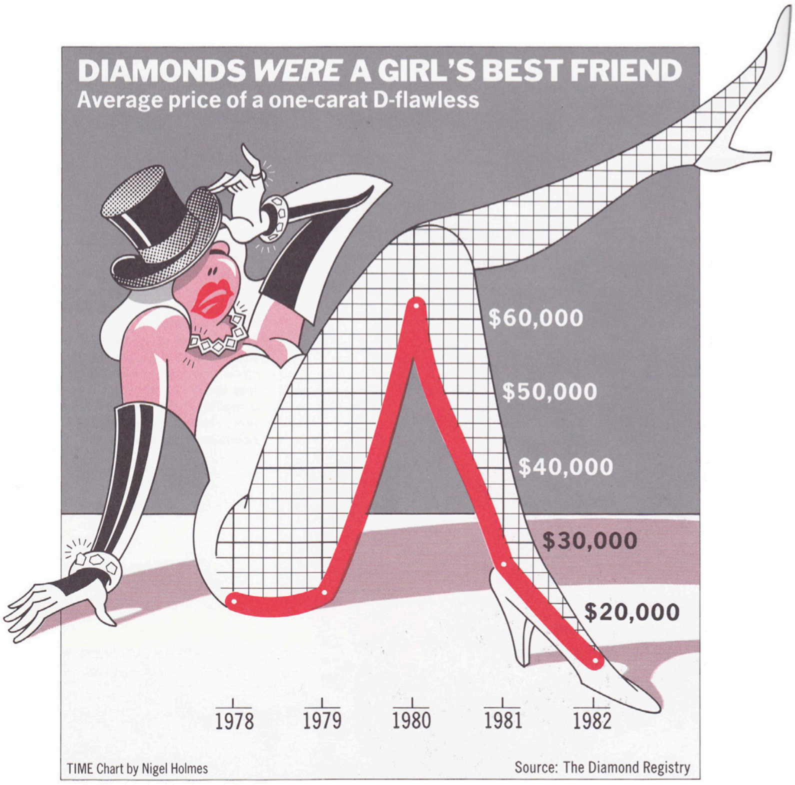

Figure 2.5 Diamonds were a girl's best friend. TIME chart by Nigel Holmes, as it appears on (Holmes 1984, p.32)

Holmes reacted to Tufte's criticism of his graph, emphasizing that it was taken out of the context of TIME magazine and its readership: "I believe his thought that everyone is going to spend as much time with a chart as he imagines is unfounded in today's rushed society (certainly in magazines, newspapers, and on the web, anyway)" (Heller 2006, p.77). He considers that Tufte's ideas of data-ink removal can lead do dryness if taken to extremes (Heller 2006, p.77). In a recent communication, Holmes clearly puts his thoughts to his approach to information design.

> I found that the best way to connect, was making people smile. Not funny, not laugh, but a smile of recognition of something. I'm trying to be friendly, to be approachable, to tap into other people's emotions, to use humor in the sense of good humor, of good feelings, of good feelings and emotions. (...) Emotions are important. (…) Some think that I'm trivializing information. But I'm not. I'm simply trying to popularize information. (…) Just because a thing is serious, does it make it an authority? I ask you. Or just because a thing is light-hearted does it mean it's not serious?
>
> —Transcript of Holmes (2015)

Stephen Few, also a proponent of minimal-style information graphics (2009), considers that at times non-data and redundant ink can be used in useful ways to support the display of information if properly chosen and designed (Few 2011b). Nevertheless, the appearance in recent literature of studies and theories that suggest that embellishments can be used in information graphics as a rule of thumb receives Few's criticism.

Bateman et al. (2010) presented a study where they measure interpretation accuracy and long-term recall of some of Holme's graphics, comparatively to more minimalistic counterparts. They found that people's accuracy in describing Holme's embellished charts was no worse than for plain charts, and that their recall after two to three-week gap was significantly better. As Few (2011b) notices, this study should not be seen as a generalizable rule of thumb on using chartjunk or embellishments in information graphics, rather, it should be seen as only stating that Holme's particular style, that supports the chart's message, can succeed as a useful mnemonic in very specific cases. He adds that Holme's style charts are "well executed examples of a particular style of embellishment (not entirely chartjunk)", and that "Embellishments that are not data in themselves can sometimes support the display of data in a useful manner" (Few 2011b).

Hullman et al. (2011) published a theoretical article suggesting that visual difficulties, in a generalized way, may benefit comprehension and recall by information graphics viewers. They found this claim in psychological and pedagogic research that suggests that "more difficult to interpret graphs and other visual materials yield better learning" (Hullman et al. 2011). Few (2011a) counterpoints this view stating that "What the authors cite (…) primarily suggests that it is engagement with information that leads to reflective thinking, which actually produced the positive outcomes that they mistakenly attribute to visual difficulties." (Few 2011a). The authors indeed, at a certain point, center on engagement but say that it can be induced by visual difficulties. They suggest that such engagement can be increased by strategies such as novelty, aesthetically appeal, tailoring and personalization, challenges and game-play. Few (2011a) notes that these aspects are not fundamentally visual in nature and "don't qualify as cognitive difficulties" (Few 2011a). Hullman et al. thesis is provocative, suggesting that chartjunk can be introduced in the form of visual difficulties to aid comprehension. According to the authors, visual difficulties can be applied in visualization by including complexity or, animation, by diminishing the data-ink ratio or by using less orthodox fonts. There is a missing link in the article on whether these aspects are intrinsically visual difficulties, or if they are rather visual aids to other aspects, namely for the purpose of increasing engagement. Hullman et al. central theme is compelling, but the conceptual choice of wording "visual difficulties" as a generalization to provoke more user engagement is unfortunate. In fact, the authors start with a pertinent viewpoint, noting that information visualization research tends to assess how a user interacts with a visualization strictly based on efficiency, seeking minimal response times and high response accuracies. This way, evaluation does not usually hinge on the assessment of important concepts or patterns learned by the users. This is another problem acknowledged by Few (2011a): "Information visualization research should test relevant outcomes, not merely measure decoding accuracy and efficiency". This way, the authors' aim of proposing a model of effective graph design as a trade-off between efficiency based approaches and concept-learning approaches is valuable, but should be separated from the general umbrella of "visual difficulties".

Borgo et al. (2012) presented an empirical study that examines the effects of visual embellishments in visualization processes. The authors evaluated the results of users on four cognitive and perceptual tasks: working memory, long-term memory, visual search, and concept grasping. The embellishments used alongside their versions of plain-graphs, mainly included figurative icons placed on familiar visualization models, such as bar charts, line charts, or bubble charts. Their results show that visual embellishments aid performance in the working memory task, having a positive impact on accuracy and response time. Furthermore, it showed that visual embellishments also have a positive impact on the speed of long-term recall. On the other hand, the increase in visual detail caused by embellishments negatively impacts the performance of users in the visual target task. In relation to concept grasping the results were inconclusive, but suggest that visual embellishments do not cause a negative impact, similarly to (Bateman et al. 2010).

Regarding a very specific use of embellishment in information graphics, Wood et al. (2012) presented a computational technique to render sketchy data graphics, mimicking data graphics drawn by hand. These stylized renderings of data not only inserted additional embellishments, but also distorted the shapes of the graphics (see figure 2.6). Their study asked participants to judge relative sizes across circles and rectangles, with varying degrees of sketchiness. The results showed that sketchiness impaired their ability to make judgements about relative sizes, being imprecise and error prone. Nonetheless, their experiments also suggest that sketchiness can be reliable to encode visual variables on an ordinal scale. In a second experimental setup, users were asked to annotate plain charts as well as sketchy renderings of bar charts, maps, area charts, pie charts, and network diagrams. As demonstrated by their positive high-level annotation tasks, the results show that sketchy renderings of data graphics can offer greater engagement, although, as the authors state, "It is not clear to which degree such responses may be due to novelty of the visual appearance as opposed to intrinsic properties of sketchy design" (J. Wood et al. 2012).

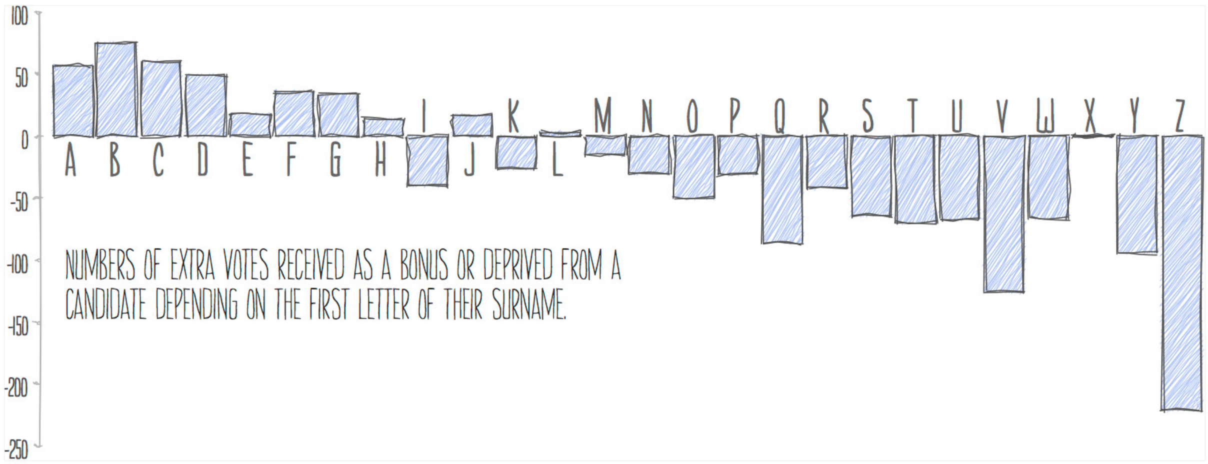

Figure 2.6 A sketchy bar chart as seen in (J. Wood et al. 2012).

Borkin et al. (2013) performed the largest scale visualization study to date, which tries to address "What makes a visualization memorable?". They found that the inclusion of recognizable human objects such as scenes, objects and people, on par with color, consistently enhance the memorability of visualizations. Similarly to previous studies, they also found that visualizations with low data-ink ratios and high visual density were more memorable than minimal visualizations. Furthermore, unique visualization types, that are more pictorial, such as grid-matrices, trees, networks, or diagrams, were significantly more memorable than more common graphs such as bar charts, line graphs, area graphs, scatterplots, or circle charts. The claims of this study regarding memorability in visualization should be considered with caution. The experimental setup involved displaying the visualization only during one second, with a 1.4 second gap between consecutive images. The tests were formulated as a game were the objective of the user was to swiftly identify and click on a visualization that appeared a second time. Few (2013) points out that the study did not determine what makes a visualization memorable, but that it addresses which visual elements or attributes in the visualization would be noticed when viewed for one second and then recognized a second time. A visualization has content, and its purpose is to communicate that content; the memorability of a visualization as a whole should take its content into consideration, and in a one second, content can be difficult to interpret and assess. More specifically, what the authors are evaluating is "stickiness" on a perceptual level (C. Heath and D. Heath 2008, p.8), i.e., visual aspects in the visualization that sticked because they stood out as memorable. Even if the authors broadly use the term memorability in this context, the article acknowledges that there are other major aspects at play in memorability, and that the study aims to understand what aspects of a visualization move it to the memorable end of the spectrum of memorability: "We do not want just any part of the visualization to stick (e.g., chart junk), but rather we want the most important relevant aspects of the data or trend the author is trying to convey to stick" (Borkin et al. 2013).

A study of Skau et al. (2015) evaluates the impact of embellishments in bar charts in absolute and relative judgements of quantitative values. The results show that triangular bar charts and bar charts with rounded tops were far worse at communicating data than standardized baseline bar charts. For example, by using thin triangles that are scaled non-uniformly to reach the bars' height, the absolute and relative judgments of quantities can be 29% to 90% worse, respectively, than a baseline bar chart. Other frequent embellishments such as rounded corners are 26% to 42% worse for absolute and relative judgments, respectively. The authors acknowledge the previous studies of Borkin et al. (2013) and Borgo et al. (2012), adding that "Changes to charts that affect the primary chart elements can reduce the communication accuracy of the chart. Increasing the memorability of a chart is certainly a worthwhile pursuit, however it must be balanced with the need to communicate information accurately in the first place" (Skau et al. 2015).

The study of Haroz et al. (2015) tests how pictographs impact working memory, speed of finding information, and engagement with information graphics. Inspired by pictorial unit bar charts (Brinton 1939, pp.121-131) (see figure 2.7) and the ISOTYPE system (Burke et al. 2014), the authors tested how pictograms perform when used as labels, when stacked as arrays of quantitative information, scaled to represent a certain variable or simply when used as decorative backgrounds. The results for several tasks were compared with simple bar charts and bars made of stacked circles. They found no strong evidence that using pictograms to communicate information impairs performance in any of their tasks, stating that "Our work adds more evidence to the claim that not all kinds of 'chart junk' are necessarily detrimental" (Haroz et al. 2015). Additionally, from their experimental results, they devised a series of guidelines.

- Superfluous pictographs are a distraction – Including unnecessary imagery, such as in the background, appears to divert from the data, and replacing text labels with pictograms makes encoding less efficient. Therefore, only pictograms that are part of the data mapping can be beneficial or at least not harmful.
- Redundantly coding small lengths and small numbers – Length-defined objects should be broken into smaller units, allowing the user to make number estimates rather than length judgements. These units can be pictograms or can result from the division of Tufte's style gridlines (Tufte 1983, p.116). It is important to notice that such should only be done for small quantities.
- Use pictographs for demanding tasks – Presenting successive visualizations of varying information puts the working memory under load. In

this case, information can be recalled more accurately with stacked pictograms than with simple bar charts.

— Pictograms engage readers – People tend, at least initially, to direct their focus towards more figurative data representations, inciting readers to inspect the visualization more closely.

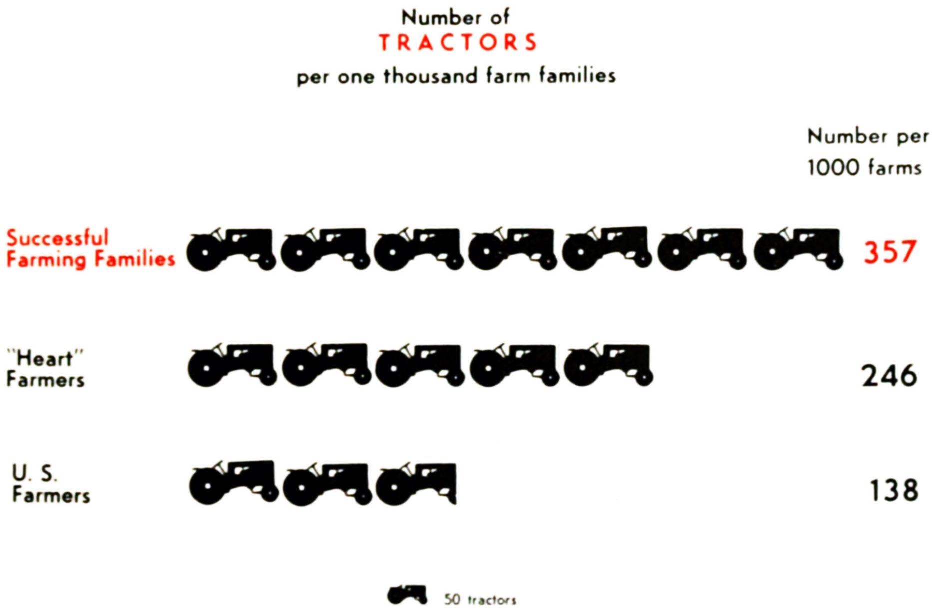

Figure 2.7 Example of stacked pictograms. “Successful Farming”, Meredith Publishing Co., Der Moines, Iowa, as is appears on (Brinton 1939, p.123).

3

Approach

In visualization, visual metaphors structure information and are suggested as a starting point for a general theory on the visual representation of information (Ziemkiewicz 2010; Cox 2006). Classical and recent techniques in visualization embody these metaphors by giving shape to information. They have been absorbed by information visualization practitioners in the sense that they are seen as conventional or literal, despite how novel they may be for the target audience. These techniques mostly deal with spatial arrangements of data, tailoring layouts and providing meaning to graphical implantations on a two-dimensional canvas. By providing a generalizable procedure on how to shape information in certain cases, these techniques can be framed as models. In this dissertation, the visualization models that embody metaphors seen in the field as conventional or literal are referred to as *structural metaphors*, since they configure information by manipulating spatial dimensions. These metaphors have varying degrees of *figurativeness*, depending on how novel they are.

But metaphors can be used for more than to structure information. Regardless of how neutral the role of the designer is within the framework provided by each structural metaphor, the designer can have an assumed role in tailoring a data-specific message by providing additional metaphors. Tailoring these additional metaphors in a way that, entangled with the structural metaphors, provide new authorial meanings to data, is an approach proposed here: *semantic figurative metaphors*.

Semantic figurative metaphors are on an exploratory side of visualization because their figurative aspects translate into novel representations. Such novelty, if not embodied by providing new metaphors, can be found in how the structural metaphor of a visualization is intertwined with new authorial cues. It is defended here that their natural applications are in the context of casual information visualization (Pousman et al. 2007). Rather than being used for strictly utilitarian purposes in the context of information visualization tools and analytics, their casual application targets a broad audience that includes a wide spectrum of users, from experts to novices. The casual use of visualization embraces momentary and contemplative usage patterns, which, aims at connecting with an audience on an emotional level and prioritizing the memorability of a visualization. Casual visualization is often found in contexts where they engage, entertain, provoke, or surprise.

Semantic figurative metaphors modify a visualization's structural model by adding specialized metaphors. These added metaphors are specifically designed not only for the structure of the data, but also for the content of the data interpreted by the visualization's author. Therefore, they substantiate the visualizer's own interpretation and message. In this sense they are semantic in the author's message. Their figurative nature consists of providing less abstract visual references to connect with the audience, and embraces novelty in their application contexts.

Implementing semantic figurative metaphors encompasses two steps: the adaption of the structural metaphor and the introduction of visual cues.

1. *Adaption of the structural metaphor*. After an adequate visualization model is chosen, this model is altered. The modification consists of evading the model's rules of spatial arrangement. Altering the model's rules is not gratuitous and each modification should have a specific reason given the data and the author's intent. These adaptions slightly restructure the information in a way that conveys new authorial messages.
2. *Introduction of visual cues*. Visual cues are added on top of the modified model. These cues aim at providing a figurative evidence of the author's message. The representation of the structure given to a certain piece of information, usually recurs to a minimalistic use of implantations. These implantations are often points with a circular shape, or lines with regular weight and a single color. Adding visual cues involves replacing such minimal implantations with more figurative representations that involve more complex shapes.

Semantic figurative metaphors intrinsically carry the weight of a subjective authorial intent. The authorial message is not necessarily or strictly present in the data itself. Nonetheless, in the context of this approach, these messages are regarded as a faithful interpretation of the data that emphasizes certain aspects of it in communication. Given their authorial intent, these messages can be partial towards the data. If this interpretation is inserted into the visualization, then it can be said that, strictly, non-data aspects are being added to it. Changing a structural metaphor or introducing visual cues with the purpose of giving new authorial meanings to data enables the insertion of inaccuracies in data representation. These inaccuracies should be employed to the advantage of implementing a dramatic depiction, supporting visual metaphors that convey the author's interpretation. Inaccuracies have the risk of subverting the data's underlying structure, and so they should be used only until a point where such distortions are negligible in the weight of the author's intent. This means that a delicate balance between enhancing certain authorial aspects extracted from data and inaccuracies in representation, has to be attained. Adding non-data aspects to a visualization can be regarded as mere embellishments. Nevertheless, it is defended here that this is not the case since these aspects have a semantic value in the author's message. They contradict the economy principle in information visualization, but depending on if they classify as data-ink or not, they can have a functional value within communication. This is not a permission or advocacy of chartjunk — employing semantic figurative metaphors should be handled soberly and with restraint. Visualization designers should make it as clear as possible what reflects the designer's own interpretation and what is the designer's message. The message, of course, should be clear and faithful to the data.

Each implementation of this approach should be tailored for each specific application. This is because an author has to create a strong message while attending requirements of the dataset and the target audience. Trying to make memorable visualizations is connected with novelty in several aspects of the visualization design, and novelty is dependent of case-by-case scenarios. For this, incorporating semantic figurative metaphors should not be seen as an add-on to the design process; they should be inherently part of the process since the beginning, including data acquisition and filtering. Figurative metaphors are tightly coupled with data and should conceptually shape how a certain visualization solution evolves. For example, the modifications on the structural model should be connected with the addition of visual cues, in light of the same metaphors. Moreover, the figurative aspect should not be extreme as it would distract from the familiar and abstract graphical languages that visualizers use to depict data – they should aim for clear, subtle and more figurative visual cues.

Semantic figurative metaphors are an approach that can be characterized by the following aspects.

— They are an exploratory approach to visualization, and not models. For this they should be tailored specifically for each application and context.
— They carry the weight of a subjective authorial message and thus are semantic in message.
— Their use is conceived in the context of casual visualization for broad audiences.
— They aim at humanizing information visualization by using more figurative depictions.
— They involve adapting a structural metaphor and introducing visual cues, based on non-data aspects.
— These modifications may involve the deliberate introduction of inaccuracies in data representation, which should be used to strengthen the author's perspective, but not to a point where they distort main data aspects.
— They take part in the design process and therefore they are not an add-on.

Within this approach, the main objective here is to show that:

The exploration of visual metaphors together with the distortion of visualization models, to the point of deliberately inserting inaccuracies in data representation, can be used to approximate the audience to the visualization while being clear, and faithful to the data.

This dissertation describes and analyzes three applications of semantic figurative metaphors. This approach is implemented by recurring to computational systems that introduce additional degrees of freedom in data representation. Instead of declaratively mapping data, these systems consist of physics-based behaviors that model data representation. Each of these systems is simulated and use their own runtime to express their behaviors based on the data. This is to say that certain elements in the system have their own rules on how they display the variation of a specific aspect of the data, through simulation time. These systems are able to interpolate through discrete data variations and hence animate data. With this it can be said that they adapt to data, as data is injected in the system, or as they query it themselves. Due to the nature of physical behaviors defined in these systems, they do not necessarily represent each quantity in data with total accuracy. These distortions in representation are used to the advantage of portraying additional visual metaphors, but only to an extent that does not impair faithfulness towards main data aspects.

4

Empires' decline and storytelling

During the 19th and 20th centuries, several imperialistic world powers lost their influence and territories. This chapter describes a visualization that shows the decline of four major empires: the British, the Spanish, the French, and the Portuguese. The visualization implements one approach to generative storytelling that is able to tell the same story with a slightly different expression.

Furthermore, the visualization implements specific semantic and figurative metaphors for disintegration, dissolution and competition. The design approach to the system is described both on a graphical and computational level. Implementing these metaphors comes at the expense of inaccuracies in the portrayal of information, which will be discussed throughout this chapter together with users' engagement.

4.1 Related work

The empires' decline visualization portrays a statistical variable for each graphical region along time. Displaying the power of empires was often done by highlighting their geographical regions on a map (see figure 4.1) or manipulating these geographical areas for political discourse, such as in figure 4.2.

More statistical variables can be shown using a scatter plot, positioning countries accordingly. Three variables, *x*, *y*, and the size of the points or circles, can be shown. Furthermore, an additional categorical variable can be shown through the circles' colors. This is what happens in the popular application Gapminder World[8], where x and y each represent a statistical variable, the circle's size represents the country's population and the circle's color represents the geographic region (see figure 4.3). Additionally, Gapminder World invests in narrative aspects, advancing time and animating the corresponding scatter plot in order to tell stories using geographic data.

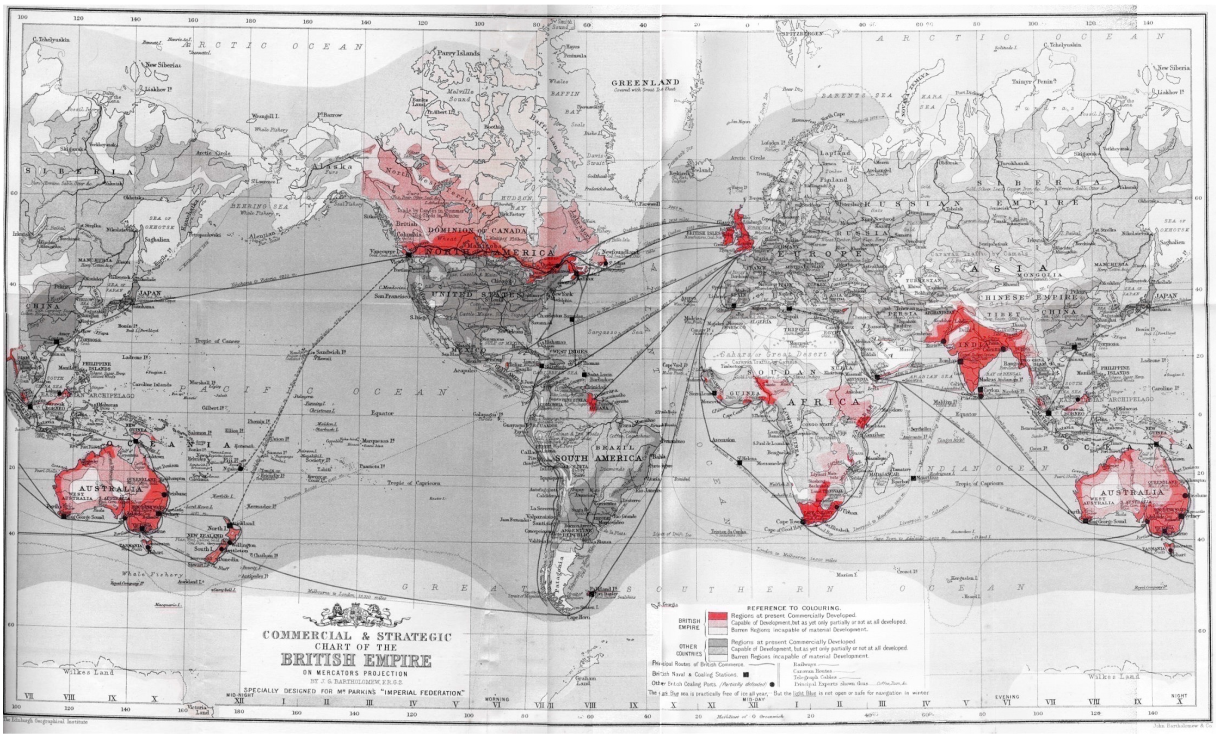

Figure 4.1 Commercial and Strategic Chart of The British Empire, by J.G. Bartholomew (Parkin 1892 315).

Sometimes when visualizing countries, their geographical component is abstracted and manipulated, or even completely omitted for layout purposes. For example, countries and several statistical variables can be portrayed using packed bubbles, abstracting their form and geographic position (see figure 4.4).

[8] http://www.gapminder.org

Figure 4.2 Portuguese overseas territories layout over Europe with the title "Portugal is not a small country" by Henrique Galvão[9]. Galvão, an active member of the regime during this period (and later on a dissident) said about the propagandistic inspiration for this map: "The men of my time were born in a small country. Fortunately they want to die in an empire."[10].

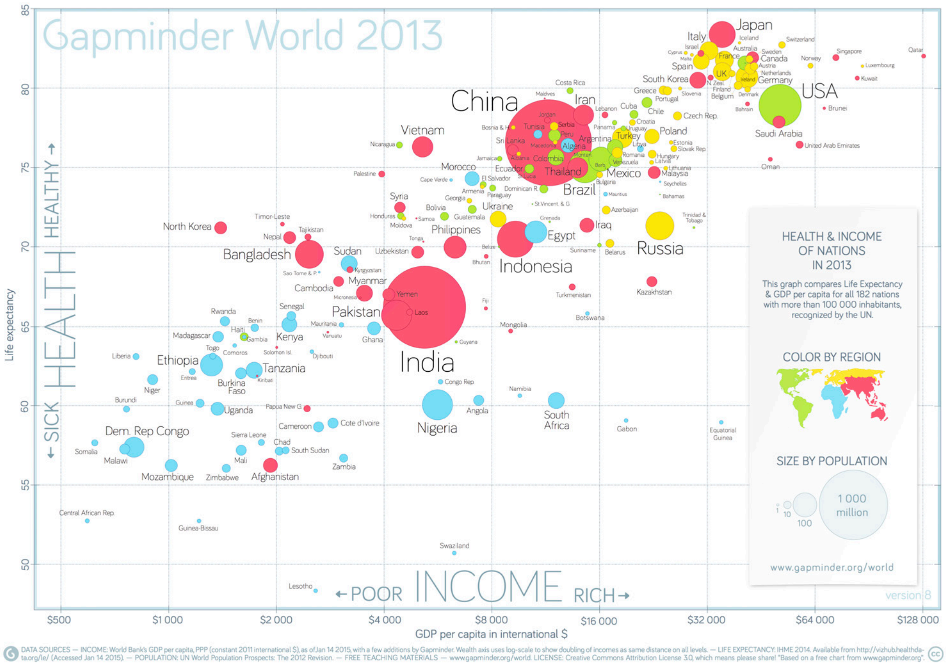

Figure 4.3 A Gapminder World of life expectancy vs GDP per capita for 2013, laying out the most and least developed nations.

---

9 "Portugal is not a small country" was a map featured in the Portuguese Colonial exhibition in 1934.

10 On the Portuguese Colonial exhibition, by Hemeroteca Municipal de Lisboa. http://hemerotecadigital.cm-lisboa.pt/efemerides/exposicaocolonial/exposicaocolonial.htm [Accessed December 8, 2015].

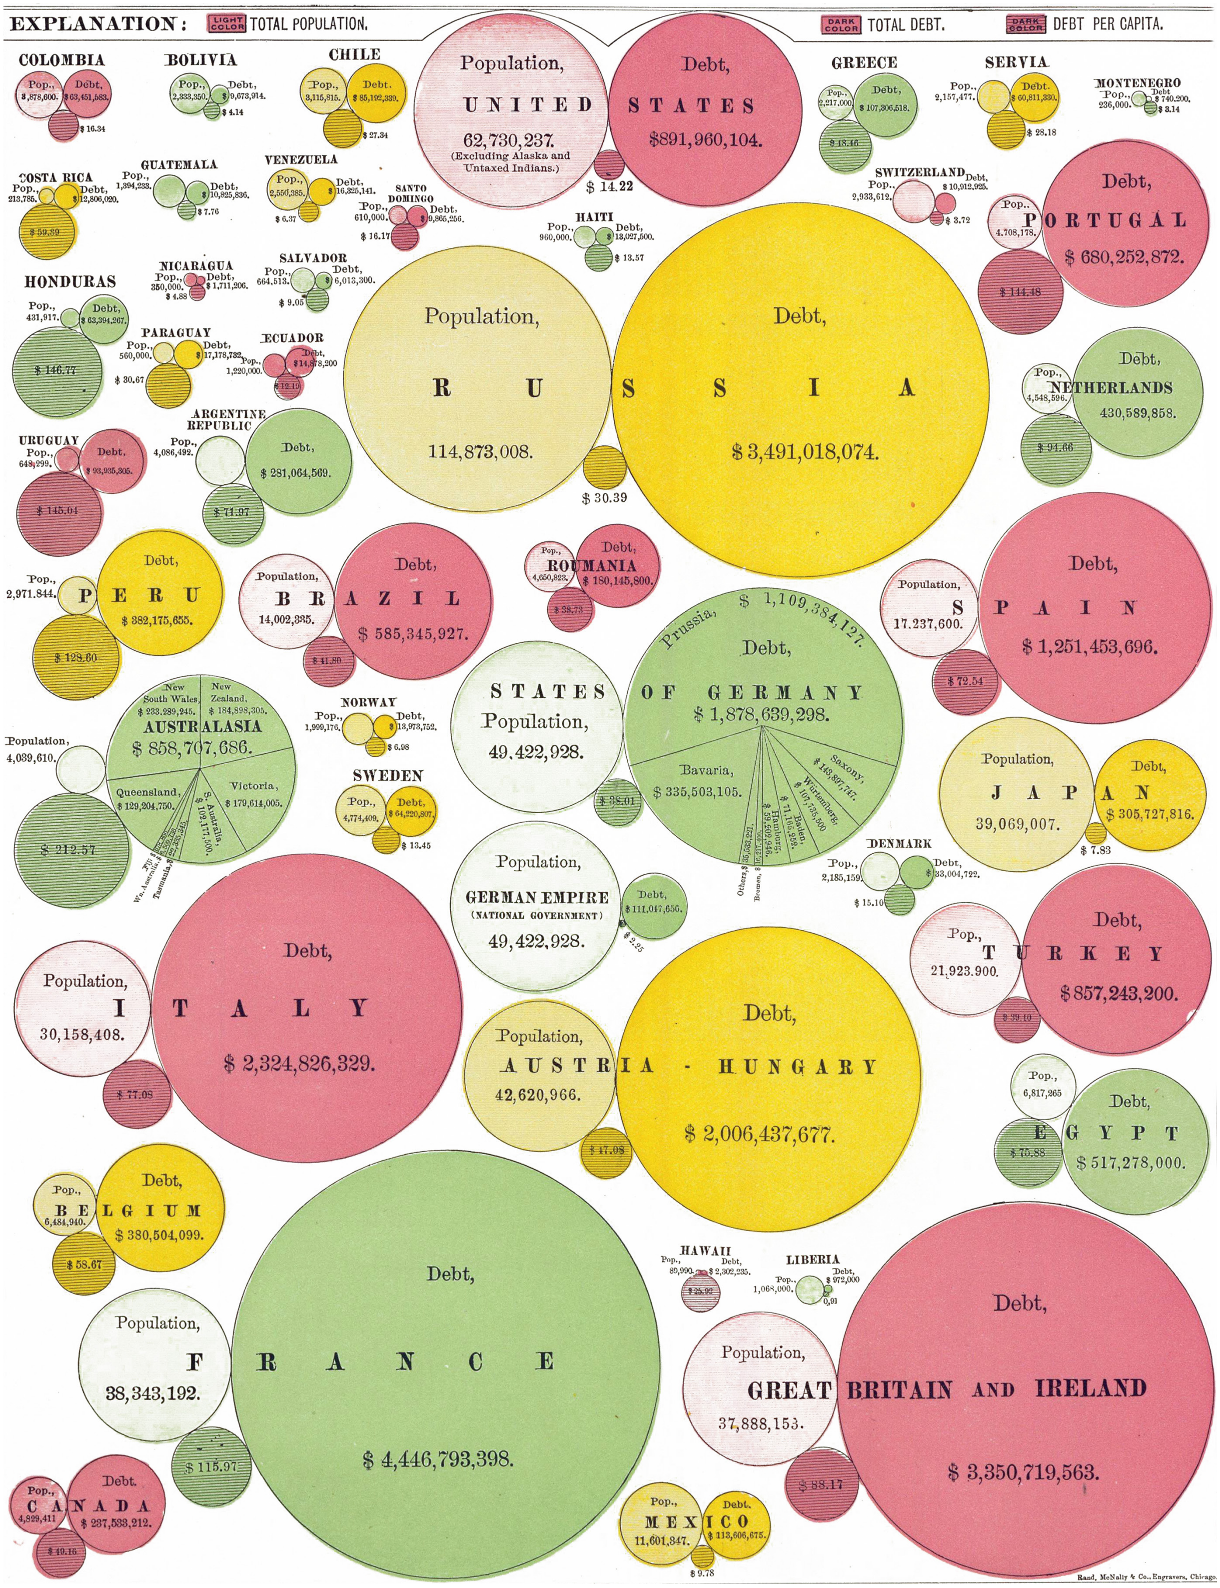

Figure 4.4 "National Debts of Foreign Countries compared with the United States, 1890". (McNally 1897, p.281) Light color: total population. Dark color: total debt. Darker color: debt per capita.

Packed-bubbles are a popular visualization model (Viegas et al. 2007), but they have been sparsely defined in the literature in a formal away. This model shows circles where their size represents an attribute of the corresponding entity, and where the color of the circles is often used to portray additional information. The circles are not positioned according with any data variable, but rather have non-data related initial positions that are then left to evolve according to the model's behavior (in the case of figure 3, the positioning of the circles is not derived from a computer model, but by the author). The model has physical properties: all the bubbles are usually attracted to the center of the canvas while not intersecting with each other. I.e., stronger or weaker collisions are often used as a resource to avoid these juxtapositions. Naturally packed-bubbles models can have additional behaviors, such as prioritizing the packing of smaller bubbles at the center, or the other way around. On moving and interacting media, they are often presented with their initial simulation stage where the circles converge to a state of equilibrium, resulting in a static artifact. Often the user is invited to interfere with this equilibrium state by resizing or moving the already packed circles. Both this possibility and the presentation of the initial simulation stage confer entertaining and ludic features to the model, since the users can observe and interact with a very simple, but self-organizing system. The nature of the model results often in aesthetically elegant solutions, since circles are packed in a bigger mass, as a way to minimize space. It also has organic tones: a bigger unitary whole made of smaller, but with different parts. Abstracting the geographical component in this type of visualization does not have to be absolute. Special examples of this are cartograms. Cartograms[11] deform shapes in maps in order to show a statistical variable. Usually they manipulate the areas of certain regions (e.g. states, countries) in order to show a demographic variable. Deforming the shapes can be done in such a way that the adjacencies of a map's regions are preserved, or they can result in non-contiguous cartograms where the adjacencies are partitioned. In these cases, the geographic position of each region is not rigorously preserved, but some properties can be, such as the general geographic arrangement and the relative positioning of each region in relation to others.

The work of Härö (1968) shows a technique for area cartograms based both on circles and rectangles. His solution abstracts the shapes of regions into more geometric renderings, while allowing the cartograms to be presented in a contiguous and non-contiguous fashion (see figure 4.5). Dorling simplified this approach by devising an algorithm that uses only circles to represent geographic regions (see figure 4.6). He named these maps "Circular Cartograms" (Dorling 1996, pp.32-36), but they have since been known as Dorling cartograms. The circles are arranged in a way that resemble the original map's topology, although they are moved from their initial locations so that they do not overlap. Doing this allows the full areas of the circles to be observed without obfuscations. A popular application of the Dorling cartograms is the world map that can be seen in figure 4.7.

[11] More can be read on cartograms in the next chapter.

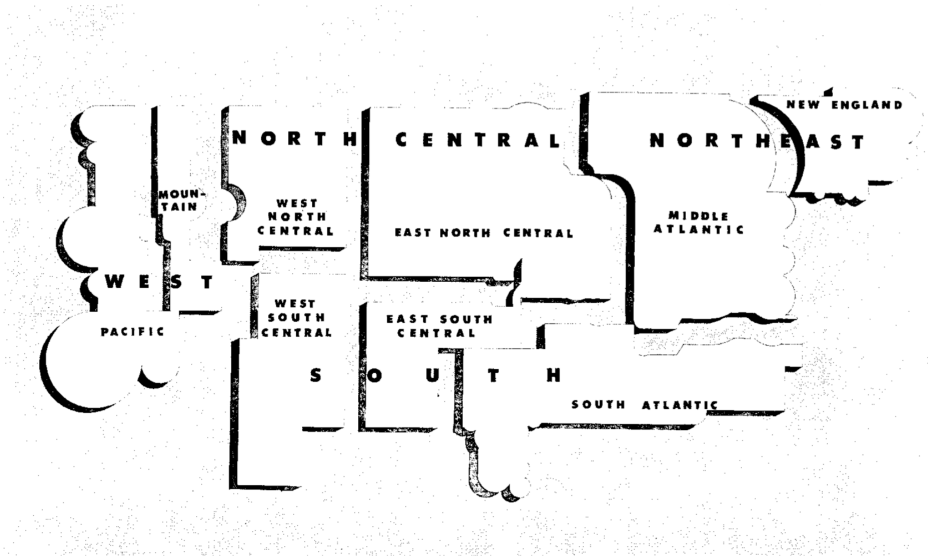

Figure 4.5 Statistical cartogram of population areas of the United States, derived from rectangles and circles (Härö 1968).

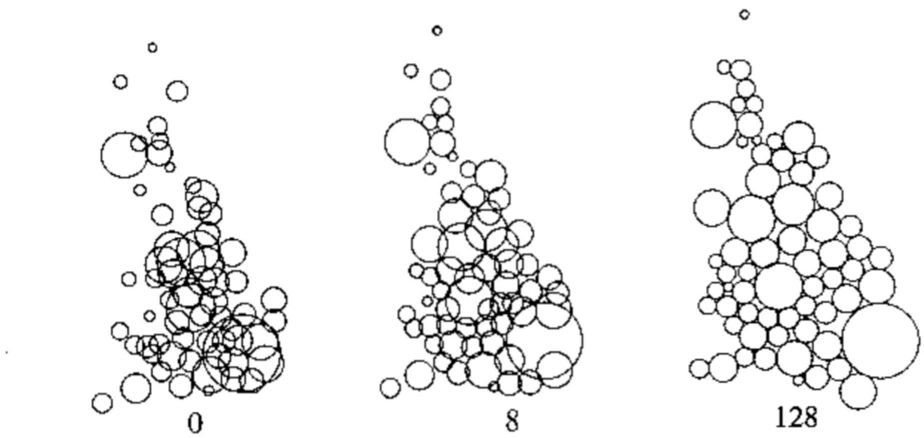

Figure 4.6 British counties scaled by population, as simulated by Dorling with his circular cartogram algorithm (0, 8 and 128 iterations) (Dorling 1996, p.60).

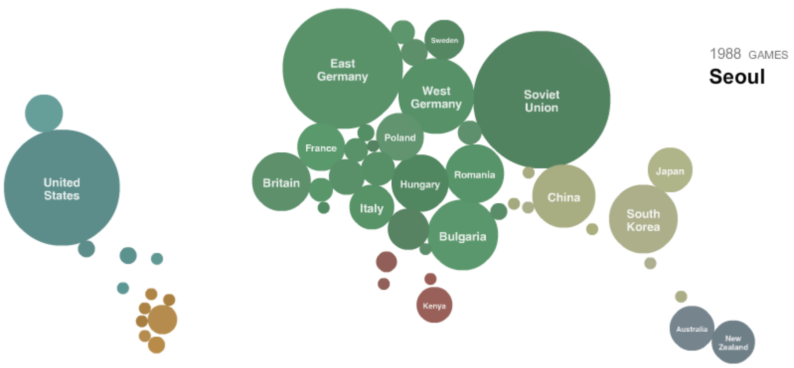

Figure 4.7 A Dorling cartogram of the Olympic medals won by each country in 1988[12]. Depending on their size, the countries can be pushed away from their geographical positions, but they tend to maintain their relative adjacencies.

4.2 Data

One distinguishing feature of the empires visualization is that data collection was part of the design process since its inception. The dataset was created in order to attain the narrative objectives of the visualization while maintaining historical accuracy.

The history of the possessions of each great empire across history is often portrayed through a sequence of static geographical maps highlighting their territorial possessions through time. Something that is often missed in these types of graphical depictions is how the evolution of an empire intertwines with the evolution of other major world players. This visualization has the main, personal, objective of portraying the apogee and decline of the Portuguese Empire and comparing it with other relevant empires. The Portuguese Empire was a modern Western overseas empire, and within this classification there were only four empires that occupied more land area than the Portuguese: the British, the Spanish and the French[13]. These were the major players in this context, sharing cooperation and disputes. Framing these empires in constant confrontation translates into a graphical depiction that shows them side-by-side.

Choosing land area to portray the influence and power of each empire is tied with the typical imperialistic maps that also portray empires in this way. Depicting land areas as visual areas in a visualization follows Tufte's principles (Tufte 1983, p.71), as he advises not to map one-dimensional data variables to two-dimensional visual variables.

The apogees and dramatic decline of each of these empires happened during the 19th and 20th centuries, making these centuries a rich source of anecdotes that brings the extinction of these empires to contemporaneity[14]. For this period, two classes of historical events were collected[15] for the four empires: growth events that involved the acquisition of new land, and disintegration events that were related to the self-determination, an often de facto way of new territories gaining independence. For the sake of simplicity these disintegration events will be called independence events. A growth event is related with an empire and is described by a year and the new area of the empire (see graph 4.1). An independence event also affects an empire, which and has the name of the new independent nation, the year, and the land area of this new nation (see graph 4.2). Determining the areas of new independent nations often involved researching by areas of territories that are no longer present in the current geo-political map. This information was formatted into XML, and was made available in English[16] and Portuguese[17]. In total, 21 and 96, growth, and independence events respectively were included.

[12] Byron, L., Cox, A. and Ericson, M., 2008. A Map of Olympic Medals. The New York Times. Available at: http://www.nytimes.com/interactive/2008/08/04/sports/olympics/20080804_medalcount_map.html [Accessed December 9, 2015].

Graph 4.1 Growth events in the data, characterized by year and area in millions of square kilometers. The British in red, French in blue, Spanish in yellow and Portuguese in green. The first growth events in data do not constitute historical events and are used to initially setup the areas of each empire. As can be observed, the Portuguese and Spanish do not have relevant historically-related growth events during this period.

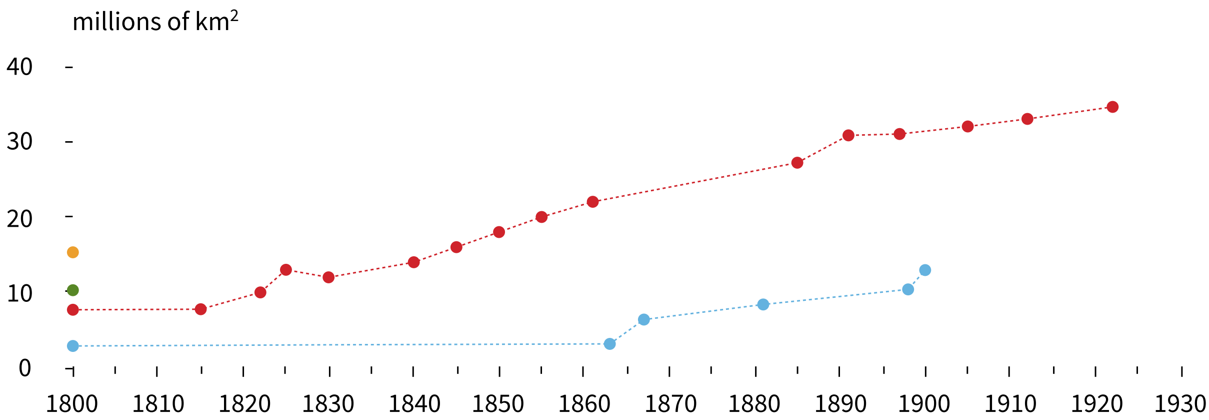

Graph 4.2 Independence events of several territories in the data by year. The land area of each new independent nation is on a logarithmic scale, in millions of square kilometers. Former British territories are in red, French in blue, Spanish in yellow, and Portuguese in green.

As can be seen in graph 4.2, the data starts by presenting two independence events from the French: Louisiana and Haiti. Several independences from the Spanish then take place and one in isolation from the Portuguese (the Empire of Brazil). The following years until the 20s of the 20th century have low oc-

[13] Another Empire that follows these same traits was the Dutch Empire. The land possessions of the Dutch Empire are sometimes hard to quantify, having had many scattered outposts and controlled lands for often fleeting and inconstant periods of time. Nonetheless, for what can be accounted, the Dutch had a smaller maximum land extent than the Portuguese and was left out of the visualization, maintaining the principle of only showing the Western overseas empires with larger area extents than the Portuguese.

[14] During this period, formally after 1830, the French Empire is referred to the Second French Colonial Empire.

[15] The research was mainly conducted by resorting to Wikipedia.

[16] http://pmcruz.com/empires_data/empires_data.xml

[17] http://pmcruz.com/empires_data/empires_data_pt.xml

currence of events, accounting for only four independences from Spain. After 1920, the British start to disintegrate, followed by the French and lastly, the Portuguese. Attention should be made to 1960 were a high number of independences from the French Empire takes place – mainly secessions of their African colonies. For the British, there are independence events only after a sequence of growth events.

The actual land area of each empire at any point in time in this period is not directly described in data but has to be derived from it. For example, in graph 1, the two first growth events of the French empire in 1800 and 1830 do not mean that the empire's area was constant during this period. As can be observed from graph 4.2, two independence events from the French occurred in 1803 and 1804, largely diminishing the area of the empire. What happened in 1830 was the invasion of Algeria by the French, bringing the empire's area to about the same size as 1800. Therefore, the two types of events have to be combined in order to extract the areas of each empire over time, as can be seen in graph 4.3.

Graph 4.3 The evolution of areas of each empire on a logarithmic scale in millions of square kilometers. British in red, French in blue, Spanish in yellow and Portuguese in green. The growth and decline of the British Empire can be observed in the graph, together with the rise and fall of the Second French colonial Empire. Only the Spanish and Portuguese's declines are shown at this period in time.

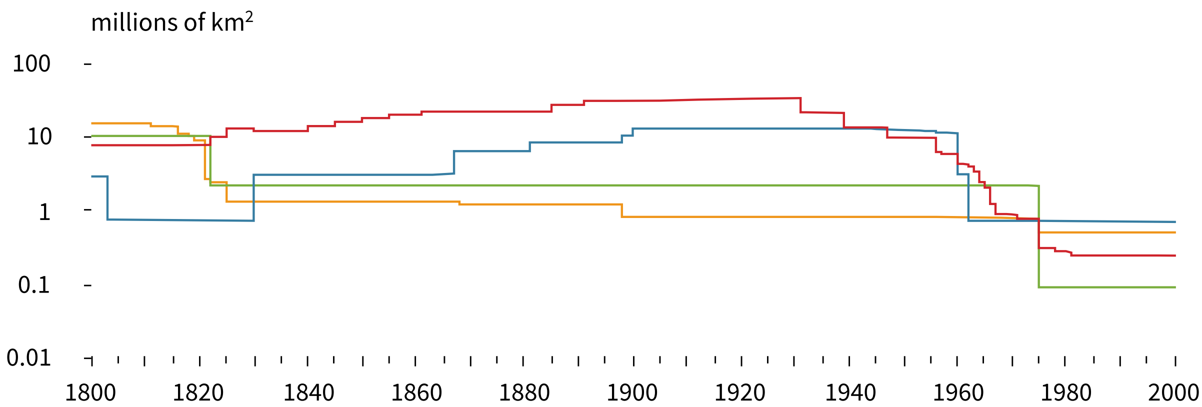

Determining the date of each territory's independence is not a univocal process. Several nations passed through complex processes that led to their independences. Independence is usually first recognized by some nations, and then only on a later date by others and the former empire. For example, in Cuba, independence was proclaimed in 1868 which initiated the Ten Years' War with Spain. Spain kept control of the island afterwards. After the Spanish-American War, Cuba was handed to the U.S. in 1898, but only gained its formal independence in 1902 as a de facto protectorate of the U.S. Sequences of events like these give space to pick several dates for the independence of a territory from its former empire. For a matter of simplicity, a common criterion was adopted in order to choose the date for each case. This criterion picked the date of the *first declarations of independence* or the *beginning of a war or rebellion that resulted in a de facto independence*. Therefore, the year chosen for Cuba's first step towards self-government was 1868. Another intricate example deals with the independence of the Portuguese African colonies in the 1970s.

Even although the Portuguese Colonial War had been fought since 1961, it did not culminated in the independence of the African territories, but rather a military coup in Portugal in 1974 ended this war. Therefore, territories such as Angola or Mozambique that did not have declarations of independence until that date only separated in 1975 when Portugal abdicated its sovereignty over these territories. On the other hand, Guinea-Bissau declared its independence unilaterally in 1973 and thus has its split in that year. In these cases, the split-date tries to correspond with the urges of a territory to become independent.

A remarkable exception to this independence criterion is the case of British Dominions, such as Australia and Canada. Even though for several years the Dominions had the benefit of de facto independence, they were seen as part of the British empire as illustrated in figure . For British Dominions, their split was only attributed to each one's ratification of the 1931 Statute of Westminster.

These criteria seek uniformity in data gathering, but they also constitute an editorial choice of the visualization designer, prioritizing urges for a territory's self-determination above the universal recognition of other established countries.

Figure 4.8 Illustration on how British Dominions were promoted as part of the British Empire in 1920 (Byrde 1920, p.1).

4.3 Generative storytelling

The data described in the previous section is a set of events on the growths and independences of four overseas empires. The objective of this work is to portray the dissolution and disintegration of these empires along time. This is to say that the visualization of this data will tell a story. Several stories with different expressions can be built from the same data. Choosing a story is an authorial exercise of the visualization designer, and in this context it is important to discuss how data can be translated from a set of events to a story. The variability of stories that can be conveyed from the same data impels the description of an approach on *generative storytelling*: a *conceptual framework* that deals with translating a dataset into the same story with varying degrees of expressiveness. This section dedicates to how storytelling for data can be seen in light of this conceptual framework. Telling a story from data consists of interpreting events, actions, and actors and amplifies the cognition of information. A story uses real-world metaphors and analogies to convey knowledge in a ludic, compact way. Storytelling is play a major role in information visualization. As Gerson and Ward (2001) point "the difference between visualizations and traditional entertainment media is the information and story conveyed in information visualization environments are usually much more complicated than those typically shown in films or the theater or on television programs and commercials." In this context, storytelling deals with the core of information visualization, structuring relevant knowledge and presenting it efficiently, coherently and economically (Gershon and Page 2001).

4.3.1 Data fabulas

Time is frequently one of the n-dimensions of a dataset, for example, historical events, financial data, biological behaviors, and model descriptions in physics. Given this temporal nature, building a story based on these datasets can be a strong device to leverage their communicative power. Following Mieke Bal (Bal 1997, p.5), a fabula is as a set of time-ordered events caused or experienced by actors. In this context, actors are agents that perform actions in the fabula's time span. This concept can be translated to a time-based dataset, which here is called *data fabula*: the set of events, agents, actions, and chronology extracted from a dataset. From now on, in this writing, "fabula" and "data fabula" are used interchangeably. For simplicity, only one fabula can be extracted from a time-based dataset, which implies that the set of events from data is immutable.

A fabula can be presented in different ways. A story is the fabula's presentation layer: it has its own structure, built from the fabula's events, actors, and actions. This structure also defines the emphasis given to the events, actors, and actions. The specific structure of a story is the narrative. The narrative's chronology frequently varies from that of the fabula in its order, pace, or rhythm. The agent that transmutes the fabula into a story is the narrator or storyteller. This transmutation is called storytelling. The narrator communicates the story through some narrative medium. In the empires' visualization, the medium is a digital short film, and the narrator can be considered for its most part as the author of the visualization, although the computational nature of the algorithm that generates the films has a degree of intervention in the expressiveness of the final story.

4.3.2 Generative storytelling and data

A time-based dataset is a collection of temporal measurements that are considered to be inherently discrete. Usually, the extraction of a fabula from a dataset results from mining the data. Clustering and aggregating the data play a central role in establishing a chronology, a set of events, and a set of actions, and in identifying the relevant actors. An event is a time-based trigger of a set of actions undertaken by the actors. The actors are visual agents that represent several data properties over time. Choosing these authors relies on semantic properties of the data as well as the personal judgment of the designer of the visualization. The actions are the variations of the actors' properties that show the data, and can be casuistically associated with time-stamps in the dataset (i.e. events).

Because a story is a fabula's presentation layer, it involves:

— the representation of the fabula's actors;
— the definition of a temporal structure that, while containing the same events as the fabula, establishes a new pace, rhythm, or chronology, therefore creating a narrative.

These characteristics form the story's identity. Generative storytelling builds the story by implementing the representation of the actors and the new structure of events. The generative nature of the actors' representation emerges from a set of rules that interprets the fabula's actions.

Although the fabula like the dataset is discrete, it can be portrayed in a perceptually continuous fashion. Continuity is one of the compelling characteristics of a story and therefore it is of interest to transform a discrete fabula into a story with a continuous timeline. This requires the portrayal of time intervals between events and gives room to the continuous temporal interpolation of the actors' representations. These interpolations or transitions over time are provided by the interpretation of actors' actions. The actions' interpretation might also lead to variations of the actors' behaviors derived from the ones the fabula imposes. This can be attained, for instance, by allowing a non-rigorous representation of attributes, incorporating behaviors that distort the fabula and that introduce collateral actions, or create new interactions among actors. Such secondary actions exist only in the story – they are an ornamentation resource that promotes the story's dramatic value without changing its semantic content. Therefore, the story does not have to be an accurate representation of the dataset, and the mapping of the fabula's actions to the story is not necessarily linear, fixed, or even deterministic. This is one of generative storytelling's most distinctive and prominent features, constituting an important narrative resource.

This approach to generative storytelling can be implemented through an engine that transforms a fabula into a story. This engine has two parts: an event model and an action model. The event model manipulates the time of the events, creating the story's timeline. It processes the fabula events and triggers the respective actions. Although the story is delivered in a continuous fashion, such continuity does not necessarily translate into linearity in the events' presentation. With this, it is possible to employ temporal manipulations of the fabula in order to achieve various degrees of expressiveness and confer dramatic tones if necessary. These manipulations not only can change the pace and rhythm but also can alter the chronology itself by introducing narrative devices such as *analepses, prolepses, ellipses, in medias res,* and *reversions*.

The action model deals with the representation of the fabula's actions, implementing a set of actors' behaviors. This model can be generalized as an adaptive system. This means that as the actions are triggered by events through time, the system tends to adapt to those actions rather than reacting directly to them. Therefore a new action can be seen as time-based disturbances to the system. This results in a set of important parameters of the model. One such parameter is the system's adaptation threshold, which determines the actions that are relevant. Another is the system's fidelity to the fabula; that is, the adaptive system might promote a loose or tight representation of the fabula. A third is the system's convergence rate to new equilibrium states; that is, the system might react rapidly or slowly to new actions.

4.3.3 Room for expressiveness

Interpreting data into a story can culminate in varying degrees of expressiveness by using the manipulations of the fabula as narrative resources. A different expressive intent can thus generate a different story. The conceptual framework of generative storytelling embraces expressiveness – with the same implementation, one can achieve different expressions by adjusting the model parameters. Furthermore, varying expressions can also result from the same parameters, using a non-deterministic model. It is desirable that this degree of non-determinism preserves the direct interpretation of the actions into a set of strong and evident behaviors in the representation, while enabling the playout of secondary behaviors of the actors that, although not directly present in data, provide a unique expressive depth to each story as a complementary devices of the main metaphors.

As previously stated, the resulting story does not necessarily represents the dataset with precision. However, the liberty taken in the generative approach can result in meaningful variations, adding a dramatic tone, permitting authorial aesthetic manipulations, and conferring a ludic nature to the resulting visualization. It is argued that this expressive playground can be used to promote cognitive amplification in the context of information visualization. Ultimately, generative storytelling should delicately balance the faithful portrayal of data and expressiveness. This is not in detriment of accuracy in data portrayal or corruption of the fabula – the manipulation of the parameters of a generative storytelling model should enhance the extraction of relevant information without distorting or occluding it.

4.4 Visualization design

The Empires' Decline visualization shows the story of the dissolution of the four major overseas empires, making use of specific figurative metaphors. There are two versions of the system generating the visualizations, which invest in slightly different messages and metaphors. The way these two versions implement these metaphors from a computational and design perspective are described in this section.

4.4.1 Message

The first implementation of the Empires' Decline visualization invests in displaying the generic struggle of independences as well as the sense of aggressive competition among empires. Therefore, the visualization shows more aggressive interactions among the bodies. The second implementation, a refined version, invests in an idea of dissolution and disintegration of the empires with more subtle, paced and detailed interactions. It still preserves the message of competition among empires as they collide among themselves, but it drastically changes the tone of an aggressive repulsion of a new nation to a fluid and tender split with a birth-like feel.

4.4.2 Implementation

The structural model for this first implementation is one of animated and packed bubbles (Viegas et al. 2007) over time. The empires visualization modifies the common properties of the packed bubbles model. For a start, circles are not perfectly circular and their shape is modified amidst collisions with other circles. Not every circle in the canvas is attracted to the center, and other additional behaviors that modify this packed bubbles model are implemented in order to carry specific figurative metaphors for this visualization.

The actors in the story are the empires, and secondary actors are the newborn countries. As time passes, two types of event occur: growth events that make the empires expand; and independence events that create newborn nations in the canvas and make the corresponding empire shrink. The independence events trigger independence actions that are described further on. The areas of the empires and the newborn countries directly signify the land area at the corresponding time.

The system simulates the actors in the canvas as soft circular bodies. This means that when these bodies move or collide with each other, their shape can distort. The soft bodies have an elastic nature since they tend to quickly regain their circular shape in the absence of external forces. The soft body behavior was chosen to confer a different expressiveness to the visualization, as well as to use this expressiveness to explore specific metaphors. For example, bodies collide aggressively with each other, and when a new country is born, it is violently projected from its former empire, staying in the canvas for a few seconds and then fading away. The disappearance of newborn nations was chosen as a way to economize space in the visualization. These abrupt and strong thrusts signify the generic struggle for each independence. This dynamic behavior of the bodies is graphically emphasized because of the soft-bodies' properties, having their shapes squished and perturbed amidst interaction and thus conveying a more dramatic tone.

Each soft body was implemented by building a skeleton of springs. Each body is formed by a central particle of unitary mass and a number of particles equally spaced over a circumference centered on that particle. These particles are connected through springs. The circumference particles are connected in order to form the body's rim, and each of the circumference particles is also connected by a spring to the central particle. It was observed that during the first simulation this skeleton was too fragile – its structure could be irreversibly compromised when external forces were exerted on its particles. This specific skeleton could collapse in on itself, radically changing the particles' relative positions among themselves and thus resulting in random shapes with unstable and irregular areas. In order to tackle this, a more robust skeleton was implemented on top of the radial springs that connect to the central particle, adding diametral springs that directly connect each pair of opposed particles in the rim (see figure 4.9). The skeletons for the empires and the skeletons for the new countries are created in a slightly different way. The empires have a higher particle density in their rims and thus possess more robust skeletons. This makes the shape of the empire less sensible to external forces than the newborn nations that can have their shape highly deformed when external forces are exerted. The implementation of the springs and particles simulation is provided by toxiclibs[18] for the Processing[19] language.

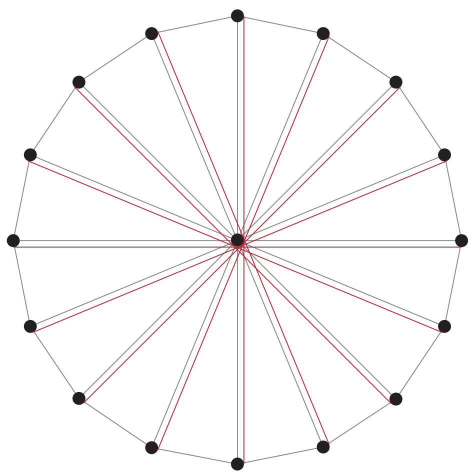

Figure 4.9 The skeleton of a soft body in the empires visualization. Each body is composed of a certain number of particles, circularly distributed in turn of a central particle. Springs connect the circumference particles in order to form the rim of the body and also connect each particle of the circumference to the central particle. Additionally, in order to enhance the skeleton's stability, diametral springs were added (in red) that directly connect opposed particles in the circumference.

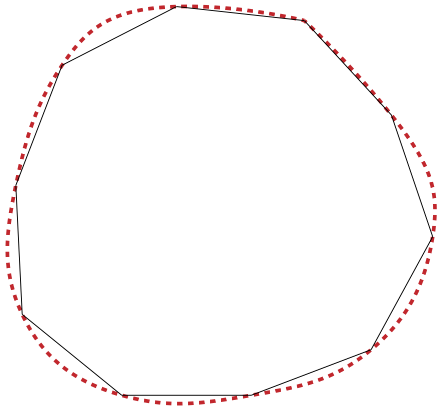

Figure 4.10 The organic shapes of the bodies are rendered using B-Splines that interpolate the vertices on the body's rim.

In order to confer organic tones to the visualization and render the appearance of the simulated skeletons as soft bodies, splines are used to interpolate the vertices of the body's rim[20] (see figure 4.10).

When external forces are exerted, the perturbations on the skeletons that are rendered through splines can result in various shapes, some radically different from the initial circle (see figure 4.11). Nonetheless, the devised skeletons are robust, and the body recovers its circular shape after forces are not acting upon it.

During simulation, these bodies will have to shrink and grow as the area they represent changes across time. Resizing a body simply involves altering its springs' rest lengths in accordance to the new area that is to be represented, based on a perfect circle. Therefore, both the diametral and radial springs are resized based on the radius of such circle. The springs that connect the particles on the rim have their rest lengths adjusted as well by computing the new circumference length and dividing it by the number of rim particles.

Springs are also used to implement the forces that act in the simulation world in order to present a behavior that includes collisions and repulsions. For this purpose, each pair of central particles of every body in the system is permanently interconnected with springs that keep a minimum rest length (see figure 4.12). This minimum rest length is the sum of the radius of the bodies that it interconnects, plus a certain threshold. This way, even new bodies (new countries) that are created on the rim of its former empire will be automatically propelled away from it. The number of connecting springs that implement repulsions grows exponentially with the number of bodies in the system (see figure 4.13). This can ultimately result in performance problems giv-

[18] The Verlet implementation was used as provided by the library and it is usually a good tradeoff between stability and performance. http://toxiclibs.org

[19] https://processing.org

[20] The used implementation is also provided by toxiclibs that specifically implements B-Splines. More can be read on B-Splines in Knott (2000, p.151)

en the high count of springs for certain simulation periods. Nonetheless performance was not an issue, provided that the number of particles that made the bodies was not high[21]. The springs are stiff enough to enable collision and repulsion behavior, as well as maintaining the bodies' structure without irreversibly compromising them during simulation. Nonetheless, such stiffness still leaves room for the bodies to alter their shapes during collisions.

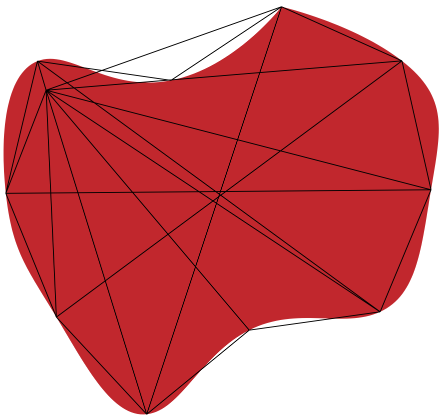

Figure 4.11 A body distorted by the exertion of external forces. The type of skeleton adopted enables this type of more organic shapes while guaranteeing that they recover their circular shape when the external tensions are minimal.

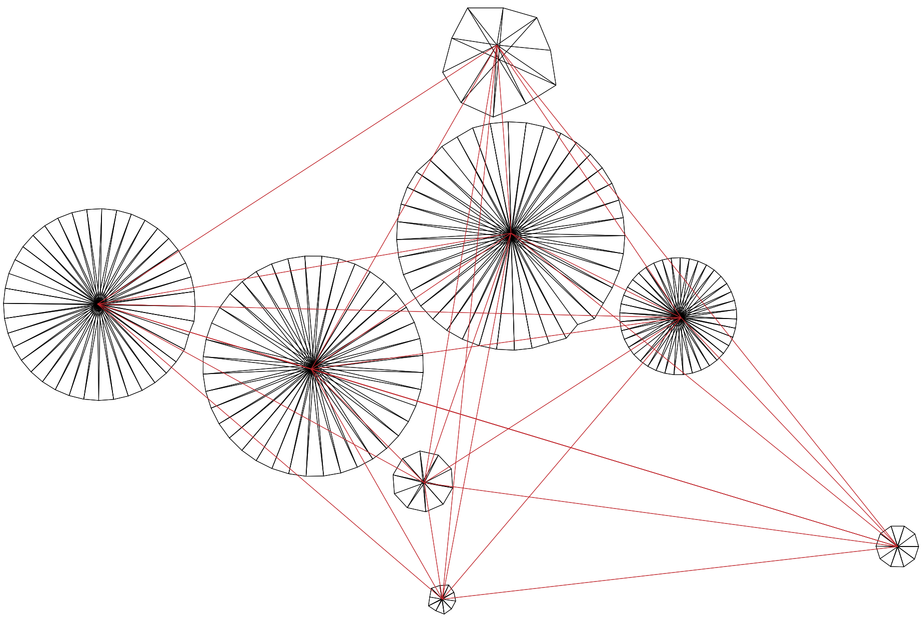

Figure 4.12 Springs (in red) connect each body to every other body in the system. These springs have a minimum rest length that enables the implementation of repulsion among bodies in the system.

[21] For n bodies in the system, the number of connecting springs for the collisions is $n \times (n-1) \times 0.5$. The simulation was able to run at a minimum of 30 fps on a 1.7 GHz Intel Core i5.

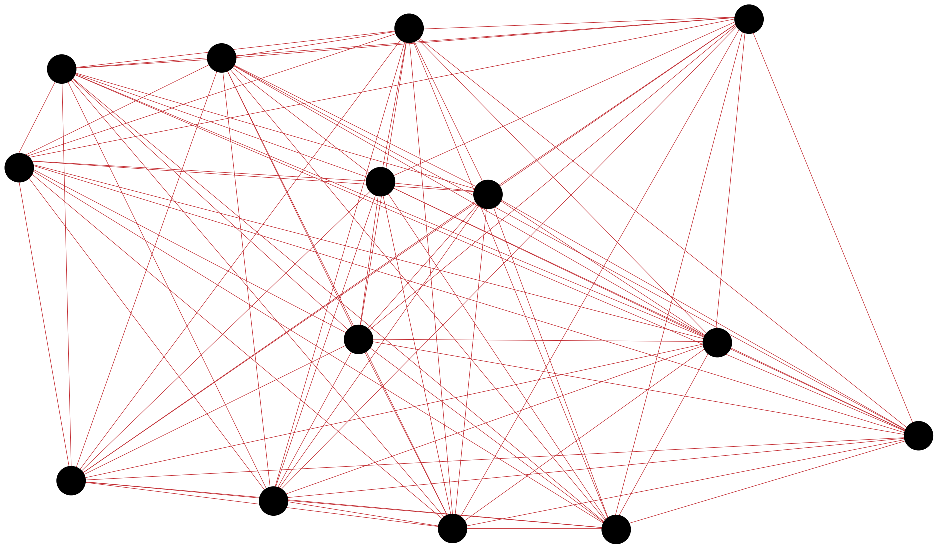

Figure 4.13 The number of springs that connects one body to any other body in the system grows exponentially with the number of bodies.

The time in the simulation is incremented linearly, triggering disintegration events. One second in simulation time is constituted by 30 simulated frames, with each having 50 iteration steps on the physics simulation. The film, starting in 1810 and finishing in 2010 has the final duration of three minutes and 30 seconds. The parameters of the simulation were adjusted in order to exhibit the behaviors and the metaphors described as follows.

— The four main empires are continuously attracted to the middle of the canvas, always colliding against each other, alluding to a sense of competition.
— Six years prior to the independence of a territory, a new circle starts growing in a random position over the rim of the empire. While growing, this new circle has no body-like properties and grows until it reaches the area of the corresponding new country. This visual anticipation of the independence translates the tension of the forthcoming separation, thus working as a narrative prolepsis and adding to the dramatic effect of the event.
— When the year of independence is reached, the circle previously mentioned acquires its body-like behavior and thus is immediately repulsed from its former empire. Therefore, the new nation immediately is projected as if "popped" from the former empire. This behavior together with the growth of the new nation has an organic inspiration from the mitosis-like splits of cells.
— At the moment of the split, the empire shrinks to its new extension. The newborn nation carries a label with its name and the color of the empire in a darker tone (this grow-pop-shrink behavior of a new nation's split is illustrated in figure 4.14).
— The new nation, contradicting the general behavior of packed bubbles, is not attracted to any position, being able to freely collide with the other bodies in the system. The new nation exhibits this behavior for ap-

proximately seven years in simulation time. After this, it fades away, giving room for other new nations to fill in the canvas.

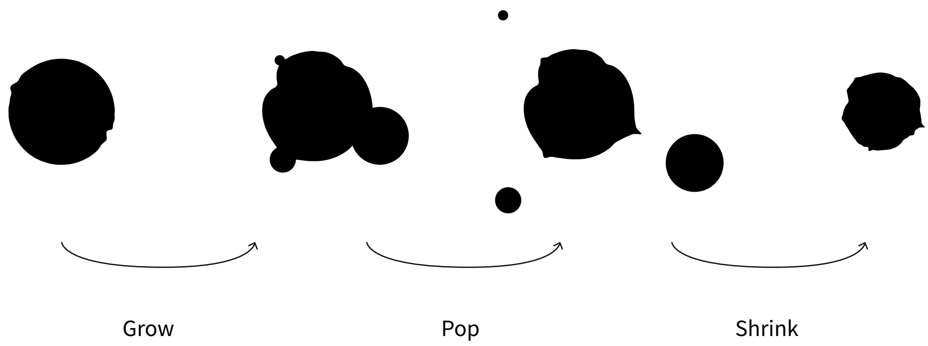

Figure 4.14 Previous to a disintegration event, a new circle starts growing on the empire's rim. When the time-event comes, the nation gain their independence, acquire body-like properties, and the corresponding empire is resized to meet its new area.

4.4.3 Color

Bertin notes that color is an excellent visual variable for the selective perceptual task, meaning that it can be used to instantly group several elements and isolate them from others (Bertin 2010, pp.67-91). As Ware also notes "using color to display data categories is usually the best choice" (Ware 2004, p.116). Nonetheless, this selective property is effective when comparing perceptually distinct colors, but not when comparing brighter and darker colors, in which case the ordering perceptual task takes precedence (Bertin 2010, p.87). In addition, Ware notes that there are only eight colors that are consistently named, and thus the space to use color for the selective perceptual task is very limited but enough to identify the four major empires at play. The color attribution had to maximize color variation in a way that avoids any type of hierarchy, but at the same time has a symbolic relation with each empire.

This color attribution is not a univocal process, but is much less arbitrary. The objective is to choose an exclusive solid color for each empire. The national colors of each former imperialist nation were based on their contemporary flags, since it is a visual reference more present in public awareness than historical flags (see figure 4.15). As a way to attain this, precedence was given to bigger empires as the colors were attributed.

The Union Jack of the British Empire has three colors. Red was chosen for the British Empire not only because it is arguably the strongest color, but because British historical maps frequently used red to highlight their possessions. Furthermore, the British East India Company extensively used red in its flag, and naturally in the context of imperialism primacy should be given to England's St. George's Cross.

If the British Empire is red, this leaves out the weld-yellow for the Spanish. Although only present since the 19th century in their flag, it has been extensively used in royal standards and its origins can be traced back to the Catholic Monarchs, namely the Crowns of Aragon and Castille. Furthermore it is a color frequently found in the flags of former Spanish colonies such as Colombia, Venezuela, Ecuador and Bolivia.

For the French Empire, the blue color was chosen as the most adequate color among the three in the current flag. In fact, blue was traditionally used in pre-revolutionary coats of arms of the Kingdom of France and it is still used today by overseas people of French heritage such as Acadians, Québécois or Cajuns.

Blue has been as well traditionally used in various Portuguese flags. The first Portuguese flag was a blue cross on white, and successive flags preserved blue escutcheons in the coat of arms. In the 19th century blue gained prominence in the bicolor blue-white Portuguese flag. The current Republican flag that appeared in the beginning of the 20th century is a bicolor green-red flag. Given the Portuguese insistence of maintaining its overseas provinces until the late seventies, together other colors being already attributed to other empires, green was chosen for the Portuguese Empire. It should be noticed as well, that before the Republican flag, at the time of Brazilian independence, the new nation took green as a preponderant national color, yielding it until today and contributing to make the green a strong choice in the sphere of Portuguese influence.

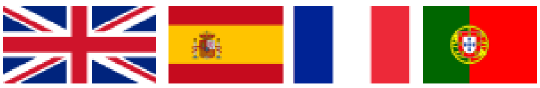

Figure 4.15 Contemporary flags of the British, the Spanish, the French and the Portuguese.

White is present in two flags, but was avoided, as it would be too light for the visualization. The chosen colors, as derived from the flags, are close to pure-hue colors and thus can be perceptually very light. Because of this, black was chosen for the background canvas in order to tone down the overall brightness while enhancing the contrast of the agents in the visualization. Nonetheless, the derived colors had to be modified in order to equalize their perceptual differences. Pure hues were not used because they can appear brighter or darker among themselves (e.g. a pure yellow will appear brighter than a pure blue) (Bertin 2010, p.85). The red on the Union Jack could not be used as it would appear much brighter than the other empires, especially against a black background and thus was toned down. The same goes for the Spanish weld-yellow as it was transformed into orange. On the other hand, the blue on the French flag would have appeared much darker than the other colors and was shifted to a cyan, while the Portuguese forest-green had to be brightened. This subjective balance was pursued in order to confer the same tone to all colors, avoiding having some more strident than others. As it can be noted in figure C, the colors chosen are perceptually very distant, and can hardly be arranged to establish any sense of hierarchy without an explicit color scale that forces this. Furthermore, in this first implementation of the empires visualization, the same color hue is used for an empire and its respective newborn countries. Nonetheless, the color attributed to the newborn countries is a darker version of the corresponding empire, establishing a noticeable visual hierarchy between empires and new nations (see figure 4.16).

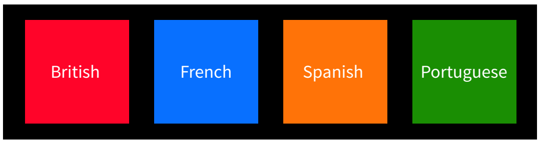

Figure 4.16 Colors chosen to represent each of the four empires against a black background.

4.4.4 Results

The simulation results in short films of three minutes and 30 seconds and can be watched here: https://vimeo.com/6437816. The story shows each major disintegration event from each of the four major overseas empires, while graphically displaying the area of each new territory and the evolution of the landmass of each empire. With this, it is possible to compare their territorial evolution along time and graphically emphasize the apogees of the empires. Furthermore, the cultural interpretations and personal knowledge of the audience, enables each portrayed event, or a set of events, to be seen as a visual anecdote (Offenhuber 2010). Visual anecdotes in visualization are small narratives tied to a data point. Several specific and more general anecdotes can be extracted from the Empires' visualization. For example, each "pop" of a new territory is related with a specific independence event that may be identified by the audience. Furthermore, more complex and broad anecdotes can be observed in the visualization such as the independence of the French African colonies in the 1960s (see figure 4.17). In fact, this specific anecdote is arguably the most dramatic of the visualization, when the French Empire "explodes" into a large set of other territories, having most of its territory lost.

Several other anecdotes can be observed. For example, the Spanish can be observed controlling more areas than any other empire in 1805. Just before 1822, there is a short Portuguese predominance prior to the split of the Empire of Brazil that dramatically reduces the area of the empire. Around this period the Spanish Empire lost most of its colonial possessions, mainly from Latin America. In 1915 British hegemony can be clearly observed, with the corresponding body occupying most of the canvas. The disintegration of the British Empire, mainly through the loss of its major Dominions, are also dramatic moments that greatly reduce the British's territory.

Probably the most climatic moment of the visualization is the period 1960-1970, where successive disintegrations events of the British and French occur extensively and rapidly. After this, another remarkable moment is when the Portuguese loses its possessions in 1974-1975. From this period onward, the disintegration events happen at a slow pace, finalizing with the handover of the territories of Hong Kong and Macau, respectively from the British in 1997 and the Portuguese in 1999. This last historical note can be interpreted with a sense of irony and closure. The story of the disintegration of these empires closes with two very small territories emerging from minimized nations. The nations once occupied most of the canvas with vibrant colors, and are now surrounded in black.

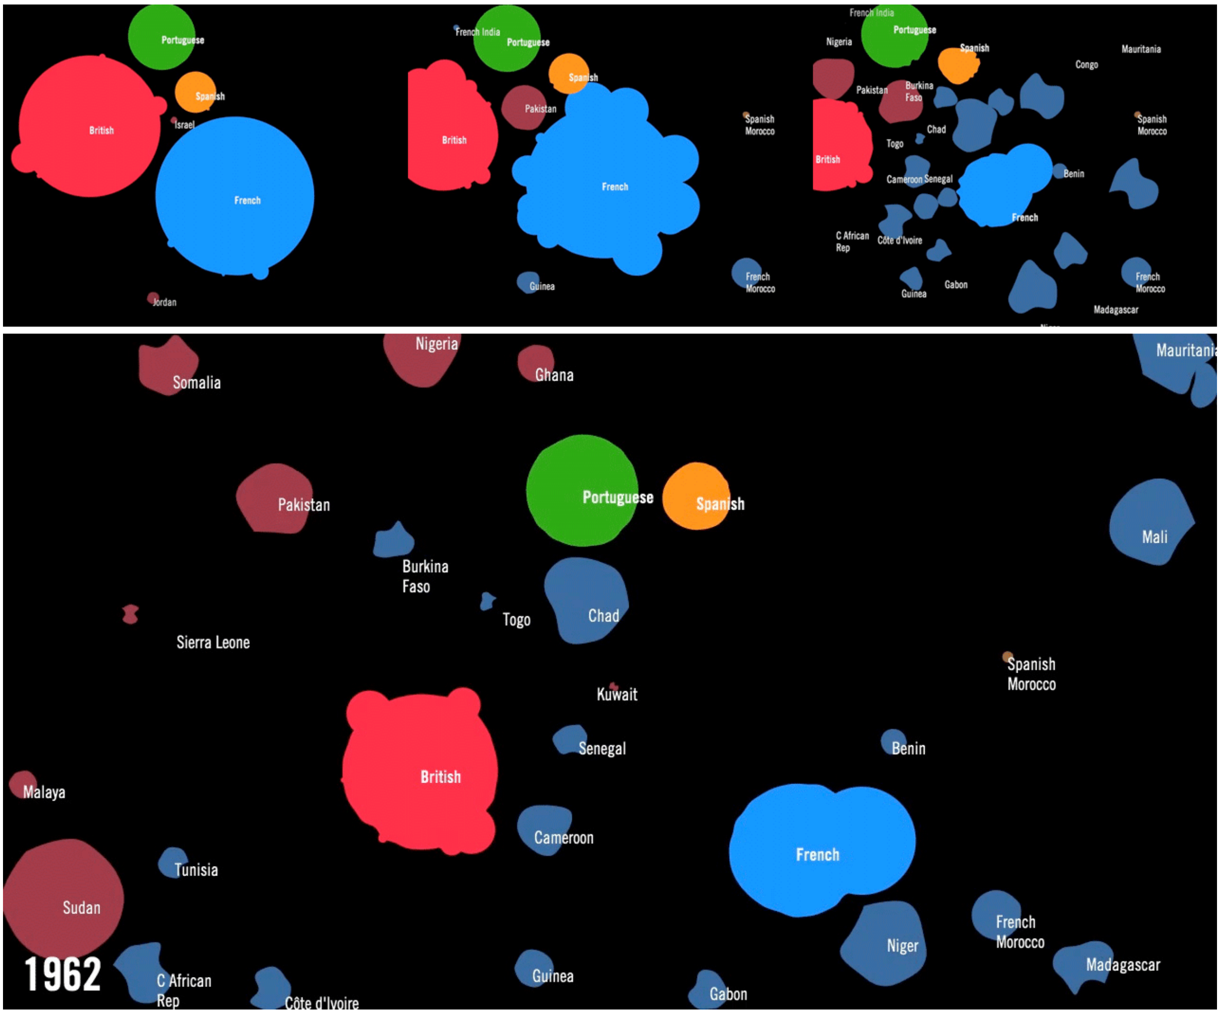

Figure 4.17 The visualization in 1960-62 when a vast quantity of French African colonies gained independence. It can be observed how all these territories build up tension by emerging on the French's empire rim and creating an organic and complex shape that call for attention to the eve of the events. The French Empire figuratively explores into a set of smaller territories, projecting them aggressively throughout the canvas, and has its area dramatically reduced.

The disintegration events can be observed in detail at the beginning of the visualization. Nonetheless, such attention effort rises as the visualization advances, with higher rates of events and higher number of bodies in the system that make the identification of each and every one a hard task, unless the simulation is stopped. This increasing rate of events is a translation of history, but also works to build momentums and narrative climaxes that confer a dramatic tone to the visualization.

4.4.5 Refinement

The second implementation of this visualization has a more sober tone, yielding a slightly different metaphor: the empires dissolve into new countries. Therefore the dissolution component is emphasized in detriment of the competition allusion that was the figurative backbone of the previous instantiation. With this, the new countries are not projected as fiercely as in the previous implementation, having more smooth and fluid interactions. The dissolution metaphor is more subtle, but is visually richer and more refined.

The violent thrusts of new nations that were observed in the previous visualization amplified the shape deformations of the bodies and caused difficulty when following some of the newborn nations. These aspects were toned down in this refined version as new bodies now leave the corresponding empire instead of being violently expelled from it. Nevertheless, the shape deformations and body interactions are still present in the visualization, since their remain a necessary core compromise when implemented these specific semantic figurative metaphors.

Refining the visualization passed through having more detailed and smooth bodies. This involves increasing the particle density of every body, and therefore the number of springs in the system, since the skeletons are more complex (see figure 4.18). The previous visualization had persistent center-to-center springs in order to implement the bodies' collisions. This fact together with the more complex skeletons for the bodies brings performance issues to the system. Since the number of springs for the bodies' collisions increase exponentially with the number of bodies, it culminates in a large number of springs that have to be simulated at the same time with different constraints. In order to tackle this issue, the center-to-center springs are not persistent anymore, and thus are only created if two bodies' centers are closer than the sum of their radius (plus a certain threshold), and removed from the system if this conditions ceases to exist. This results in a less computing intensive solution that brings the quantity of springs to a number that can be simulated at 30fps, and even enables the creation of other new types of springs that brings new behaviors to the system as described ahead.

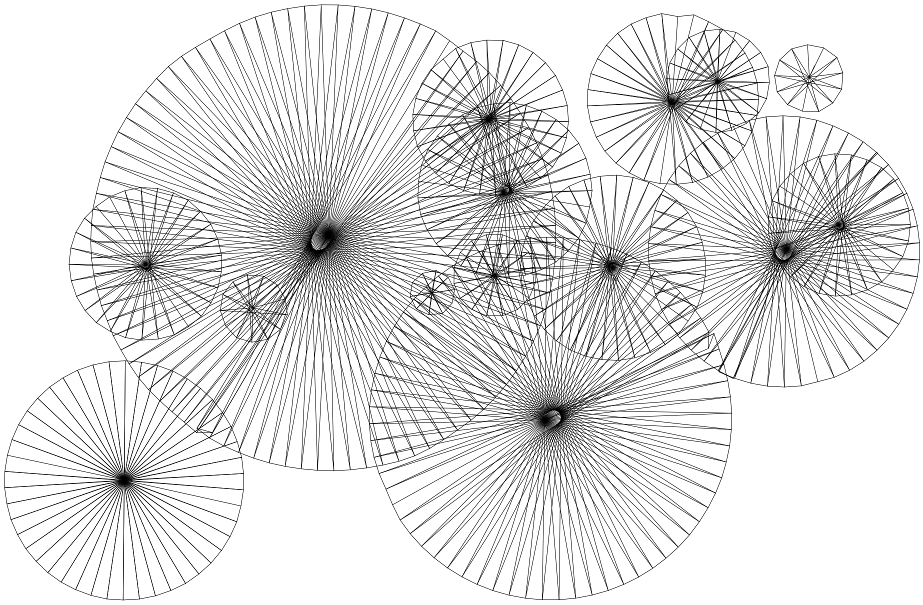

Figure 4.18 In the refined implementation of the Empires' visualization, the particle density of each body was increased, yielding more complex and robust skeletons.

In order to add finer details to the bodies' appearance and better show their soft-body properties, repulsion springs exist not only between center-to-center particles, but between the particles of bodies' rims. This way, two particles from different bodies' rims that are too close, have a spring added that pushes these particles away, generating new perturbations in the colliding bodies and contributing together with center-to-center springs to avoid juxtapositions. The bodies' rims are now jigglier and more fluid, but given that the skeleton's bodies are overall more robust, they resist radical alterations of the bodies' circular shape (see figure 4.19). In fact, having strong repulsions with only center-to-center springs as in the previous implementation often contributed to the radical distortion of the bodies as can be observed in figure 4.11.

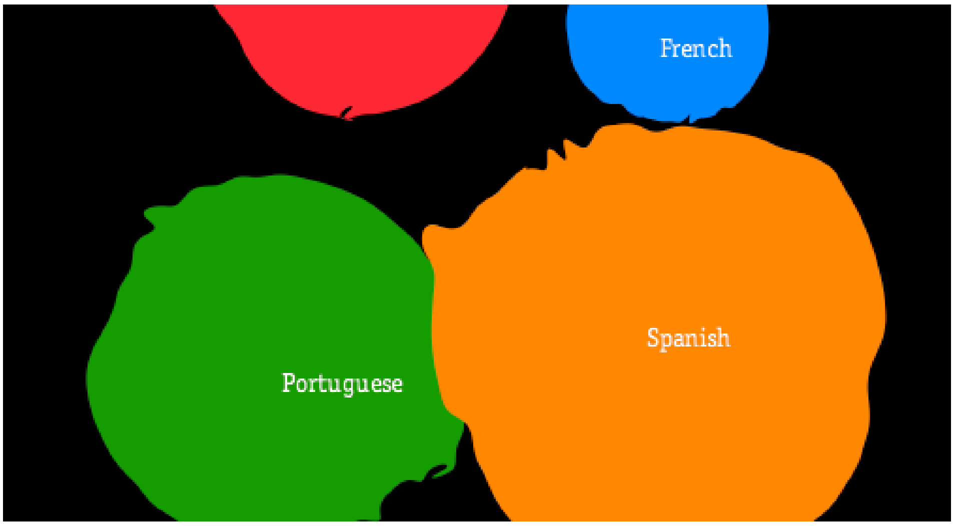

Figure 4.19 The new interaction in the refined version result in a more detailed graphic expression. In this case rim-to-rim repulsion forces cause jigglier deformations in the bodies' shapes, with more organic and fluid tones.

Another detail that was added to the visualization is the figurative deformation that a growing territory causes on the corresponding empires. When a new nation is growing on the rim of an empire, the associated rim's particle is pushed away from the center as well. This causes the whole skeleton's body to deform in order to accommodate the new constrain, pulling the opposite diametral particles and altering the rim's shape (see figure 4.20). This new behavior was implemented in order to attain a more figurative mitosis-like deformation, as if empires were splitting like cells or giving birth to new nations (see figure 4.21).

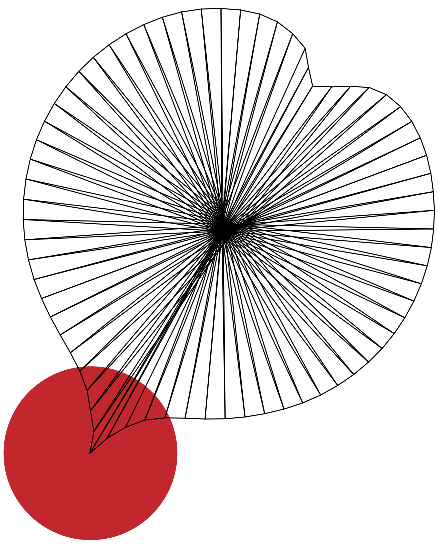

Figure 4.20 When a new nation grown on the empire's rim, the associated particle is pushed away from the empire's center, thus deforming the empire's skeleton and providing a more figurative mitosis-like split.

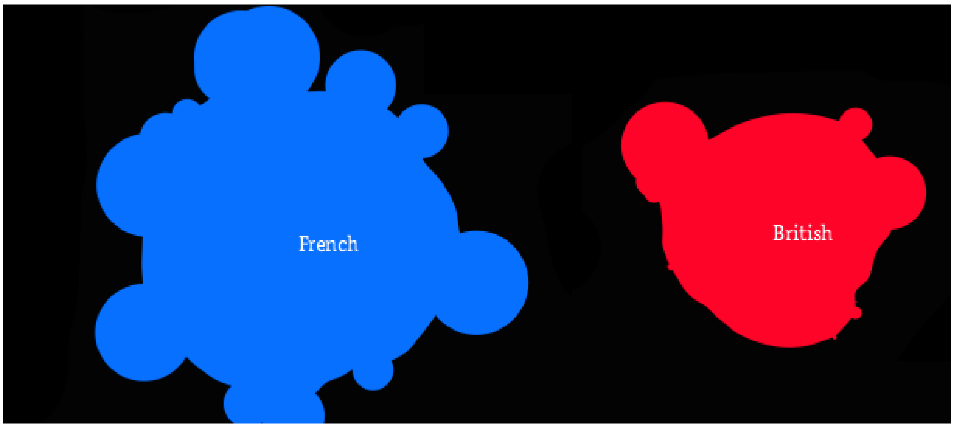

Figure 4.21 The British and the French Empires growing new nations. Each new soon-to-be nation pulls the bodies' rim, generating shapes with more organic expressions.

The structural metaphor upon this visualization is based was also altered in this refinement stage. Instead of positioning the bodies in the system based on a packed bubbles strategy, this implementation recurs to a closer approach on Dorling cartograms. Therefore, new geographic information is included in visualization through a new behavior: each empire and territory tends to be attracted to its geographical position. These visual positions were determined by visually and manually extracting the centroid of each nation on the simulation canvas, using a Mercator projected world map as a reference. Naturally, the centroids of an overseas empire were not determined by taking into consideration all of its territories, but rather used the corresponding metropoles' positions as reference. Using the strict definition of a centroid in these cases would result in an unfamiliar positioning of each empire.

In order to implement a variation of a Dorling cartogram in the Empires' visualization, the following was attained:

— Each central particle of a body is connected to its geographic position through a spring that pulls it into that position. Therefore, when a new nation splits from an empire, it travels to that geographic position.
— Each newborn nation or territory only collides with empires, but not with other nations. This diminishes the amount of interactions in the system that, when exaggerated, can be distracting and confusing. Furthermore, it allows new nations to superimpose themselves and economize space in the canvas.
— Contrary to the previous implementation, new nations do not fade way but stay in the canvas across simulation time. This is possible since they can now be juxtaposed and have more space-economical arrangements. Nevertheless, there is such an large number of bodies that accumulate over the simulation time that they have to be perceptually worked in order to not distract from the main actions and events. Therefore, new nations are only represented by their silhouettes or rims, in order to make them perceptually more silent and to allow for the readability of their juxtapositions. This solution creates a new type of map with a more geometric expression that iteratively conveys how much of the world was once part of each of the portrayed empires (see figures 4.22 and 4.23)

— When a new nation starts growing on the empire's rim, the empire is also attracted to the geographic position of the soon-to-be new nation. The empire tries to travel, amidst collisions with other bodies, until halfway to the new country's position. This way, when a split takes place, the new nation does not have to travel all the way to its position. This behavior brings the analogy of an empire traveling to deliver its newborn nation, as if it was "giving birth". When there are several nations growing on an empire's rim, the empire travels in the average direction of these nation's positions. For example if two nations are growing and will have to travel East, the empire starts traveling East. On the other hand, if two nations face opposite directions the empire will not forego this traveling behavior.

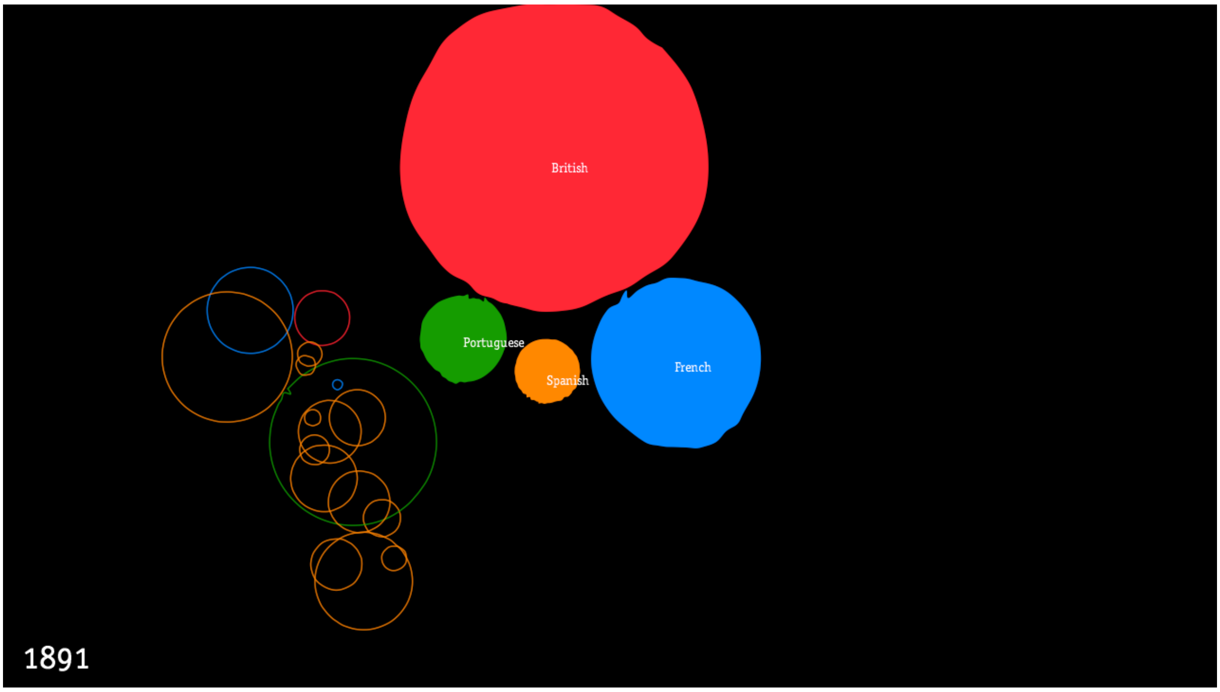

Figure 4.22 In this stage in the visualization, a forming Dorling cartogram of the Americas can be observed. It can be noticed that at this point in time most of the South America has gained its independence from the Spanish and Portuguese Empires.

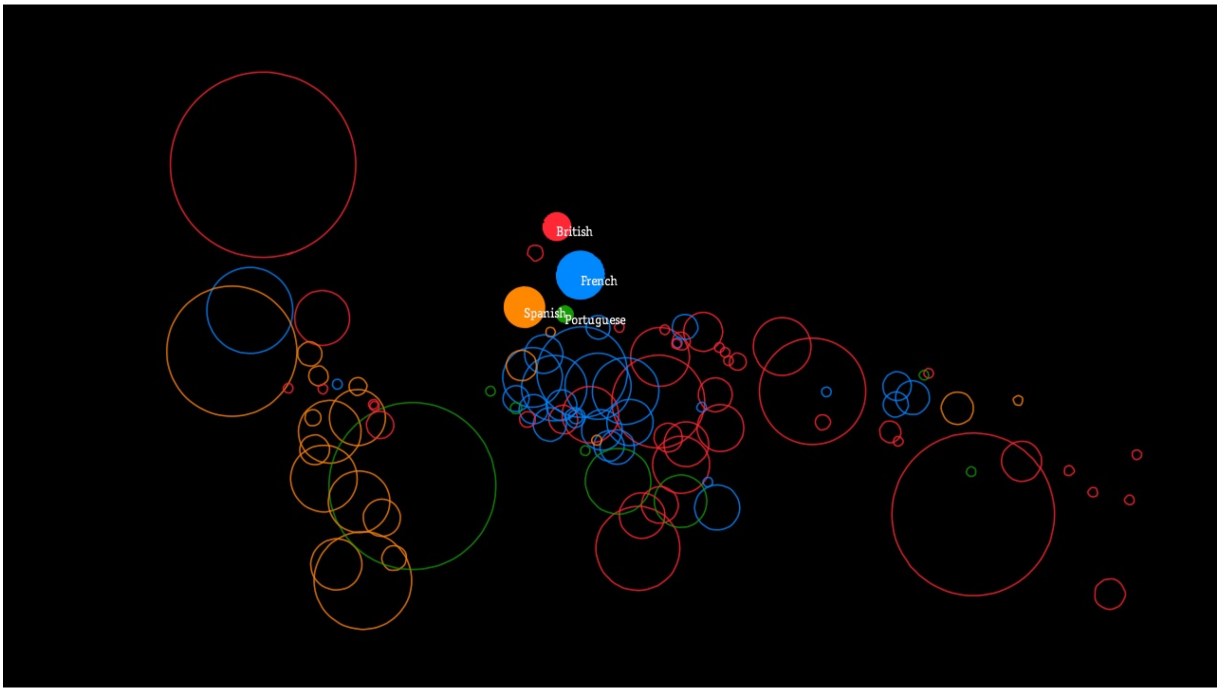

Figure 4.23 One of the final frames of the refined visualization, showing a modified Dorling cartogram, portraying how much of the world was once part of one of the four depicted empires.

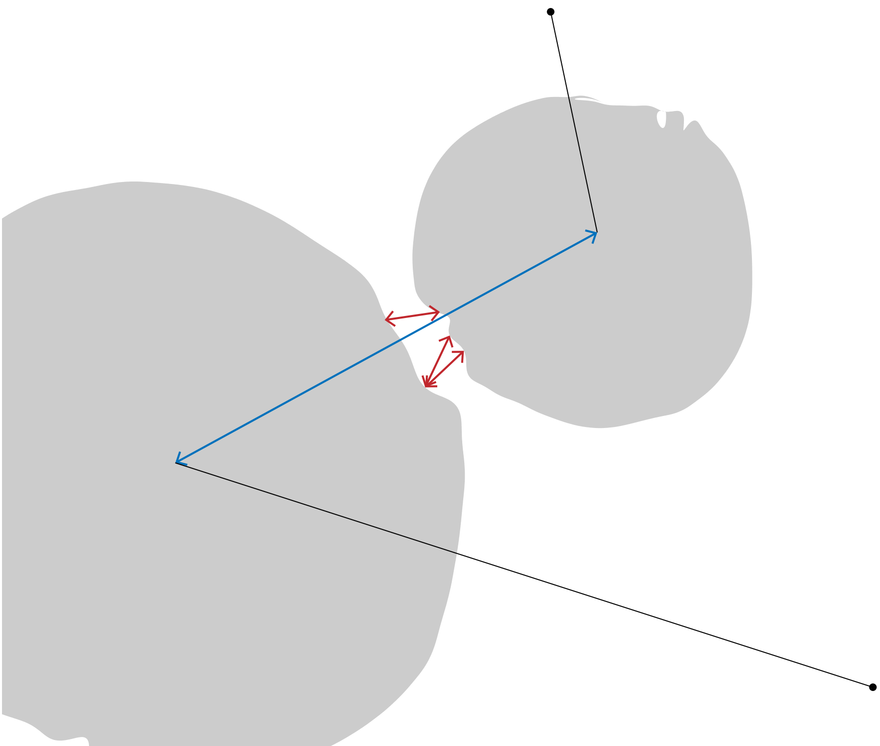

Figure 4.24 Two colliding bodies and the springs that act on the system. In black, the springs that attract each body to its geographic position that are persistent across simulation time. In blue, the center-to-center springs that repeal the bodies away, are created and removed dynamically as necessary. In red, springs that are added dynamically as well, repeal particles on the bodies' rims.

The springs that attract each empire to its corresponding position as a way to implement Dorling cartograms, together with the springs that create the collisions in this refined version of the system are illustrated in figure 4.24.

During the linear iteration of time in the visualization, it was observed that there are eventless periods in the narrative. Most notably, when the British Empire is expanding at the end of the 19th century, independences are non-existent in the data. In order to add more dynamism to the narrative, the program looks ahead to the near feature, and if there are no independence events, the year iteration rate is increased. This is to say that the timeline is no longer linear, speeding up during eventless periods, and slowing down when there are important events to portray. This constitutes a way to compact the visual story without compromising the observation of complex periods. The new simulation now starts in 1770, displaying other independence events such as of the Thirteen Colonies from the British, and has a new more compact duration of two minutes and 56 seconds.

4.5 On distortion

Tufte discusses how several graphs distort data when the visual representation of the data is not consistent with the numerical representation (Tufte 1983, p.55). Therefore, he advises against any non-direct encoding of numerical quantities. He devised a "lie factor" that tries to quantify how much a graph is lying by calculating the ratio of the size of effect shown in a graphic and the size of effect in data (Tufte 1983, p.57). The Empires' visualization inserts certain distortions in area representation in order to implement semantic figurative metaphors. These distortions can cause the misrepresentation of area quantities, and therefore it is important to assert to what magnitude such misrepresentations are taking place. Having such errors visualized and quantified is an analysis that can show if these errors are acceptable or negligible, or if the system, in spite of its constrained behavior, is distorting data in unacceptable ways.

Numerous experiments in psychological literature have shown that the perceptual judgment of circle's areas is not proportional to the area itself, but are rather underestimates of the actual area (Cleveland et al. 1982). These research studies suggest the perceived area in the form a power law: reported perceived area = (actual area)^x, where x is an exponent usually between 0.7 and 1.0. One of the earliest and popular reports on such power laws and its respective exponent is Stevens' Psychophysics. The book discusses the perception of several stimuli such as loudness, vibration, brightness, taste, and among other, visual areas (Stevens 1975, p.15). The exponent that Stevens suggests for areas is 0.7. Other studies suggest other exponents such as Flannery (1971), who suggests an exponent of 0.8747, adding that "underestimates were a fact" when interpreting the area of circles.

Divergences exist in the exponents reported by psychophysics experiments, which are in fact based on averages and do not apply for all test subjects. This situation impels Tufte to inquire: "… But what is a poor designer to do? A different graphic for each perceiver in each context? Or designs that correct for the visual transformations of the average perceiver participating in the average psychological experiment?" (Tufte 1983, p.56). He finalizes by adding that "At any rate, given the perceptual difficulties, the best we can hope for is some uniformity in graphics (if not in the perceivers) and some assurance that perceivers have a fair chance of getting the numbers right. (…) The representation of numbers, as physically measured on the surface of the graphic itself, should be directly proportional to the numerical quantities represented". Tufte's discourse is in line with the common practice of not applying these perceptual adjustments when designing visualizations, which was the case of the Empires' visualization. Tufte also adds that he considers distortions in relation to the size of effect in data to not be substantial if they are below 5%.

Cleveland et al. (1982) found an average exponent of 0.95 when judging circle's sizes comparatively with other circles. With these results, they adopt a similar discourse with Tufte's, recommending that areas should be encoded directly, since 0.95 is much closer to 1.0, and in fact falls within their 95% confidence interval. In spite of this advice on using direct mapping instead of the exponent function, the perceptual misinterpretation for areas exists, and its subjective nature has impeded a numerical consensus on the exponent. If the Empire's areas would to be perceptually adjusted the most conservative exponent of 0.95 (Cleveland et al. 1982) would be used, and using the same units of that experiment: mm^2. Therefore, each area of an empire would be scaled up in order to compensate for the perceptual underestimate and would possibly result in a more precise judgment of the area. So, by inverting 0.95, the adjustment to be made would be $A^{1.0526}$ for each circular area. If by chance the actual areas resulting from the Empire's simulations lie between the expected areas and the adjusted areas for perception, it can be said that their deviations are negligible in the sense that they tend to naturally correct for judgment errors.

The generative characteristic of the Empires' visualization yields in a singular aspect that is random throughout the simulation. The new nations erupt from the edges of the empires, but the exact position on such a circumference is computed in runtime using pseudo-randomization. Therefore, each time the simulation is run, the new nations indeed detach from the main empire, but the position of eruption is random. This situation sometimes incurs more complex outcomes from the visualization, where collisions are maximized and, consequentially, their areas can be vastly impaired beyond what could be deemed acceptable, even if for short periods of time. Having this in mind, it is of interest to observe how vast and frequent these distortions are. The simulation was run one thousand times in different random seeds in order to have different eruption locations for the new nations. Graph 4.4 shows the total desired areas for the four empires over simulation time, together with the average of the actual total areas of the one thousand experiments, as rendered by the visualization. On first analysis the distortions applied by the system seem to accompany the desirable profile. When such is not happening, usually the areas are being overstated or positively exaggerated. This comes from the natural expansions of the empires when land acquisitions are happening – giving their spring-based nature, tend to over-expand until they reach their equilibrium state. Several understatements of areas are also present and come as well from their elastic nature: when there are great land losses, each empire quickly shrinks to a much lesser area, incurring less than desirable areas for a sparse amount of time until it reaches its equilibrium state.

Graph 4.4 Totals of the desired and actual areas of the empires for the one thousand simulations in mm^2 (— sum of the empires' desired areas, — average sum of the empires' actual areas). Several major land losses are highlighted in the graph: Louisiana (1803), Mexico and Brazil (1821-1822), Canada (1931), Australia (1939), French African colonies (1960), and Portuguese African colonies (1975). The areas vary in consonance with the system's oscillatory nature, with, generally, land gains resulting in positive exaggerations and land losses resulting in negative area exaggerations. In simulation time, those exaggerations are constrained to short time spans.

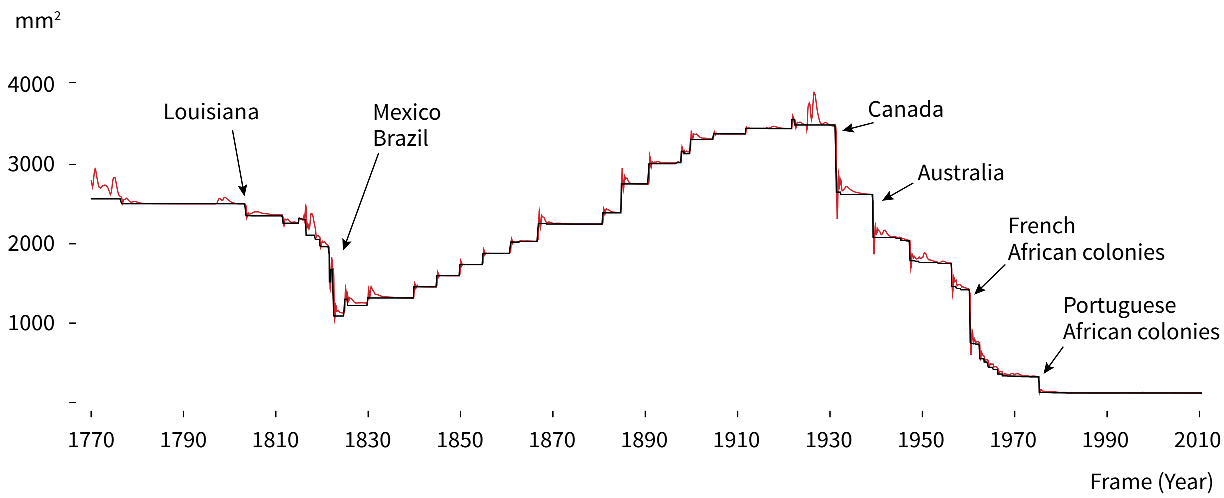

Graph 4.5 shows the total desired areas for the four empires during simulation time, along with the range of values for total actual areas, considering the one thousand experiments. It also shows the perceptual adjustment based on Cleveland et al.'s (1982) exponent of 0.95, if applied to each circle individually and then summed. Except for rare understatements in the areas, the visualization generally overstates the areas below the suggested perceptual adjustment, which can suggest that the exaggerations are acceptable and go in the direction of correcting areas for perceptual interpretation, even if their magnitude is small. Even if this seems to be the general case, graph 4.5 already provides hints when such is not true. Take, for example, the period between 1810 and 1830 where the perceptual adjustment seems to intercept with the range of maximums and minimums, or from 1960 onwards where the maximums are clearly above the adjustment line for perception. In order to investigate these cases, a more sensible approach to measure the deviations from the desirable is devised.

Graph 4.5 Totals of the desired areas, the adjusted areas for perception and the range max-min of the actual areas, for the four empires for the one thousand simulations in mm^2 (— sum of the empires' desired areas, ■ range min-max for the empires' actual areas sum, ····· sum of the areas adjusted for perception). Generally, the actual areas lie between the desired and the adjusted lines, meaning that they overstate areas in a way that compensates for perception.

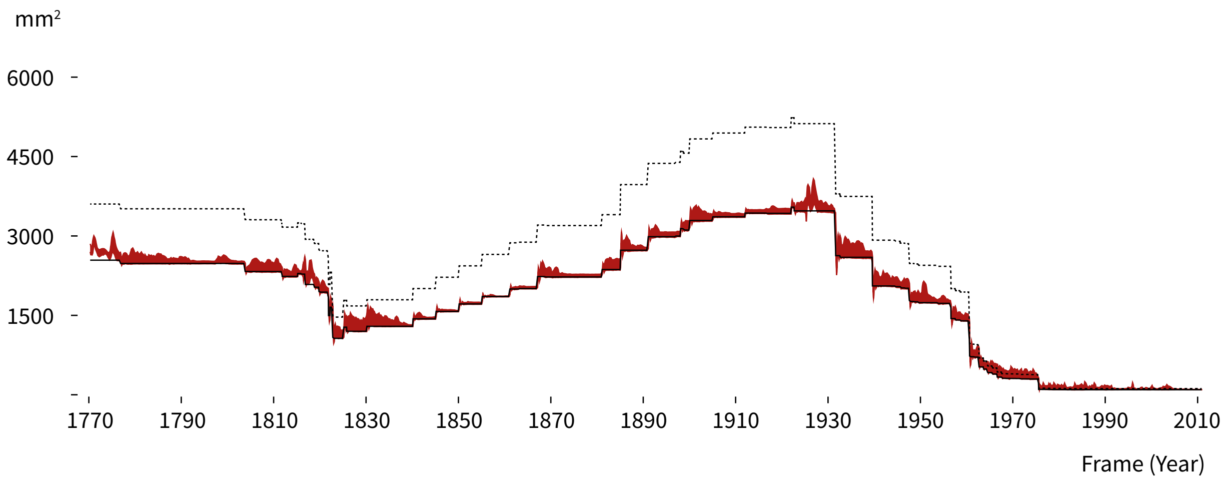

The deviation of the actual areas from their desirable errors can be seen as an error in the visualization. Each empire $i \in$ {British, Portuguese, French, Spanish} has both a desirable and an actual area at some point in visualization time – $A_{desired}$ and A_{actual} respectively. The error on an empire e^i is the absolute deviation from its actual area to its desirable area, relatively to its desired area. The error E considering the four empires for each frame of the simulation is then the average of the four errors of each empire e^i, weighted by each of their desired areas. This way, errors in bigger areas will have a larger weight in the final E than errors in smaller areas, as they in fact account for higher absolute deviations. The errors are calculated with the following formulas.

$$e^i = \frac{\left|A^i_{desired} - A^i_{actual}\right|}{A^i_{desired}} \qquad E = \sum_i \left(e^i \cdot \frac{A^i_{desired}}{\sum_i A^i_{desired}} \right)$$

The perceptual adjustment of the areas, although not being an error, can be seen as a relative distortion ratio that could be applied to the areas. Therefore, in a similar way to the errors, the perceptual adjustment ratio D of each frame is calculated in function of each of the adjusted ratio for each of the empires (d^i) with the following formulas.

$$d^i = \frac{\left|A^i_{desired} - A^i_{adjusted}\right|}{A^i_{desired}} \qquad D = \sum_i \left(d^i \cdot \frac{A^i_{desired}}{\sum_i A^i_{desired}} \right)$$

Graph 4.6 shows the one thousand visualization runs, the average error for each frame $\overline{E}$, the maximum error for each frame, the 95th percentile of the error E, and the distortion ratio for perception adjustment D. This graph shows that the error is close to zero when there are no independence events or expansions. Most of the time, it is argued that the error is well controlled, rising momentarily in the advent of independences. Nonetheless, losing large territories causes the error to spike, even if for a few frames, and beyond the perceptual adjustment ratio. One of such events is the independence of Brazil and Mexico (1821-1822). This is caused by the independence of large new nations that are expelled with more thrust from the corresponding empire, creating a higher disturbance on the system that amplifies the collisions with other empires and therefore might alter their represented areas, while also creating a large area gap to which the empire has to adjust in a short time-span, causing a big error during that adjustment period.

The same can explain the spike in 1960. In 1960 the French Empire loses much of its area into a large number of bodies that greatly amplify the number of collisions in the system. The largest average error in the system is in 1975 when the Portuguese lose their African territories. The land loss is so great relative to the area of Portugal afterwards and the areas of the other empires, which already had lost most of their possessions, that smaller oscillations in the area of Portugal and the other empires cause a dramatic increase in the relative error. By no means are errors of such magnitude acceptable, but the duration of such a large error is of only two frames. The small dimensions of the empires in the late 20th century, as well as the increased number of new bodies in the system, can brutally increase the errors. As can be seen on graph 4.6, after 1960 the maximum error of the one thousand experiments can largely go beyond the perceptual adjustment ratio to unacceptable magnitudes. This shows how this simulation time period can be problematic, and that it could not be clearly spotted in graph 4.5. Nonetheless, since one can also observe the 95th percentile curve on graph 4.6, most of the errors are well controlled, close to the average error, and only less than 5% of the one thousand runs presented such high errors for the referring period. Such deviations are a natural consequence of the system's generative properties that, by creating new bodies in cluttered positions, maximizes collisions and can cause great disturbances in the empires' areas. These errors have the tendency to increase when the empires are small, since they are constituted by fewer particles, have less mass and fewer springs to maintain their bodies' shape, making them more vulnerable both to collisions and the natural oscillations that are caused by having their sizes altered.

Graph 4.6 Distortion errors for the one thousand visualization runs for each frame of the periods 1770-1920 and 1921-2010. (— average error $\overline{E}$, — maximum error $max(E)$, — 95th percentile of the error, ····· distortion ratio for perception adjustment D).

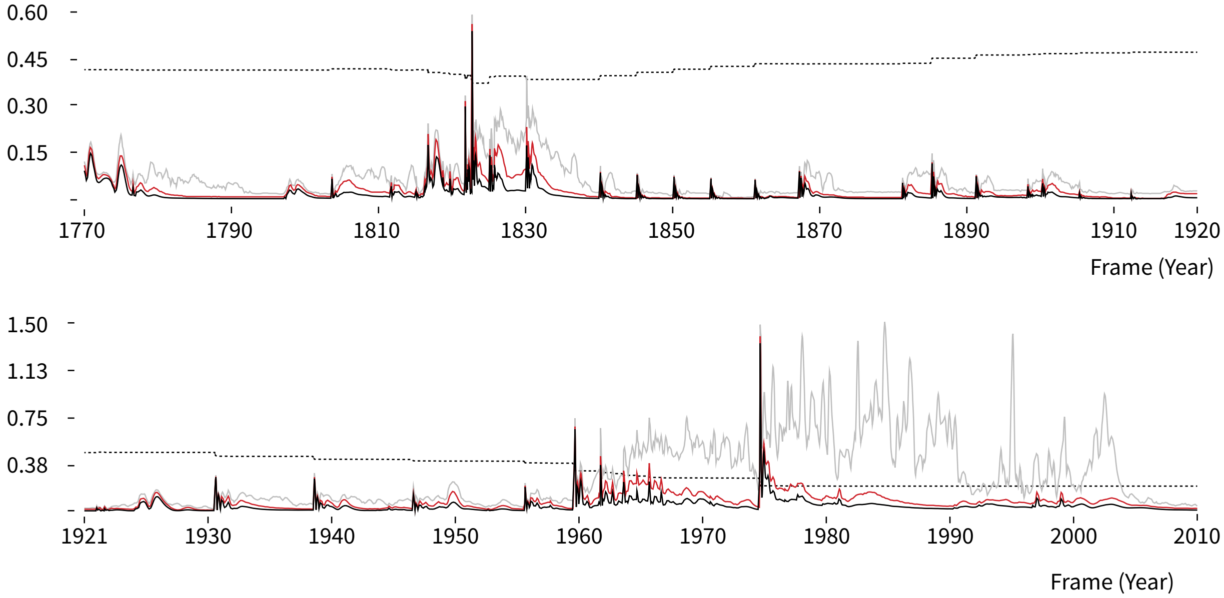

The design decisions that led to expression of more figurative metaphors in this visualization came at the expense of an error in the areas' portrayal, or using Tufte's terminology, a "lie factor". The computed errors can roughly be interpreted as percentages, being an error of 0.15 a 15% deviation of area of what the visualization was supposed to display. Tufte suggests that errors below 5% are acceptable (Tufte 1983, p.57), and here it is argued that such an amount would be negligible. Nonetheless, even if momentarily, the error easily surpasses 0.05. The average of $\overline{E}$ throughout the simulation time is 0.027, and is below 0.093 for 95% of the frames. The 95th percentile of the error is 0.049 on average throughout the simulation time, and for 95% of the frames it is below 0.16. These error ranges are excusable in the context of this visualization for the following reasons:

- The premise of implementing figurative metaphors makes room for these errors with the purpose of having an authorial expressiveness intent and a stronger connection with the audience. In order to attain this, the empires are often depicted as irregular circles, with occasionally undulating edges. These deformations are neither suitable nor adequate to be quantitatively judged in terms of area, and they were not designed for this.
- The use case of this visualization is to extract ordinal judgment of areas rather than precise and quantitative judgments. It provides a synthetic way, through a non-interactive animation, to depict more than two hundred years of independence events. The visualization is ludic and informative, but was designed for a casual context in information visualization, and therefore less utilitarian for data exploration. The errors observed do not affect this type of ordinal judgment task among empires.

— The errors are not hidden in the visualization, being portrayed as natural distortions of the system. The visualization clearly exhibits its physical behavioral characteristics and thus shows area deformations as a consequence of physical events, clearly prompting the viewer that the bodies are in a process of adaptation and stabilization.
— Throughout each simulation run, and among simulations, the error is generally below the perceptual adjustment ratio for circles, indicating that the deformations might indeed compensate for this perceptual adjustment, and that the errors fall in judgment error ranges contemplated by the literature (Cleveland et al. 1982).

4.6 Adoption and reaction

> Oh, this could be so expanded upon... let's include ALL the empires... the Soviets... the American Empire... the Chinese Empire... Unfortunately, the author of this wonderful clip was a byproduct of the anti-British-imperialist prejudice... but he nibbled on something that is worthwhile.
>
> —Anonymous YouTube user

The first implementation of the Empires' decline visualization was launched on Vimeo on September 4, 2009. In its first four months it was able to gather 320,284 plays. The visualization was submitted to specialty websites such as datavisualization.ch[22] and flowingdata.com[23], and through there it has spread through the blogosphere and highly popular websites such as Boing Boing[24] and The Huffington Post[25]. After this period, the visualization counted a total of 464,680 plays by November 20, 2015, and had 1,184 likes, with an average of 2,033 plays per month. The total number of people watching from start to finish is 267,452, equaling 58% of plays in total. The video was also downloaded 390 times.

The video was also published on YouTube[26] on October 5, 2009. It gained most of its views during the first four months online (98,575). By November 20, 2015 the visualization received 229,055 views, with an average of 1,864 views per month. The views were mostly fed by YouTube itself (76%), being incorporated in suggested videos lists, users' lists, and available through YouTube's search. The video had a duration of three minutes and 30 seconds. Users watched two minutes and 14 seconds of the videos on average, so 64% of the total video duration. The video gained 959 likes and 54 dislikes, which is 5% of the users who voted.

The Vimeo and YouTube dissemination of the first implementation of the Empires' Decline visualization accounts for a total of 693,735 views. The magnitude of such results shows that the video went "viral". The geographic distribution of the users can be seen for Vimeo in table 4.1 and for YouTube in table 4.2. It can be seen that in both tables France, Spain, Portugal, and the United Kingdom are among the top proveniences, but the United States dominates the positions, both because of their users' strong web presence, and arguably because of their interest in the history of Western colonialism. Some territories that were once from the French, Spanish and British are now incorporated into the United States.

On Vimeo and the personal website pmcruz.com[27], the visualization received a total of 132 voluntary and spontaneous comments. Ninety-seven of these comments complimented the visualization with adjectives such as "beautiful", "original", "hypnotic", "attractive", "dramatic", "organic, "visually stunning", "historically compelling", "efficient", "instructional" or "mesmerizing". Twenty-eight comments suggested that other countries' empires could be included, such as the U.S., Russia, China, Germany, the Netherlands, Austria-

Hungary, Japan, or the Ottoman Empire. This naturally would involve a change to the time-span covered by the visualization in some cases. Twenty-six comments discuss the dates of independence events, pointing out that in their opinion the years of some of the independences are inaccurate and that some independences are missing. The way these dates were chosen are justified on the Data section, as a way of making uniform the procedure of choosing independences for every territory. It is natural that some of the dates – using the principle of setting the date based on the first sign of an independence struggle – is not universally accepted, varying with each viewer's cultural and educational background. Six comments suggest that other statistical variable could be used to translate an empire's influence, for example GDP or population, instead of territorial area. Five comments pointed out that the visualization is hard to watch, being either too quick or too slow, and that it should be interactive in order for the users to control the timeline. Three comments explicitly said that the viewers did not understand what the area was representing in the visualization, and two comments presented confusion as to what the "wobbles" in the bodies represent, inquiring if it was internal conflicts or if the collisions represented anything. One comment explicitly stated that the user could not understand the visualization, and another comment classified the piece as "infoporn". Three comments suggested that the film should have an accompanying soundtrack.

Table 4.1 Top 13 countries with the most plays on Vimeo.

Country	*Views*
United States	168,789
Spain	84,849
Canada	19,501
United Kingdom	18,349
France	17,655
Japan	13,827
Russia	13,204
Brazil	12,278
Germany	10,729
Mexico	7,578
Portugal	7,134
Australia	6,788
Netherlands	4,907

[22] https://datavisualization.ch/showcases/visualizing-empires-decline/
[23] http://flowingdata.com/2009/12/24/the-decline-of-maritime-empires/
[24] http://boingboing.net/2009/11/17/visualizing-the-decl.html
[25] http://www.huffingtonpost.com/2009/11/27/the-decline-of-empires-vi_n_372259.html
[26] https://www.youtube.com/watch?v=EwOA8AfeHM4
[27] http://pmcruz.com/visual-experiments/visualizing-empires

Table 4.2 Top 10 countries with the most plays on YouTube.

Country	*Views*
United States	70,340
Russia	18,990
Portugal	15,654
United Kingdom	12,042
Poland	9,798
Canada	8,548
Sweden	7,239
Brazil	6,244
Latvia	5,975
Spain	5,673

The film on YouTube received 345 comments, but with a different distribution of subjects. In fact, 68% of these comments involved a long discussion among users as to whether the United States could be considered as an empire or not. Similarly to the Vimeo's comments, 42 also suggested the inclusion of other empires in the visualization. Like the Vimeo's comments, 30 comments used similar complimentary adjectives to refer to the visualization. Additionally, 45 comments discuss the dates chosen for independence events. One comment also questions what the "wobbles" signify, and another suggests an interactive visualization to deal with the timeline incrementation. Two comments suggest a soundtrack and three comments did not understand the visualization and classify it as "pointless". On a side note, six comments compare the Empires' visualization to the multiplayer online game named agar.io[28], where each user is a circular mass that can absorb other smaller bodies, but has to escape bodies with higher mass. The main categories of comments between Vimeo and pmcruz.com and Youtube are synthesized in table 4.3

Table 4.3 Categories of comments for the Empires' visualization on Vimeo and pmcruz.com compared with Youtube.

	Vimeo and pmcruz	*Youtube*
Total comments	132	345
Complimentary	97 (73%)	30 (9%)
Include other empires	28 (21%)	42 (12%)
Independence dates	26 (20%)	45 (13%)
Discussion on the U.S as an empire	—	234 (68%)

The refined version of the visualization was only published on Vimeo[29] on May 5, 2010. By November 20, 2015 it had 26,683 plays, with a peak of 6,349 in August 2010 and a steady rate of 313 plays per month since then. The number of plays start to finish is 7,936, which is 30% of total plays. This visualiza-

[28] http://agar.io
[29] https://vimeo.com/11506746

tion, being only a refinement of the first implementation was not disseminated through visualization-focused websites and thus did not gain the same attention neither in the special nor the general media. An exception can be made in the peak of August 2010 when the University of Coimbra disseminated[30] the work through Portuguese general news websites as a recipient of the ACM SIGGRAPH Student Research Competition 2010.

The refined version addressed the following concerns expressed in the comments to the first implementation.

— The visualization being too slow or to quick through the non-linear time iteration.
— The introduction of a small in-video caption and title that clarifies what the areas signify.
— The introduction of a soundtrack composed specifically for this visualization[31].
— The slight extension of the time period as the visualization starts now in 1770, thus including the independence of the Thirteen Colonies of United States from the British.
— The addition of Ireland that was a country that was indeed missing from the dataset.

[30] http://www.publico.pt/tecnologia/noticia/estudante-portugues-premiado-em-conferencia-mundial-de-computacao-grafica-1449903#comments
[31] Music composed by CHOP WOOD – chopwood.eu.

4.7 Conclusions

> I like how the subaltern turns dark and disappears. Good metaphor for their actual role in global politics.
>
> —Anonymous Vimeo user

The visualization of the Empires' decline is an exploration of semantic figurative metaphors and employs a conceptual framework for generative storytelling. Building a system to tell the story of the decline of the Empires involved translating the countries into actors, the growths and disintegrations into time-based events and the set of behaviors that arise from the system's interactions. Each time the simulation is run, it tells the same story with slightly different expressions. This results from the randomness of the new nations' positions over the empire's rim. Different positions mean that they can be expelled in different directions, colliding with other bodies in the system and originating perturbations that are different each time the simulation is run. This property of the system enables different expressions for the same narrative, yielding the generative characteristic of the stories. The visualization invests in the anticipation or narrative prolepses of disintegration events, growing new bodies on the empire's rim prior to the event's time, therefore directing the audience's attention to that area of the canvas while building tension in the narrative.

The Empire's visualization had two implementations. The first invested in the structural metaphor of packed bubbles and added visual cues that allude to a sense of disintegration and competition. The second implementation is a refinement of the previous, basing its layout on Dorling's cartograms and displaying visual cues for a sense of smooth dissolution with more organic tones. This refined implementation is able to show which parts of the world were once part of the chosen empires, alluding to its contemporary heritage.

It was observed that by implementing more figurative metaphors with system, one is introducing errors in the portrayal of information. These errors were measured and it is argued that they are acceptable since they fall in line with common judgment errors.

The visualization went viral, having a total of 693,735 views since launch. The viewers engaged with the videos by actively watching them and leaving comments. Most of the comments were eulogistic, but the users expressed concern in how the independence dates were chosen and suggested more empires for the visualization. On a smaller scale it was suggested that other variables rather than land area should be used to express an empire's influence, and that the "wobbles" in the visualization are prone to confusion. Choosing other statistical variables to portray the influence of an empire such as population or GDP would be a natural second step in evolving this data and visualization. Using a land area to depict an empire's influence is a familiar technique for mapping empires and thus is a natural first step when exploring semantic figurative metaphors. The non-data movements that can cause confusion to some users are used to strengthen the metaphoric value of the visualization and its dramatic tone. In fact, most of the comments on the visualization praise it as a whole, but they miss to verbalize which specific metaphors were observed. Even if the metaphors of competition, dissolution, and disintegration cannot be directly unmasked from the viewers' reactions, it can be observed that certain aspects of the visualization can be interpreted with unintentional metaphors, giving some space for viewers to build their own speeches on the data.

5

City traffic and cartograms

This chapter describes several metaphorical mappings of city traffic. The objective of this work as whole is to explore a new type of cartogram: edge-based cartograms. These cartograms formulate a new model to depict traffic information in a dramatic way, by distorting the city's road network. In the model, roads that have higher velocities become compressed as if distances were shorter, and roads with lower velocities expand as if traffic congestion makes perceived distances larger. The model adapts dynamically to data in simulation-time, presenting smooth animations between different data states that show a continuous narrative of the evolution of traffic in a city.

The first dataset to experiment consists of GPS traces for several vehicles in the city of Lisbon. This data is described and filtered in order to create a visualization of trajectories where color mapping has a semantic meaning. An additional visualization model that is more figurative has also been developed, which was inspired by the metaphor of the "city as a living organism", and depicts it as a series of intertwined blood vessels that pulse according to traffic conditions.

The edge-based cartogram is implemented and tested using Lisbon's data. The impact of the model's parameterization is studied. As with any contiguous cartogram, this model brings errors during data representation. These errors are reported and approached in order to obtain the parameterizations that present the best compromise among several types of errors, while also observing how it influences the dramatic effect of a cartogram's expansion and contraction.

The model provides a skeleton to distort the city in simulation time, but more information can be provided if additional visualizations are applied over the cartogram. The applied visualizations are two: the visualization of trajectories and the "blood vessels" visualization. The blood vessels visualization over the cartogram shows semantic figurative metaphors at its fullest: while distorting the structural model using a cartogram, it also provides non-data related visual cues with a figurative and semantic representation, in the context of the chosen metaphor.

Finally, the cartogram model is tested again using similar data but for the city of London, and its results are compared with Lisbon's. The cartograms produced by London's dataset are rendered using the blood vessels visualization in order to show how this visualization can also be applied to other cities.

5.1 Related work on cartograms

Cartography has long since mapped the world and our surroundings to graphically amplify our cognition of what is not directly tangible. Cartography has moved to map more than geographical information and can present forms of statistical information. Statistical maps can depict a multitude of information: sociological demographic or even time-constrained events. Nowadays, cartography is intrinsically close with computer graphics, integrating its tools and techniques to map information. Moreover, the frontier between cartography and information visualization is more blurred than ever, considering the advent of new cartographers and techniques, as well as the reinterpretation of old techniques to dynamically layer various types of information on a map.

Information visualization, in its frequent applications in geographical contexts, has similar objectives to cartography, but has driven its techniques from an amalgamation of fields such as computer graphics, interaction design, statistical graphics, and cartography. General maps are driven by a scale and a system of projection, but more special maps such as cartograms bring more variables than scale and projection to the mapping of space. While cartograms can be highly abstract or diagrammatic, they have historically been a legitimate tool in cartography (Raisz 1948, pp.11-12).

The term cartogram have been used throughout history to refer to several types of maps (Tobler 2004). Funkhouser (1937) refers to choropleth maps[32] as cartograms and Raisz (1948, p.256), admitting that "some authors in Europe call every statistical map a cartogram", states that a cartogram may be defined as a diagrammatic map. By diagrammatic he means any abstraction, conventionalization or selective use of the elements of a map. The first cartogram is widely attributed to Pierre Levasseur in 1870, who used one in his classes and geography textbooks (Funkhouser 1937, p.356). Levasseur's graphic, as seen in figure 5.1, is a map of the countries of Europe represented by squares, with size being proportional to the area of the country. The squares are grouped in a way that approximately reflects their geographical positions and adjacencies. Tobler (2004) points out that such a map could be called an "equal-area cartogram". In fact, Levasseur's map distances itself from the contemporary convention of cartograms since it is more a diagrammatic representation of space than a distortion of space that illustrates an additional statistical variable.

[32] A choropleth map is a thematic map in which areas are colored or patterned in proportion to the measurement of a statistical variable.

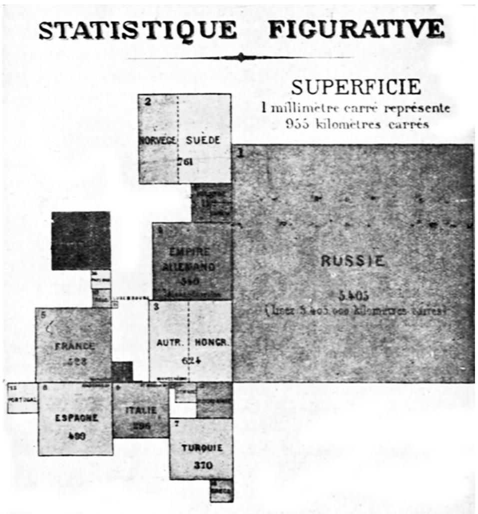

Figure 5.1 Levasseur's cartogram from 1870, representing countries as squares with an equivalent land area as seen in (Funkhouser 1937, p.356).

Nowadays cartograms are viewed as distorted maps that resize their regions according to a geographically-related parameter (Keim et al. 2004; Krygier and D. Wood 2011, pp.192-193). In 1911, Bailey (1911, p.722) published a U.S. map that tessellates semi-rectangular states, with size being based on their population (see figure 5.2). All the states fit inside the U.S.'s continental boundaries, and while topology is not strictly preserved, it was one of the first cartograms to distort space in order to represent population. Bailey called his cartogram "apportionment map", a less vague term than that came to be known as a "cartogram".

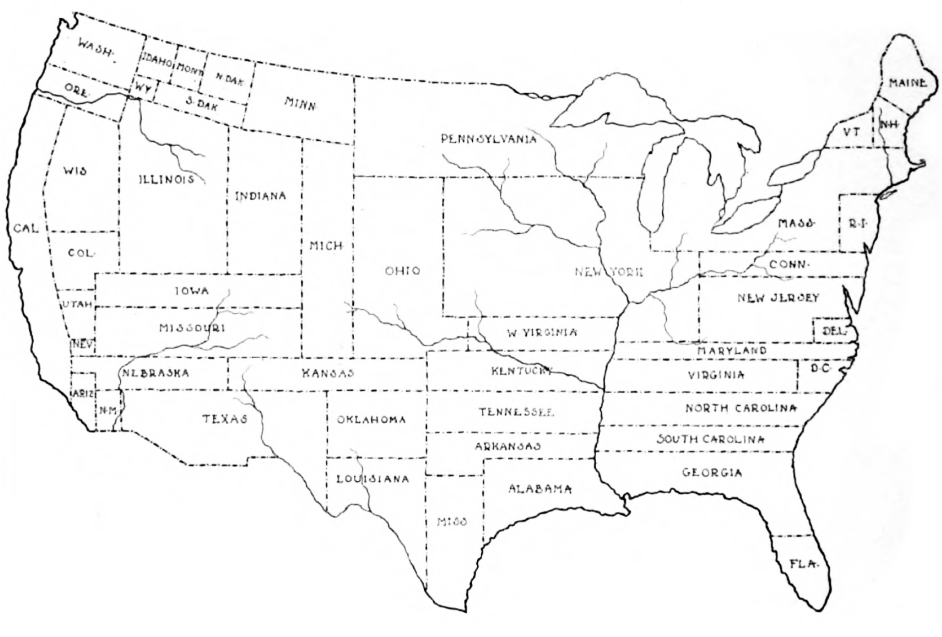

Figure 5.2 Apportionment map of the U.S. (Bailey 1911, p.722), one of the first cartograms to use population as a statistical variable to distort areas.

In 1921, “Electronic World” published a U.S. map that resized states based on their electricity consumption while trying to maintain topology and the original shape of the states (see figure 5.3). In 1923, Karsten (1923) published a methodology for cartograms that he called “population projection”, applying it to a map of the U.S. that resized the states accordingly to their population while trying to preserve shape. The technique was patented later on (Karsten 1925). Another notable cartogram that preserves shape and topology is Grundy’s map of the U.S. that resized each state based on population and taxes, which was published by the “Washington Post” in 1929 (Tobler 2004).

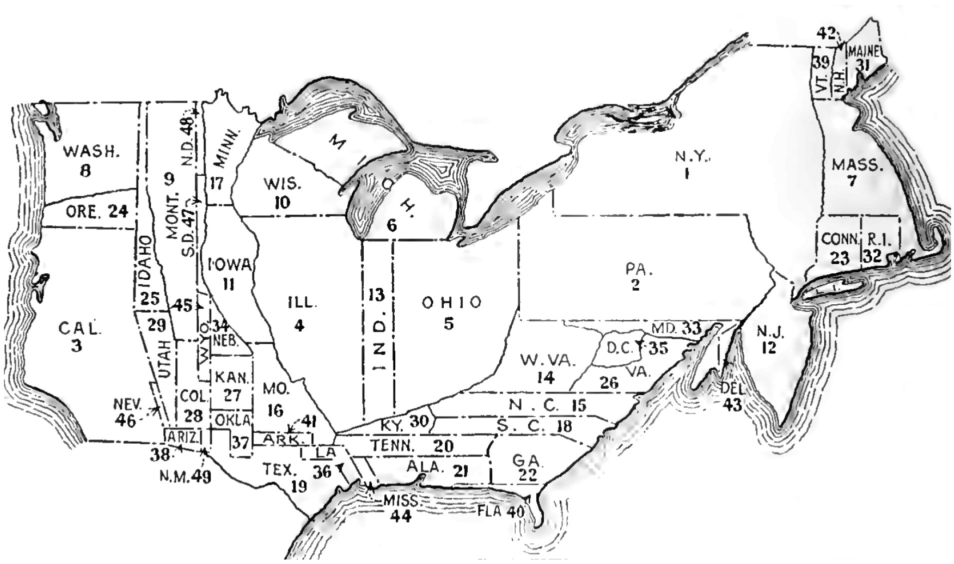

Figure 5.3 *Relative Size of States Based on Electrical Energy Sold for Light and Power* (Onken and Braymer 1921), was one of the first cartograms that tried to preserve topology while distorting area according to a statistical variable.

It was not until a later date that cartograms were popularized by Raisz[33] in its rectangular variant. In 1934, he published some of the first statistical cartograms of the U.S. (Raisz 1934), one of which can be seen in figure 5.4. Raisz understood that value-area cartograms implied dividing regions or countries in rectangles and resizing them according to a statistical variable (Raisz 1948, p.257).

[33] Erwin Raisz was a professor of cartography at the Institute of Geographical Exploration at Harvard and was most well known for his hand drawn physical relief maps.

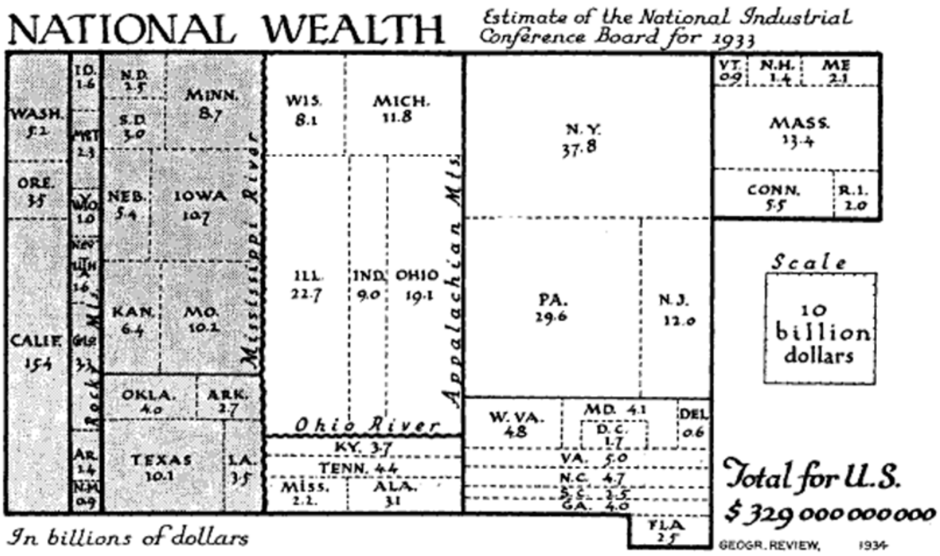

Figure 5.4 Rectangular statistical cartogram of the U.S. with sizes proportional to the distribution of national wealth among the states (Raisz 1934).

Cartograms can be *contiguous* or *non-contiguous*. Contiguous cartograms try to preserve the shape, orientation, and contiguity of the resized regions. Contiguity here refers to the specific connectivity and adjacencies of regions in the original map, and hence is a matter of preserving topology. Non-contiguous cartograms, on the other hand, do not have this constraint and thus can free regions from their adjacencies, appropriately resizing them while accurately maintaining their shapes.

A special type of non-contiguous cartograms includes the Dorling cartograms, which were named after its inventor. The "circular cartogram", as Dorling (1996, pp.32-36) named it, replaces major regions on a map with appropriately-sized circles and hence preserve neither shape nor topology. The circles are organized in a way that resemble the original arrangement, but moved from their initial locations so that they do not overlap. This way the full areas of the circles can be observed without juxtaposition. Dorling cartograms are also referenced in the related work of chapter 4 together with other techniques that abstract geographical positioning.

5.1.1 The communicative power of cartograms

There is rare empirical evidence that the more unfamiliar cartogram can be as effective and efficient as traditional statistical mapping techniques, such as choropleth maps and circle maps. For simple questions, cartograms seem almost as effective and efficient for inference as choropleth and circle maps (Kaspar et al. 2011). Furthermore, for qualitative themes, cartograms seem more effective and more preferred than traditional thematic maps (Sun and Li 2010). While cartograms appear unfamiliar or even provocative, they are also readable and understandable. The strength of cartograms lies in their unfamiliarity, making them more evocative and thus delivering stronger messages. Tao[34] (2010, p.243), addresses the unfamiliarity of cartograms by suggesting that, although techniques for cartograms are advancing, the construction and usage of cartograms is shackled by conventional attitudes. She urges for a change in attitudes:

> If both cartogram makers and readers stop thinking of cartograms in the same way they do traditional maps, some awkwardness and confusion may be avoided. The main reason for using cartograms is not only to tell people where they are, but also how they live: to focus on what happens in those places. Once the observer is freed from the shackles of conventional attitudes, the world may be observed from a brand new angle.
>
> —(Tao 2010, p.243)

Dorling's cartograms in general have the ability to "shock" readers by displaying a dramatic picture (Dorling 1996, p.4). For instance, when considering population representation, they provide a "more equitable representation", because they contrast with the more familiar representation of land areas, creating an unusual and provocative perspective on a common topic.

> Cartograms are produced for a variety of purposes. They can be used, like the London Underground map, to help people find their way. In atlases they are often used for their *ability to shock*; cartograms where area is drawn in proportion to the wealth of people living in each place show a *dramatic picture*. A major argument for the use of equal population cartograms in human geography is that they produce a more socially just form of mapping by giving people more *equitable representation* in an image of the world.
>
> —(Dorling 1996, p.4) italics added

Interpreting a cartogram is always dependent on the existence of a familiar mental image of the undistorted map. This contrast and this ability to "shock" make cartograms more evocative, enhancing the communicative power of a map and making them more dramatic. Cartograms have this effect because instead of using color, value, or shape to show information, they manipulate positions which are systematically rated as the most accurate and strong visual variable to show quantitative, ordinal, and nominal information (Bertin 2010, p.69; Mackinlay 1986; Cleveland and McGill 1984). Furthermore, their evocativeness can be a consequence of a well-known psychological phenomenon: the peak shift effect – the tendency to respond more to exaggerated versions of a reference stimuli than to the reference itself, provided that the exaggeration makes the reference more perceptible.

Ramachandran (1999) explains the peak shift effect as a principle in animal discrimination learning: if a rat is trained to discriminate a square from a specific rectangle and is rewarded for choosing the rectangle, it will respond more often to it after the training. If the rat then has to respond to a series of rectangles, including the training rectangle and the training square, it will respond more frequently to exaggerated rectangles that are longer and thinner. That is the effect where a shift in the peak of responses to certain stimuli can be observed not exactly to training stimulus, but to stimuli that exaggerates discriminating features of the training stimulus. The peak shift effect is then suggested to explain not only caricatures, but aesthetic preferences and many other aspects of art.

Lynn (2010) emphasizes the peak shift effect as a directional behavioral bias, adding that subjects respond most frequently to stimuli when they have never encountered it before, but nevertheless share a resemblance with training stimulus[35]. This effect can help explain why cartograms are described as shocking, intriguing, provocative, and dramatic. The hypothesis is that if we are trained with common maps and projections, learning the shape of each region and its connectivity, we would respond more often to novel and recognizable distortions of those familiar shapes. Cartograms distort familiar shapes in order to display statistical variables while trying to preserve recognizability, either through shape or topology. Cartograms embody stimuli, such as the unfamiliarity of the new shapes, that can explain a peak shift. Although unfamiliar, they are recognizable either by their distorted shape, their surrounding shapes, or a familiar context such as a world map or a known geographical region.

5.1.2 Computer cartograms

Waldo Tobler devised the first computer programs that produce cartograms. His first technique was specifically invented for population districting, and he describes his approach with reference to a rubber map:

> Suppose that one could stretch a geographical map so that areas containing many people would appear large, and areas containing few people would appear small. On a rubber map, for example, every person might be represented by an inked dot. We now imagine the rubber sheet to be stretched so that all the dots are at an equal distance from each other.
>
> —(Tobler 1973, p.215)

This method starts by dividing the map into a regular grid. Each grid cell has a population density value and the objective is to resize every cell so that they have the same population density. For each grid vertex, a displacement is computed that tries to minimize the density error of its adjacent cells. Before any vertex is adjusted, a topological test is performed to keep vertices from crossing the boundary of any cell. This process continues until it converges so that little or no improvement can be made. This first method for computing cartograms has extensive overall errors in the area distortion of the re-

[34] In her PhD dissertation supervised by Dorling.

gions and does not accurately depict the statistical variable in question. Experiments in applying the rubber map method by Tobler resulted in errors of 68% and 45% (Tobler 1973). Later on, Tobler (1986) introduced another method called "pseudo-cartograms" that are false cartograms or approximate cartograms, since they also introduce extensive errors. This method moves the latitude and longitude grid lines on the map in order to minimize the error. The method was devised as a convenient starting point for his rubber map algorithm as well as for cartographers creating contiguous cartograms by hand.

The last three decades brought substantial contributions and new techniques that reduce error and increase the efficiency in computer cartogram generation. For example, Dorling (1996, p.29) invented a cellular automata method inspired by Conway's Game of Life[36]. In his method, Dorling created a variant of the game that can grow cartograms. A grid of cells is created and each cell is assigned to one of the regions on the map, with its area proportional to its population. Regions that have too few cells will gain cells from neighboring regions that have too many, until each region has correctly obtained the desired amount of cells. This technique is accurate on resizing regions according to their population, but the contours of the regions are often lost, giving regions unrecognizable shapes (see figure 5.5).

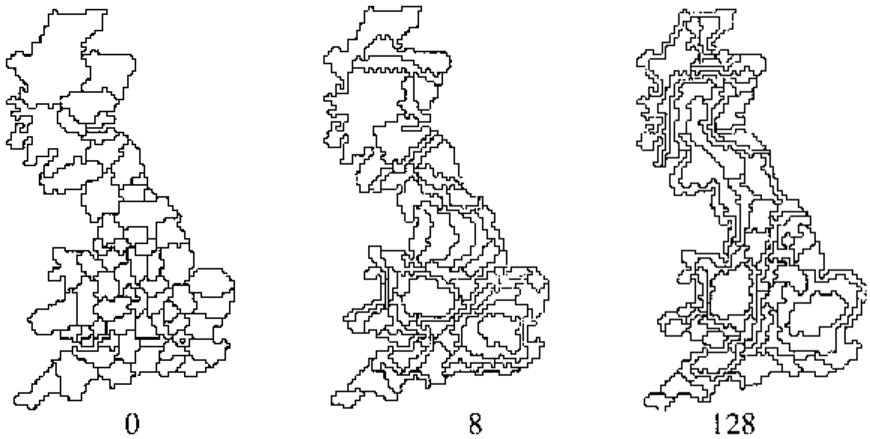

Figure 5.5 Contiguous population cartogram of Great Britain, after 0, 8, and 128 iterations, as simulated by Dorling (1996, p.29).

Other techniques for computer cartograms are inspired by physical mechanisms, such as springs that relax for regional shapes (Kocmoud 1998; House and Kocmoud 1998) or – perhaps the most popular for computer cartograms – a technique based on the diffusion process of elementary physics that allows a population to flow away from high-density areas into low-density neighboring areas until density is equalized in the whole map (Gastner and Newman 2005; Gastner and Newman 2004). An application of this technique can

[35] Lynn's experiment (Lynn 2010) to assert the peak shift effect involves a set of bumble bees (treatment set) being positively reinforced to hue A and punished when responding to hue B. In an ordered scale, A < B. When exposed to a multitude of hues between and beyond A and B, the bees respond more frequently to hues close to A and move away from hues close to B. These peaked stimuli may be considered exaggerations of A that share a sameness.

be seen on figure 5.6. These techniques usually present good results, bringing greater accuracy to computer cartograms construction, while preserving topology and maintaining shape recognizability. During the last ten years there has been an increase in the production of new computer cartogram models, with several patents being produced in 2005 (Fox and Daily 2005; Keim, North and Panse 2005a; Keim, North and Panse 2005b). Two recent techniques use self-organizing maps (Henriques et al. 2009) and flood simulation (Sagar 2014) in order to equalize densities on a map. This last one has the property of preserving the global boundaries of the map in question, but at the expense of the original form of each region. For a comprehensive overview of current computer cartogram construction techniques, refer to Kocmoud (1997, pp.15-25), Tao (2010, pp.45-55) and Henriques et al. (2009, pp.485-487).

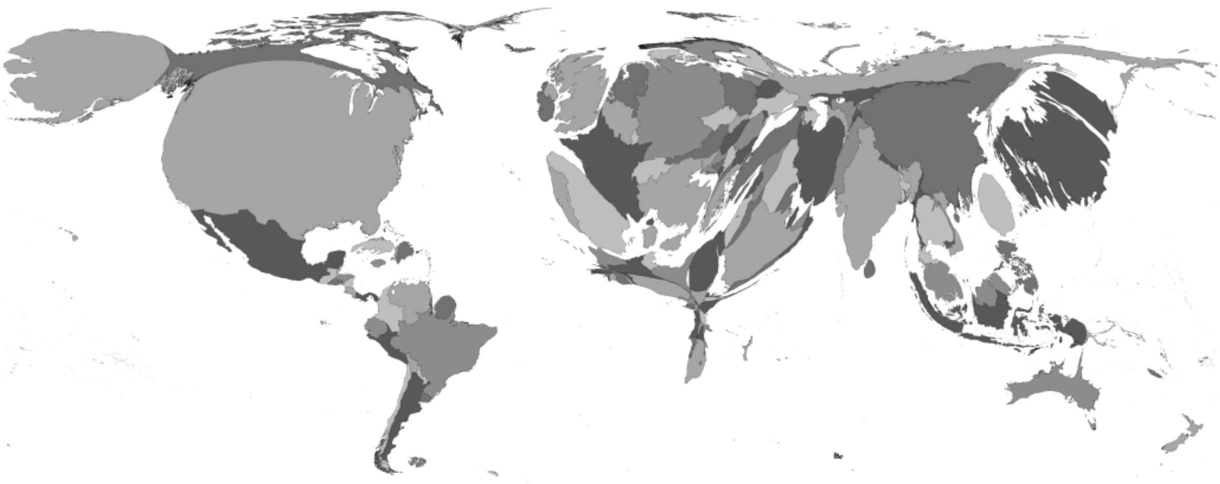

Figure 5.6 *Images of the social and economic world* by Newman[37] is a set of world cartograms that depict several demographic and economic indicators. This figure displays the world in terms of energy consumption and was created using the diffusion algorithm of Gastner and Newman (2005)

5.1.3 Cartograms for traffic visualization

A special type of cartograms is often used to display travel-time from a certain location. These cartograms have their roots in what Raisz (1948, p.262) calls "isochronic maps". The most common type of isochronic map connects points of equal time-distant with lines, called isolines or isochrones, showing the progress of travel in all directions in certain specified time-intervals. An early example of an isochronic map can be seen in figure 5.7.

---

[36] Created by John Conway, the Game of Life was first described by Martin Gardner (1970). The Game of Life is initialized on a square board where some cells are alive and others are dead. A new life is created when a cell is surrounded by exactly three live neighbors, and live cells that are surrounded by more than three or less than two other live cells, are deemed to die from overcrowding or exposure. Over several iterations these simple rules create incredibly complex patterns.

[37] Energy consumption (including oil) in Images of the social and economic world. Available at: http://wwwpersonal.umich.edu/~mejn/cartograms/ [Accessed November 1, 2014

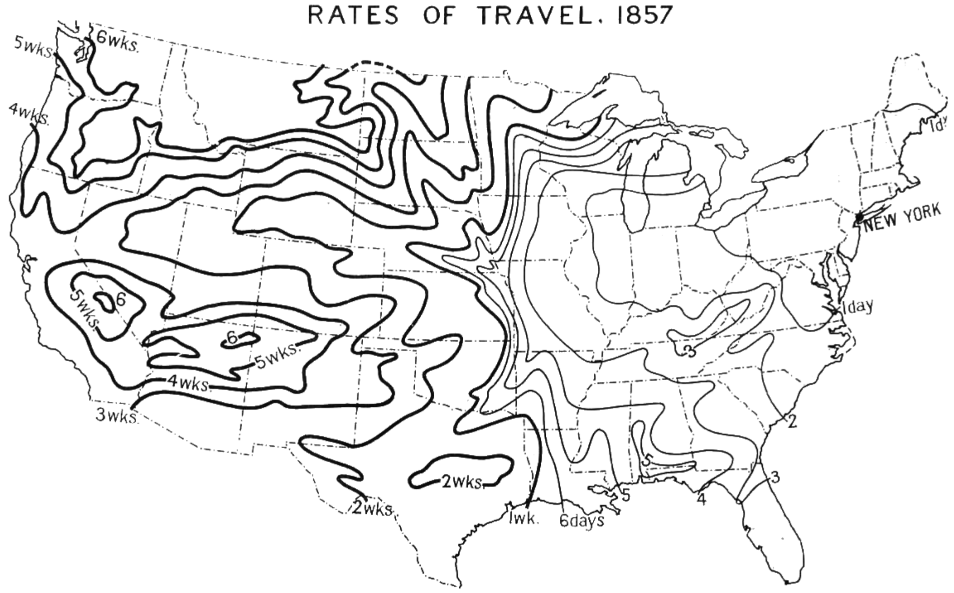

Figure 5.7 Isochronic map showing the rate of travel from New York City to the rest of the U.S. in 1857. Isolines represent several travel durations at time-intervals of days and then weeks (Paullin 1932, p.133 plate 138C).

Cartograms generated from isochronic maps are called "isochronic cartograms" or "distance cartograms". Isochronic cartograms distort space in order to position points with the same travel duration at the same distance. Therefore, they transform isolines into concentric circles based on a reference point. Isabel Meirelles (2013, p.167) points out that in these cartograms, time is used as a distance metaphor.

> We often use the metaphor of time as distance in our daily lives, such as when we provide temporal measures for giving directions. (…) There are many instances in which the measure provided by "how long it takes" replaces the spatial distances between places. *Isochrone lines* and *distance cartograms* are two common techniques using times distances.
>
> —(Meirelles 2013, p.167) italics added

As pointed by Lima[38], one of the first to experiment with an isochronic cartogram for travel durations was Sugiura (1973), a graphic designer that distorted Japan's map using Japan National Railways timetables in the 1970s (see figure 5.8). Later on, Clark (1977) devised and implemented a transformation algorithm computationally that distorts a portion of Seattle's road network according to the travel time from a certain location. Chen (2011) applied a time-distance transformation for Paris and Singapore using travel times from central locations. It is important to note that even if isochronic cartograms guarantee that the distances of the nodes of a travel network to a certain location

[38] Japan Travel-Time Map in website Visual Complexity. Available at: http://www.visualcomplexity.com/vc/project.cfm?id=204 [Accessed December 2014].

represent travel duration, they do not guarantee that the distances between the nodes in the network have any significance in travel time. Positioning nodes are only constrained by the distance to the reference node and hence other spatial properties of the transformation result from the specific technique being used. These techniques can take into consideration constraints such as the topology of the network or the overall shape of it.

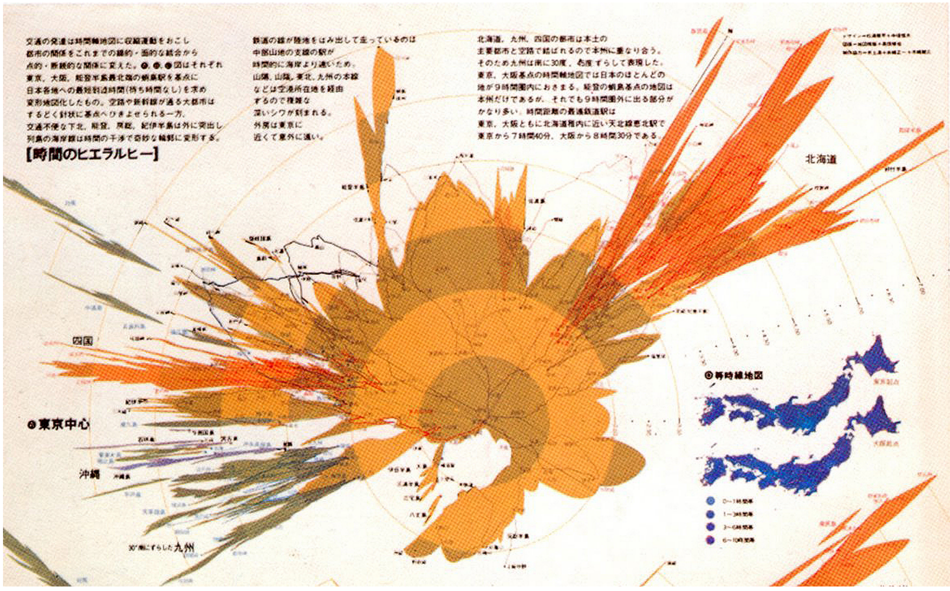

Figure 5.8 Detail of one of the first isochronic cartograms by Kohei Sugiura (1973), displaying a distorted map of Japan according to travel times from Tokyo.

Traffic information can refer, for instance, to travel times, congestion levels, or vehicles' velocities. Such information can be related with a specific point on a traffic network, but it can also be intrinsically related with specific road segments on the network. When having a traffic network in the form of a set of edges interconnected by a set of points, the problem of modeling a cartogram for traffic information should start by considering the network, its segments, and its interconnected nature. Offenhuber (2002) presented a contiguous cartogram that instead of manipulating areas or regions on a map, altered the lengths of road segments to represent a specific variable in traffic conditions. He used velocities in two sets of roads in Los Angeles to distort the road network according to time-distances. The roads are represented by "rubberbands", where the rest-length of each band varies with the temporal distance of the road. This strategy was able to deform the network and generate an edge-based cartogram that preserves topology. Nonetheless, it was only applied to two small portions of L.A.'s road network: Sunset area / Western Av. and Glendale intersection.

The "non-connective linear cartograms" of Wu and Hung (2010) took another approach to this problem. The authors disregard the connective nature of the network, and hence altering the lengths of the segments is neither constrained by their initial size nor by their connectivity. They exemplified this approach by representing traffic conditions for a portion of Salt Lake City's road network. Figure 5.9 shows the result of the approach by mapping travel speeds that increase with the lengths of the segments. This is counterintuitive to Meirelles' time-distance metaphor (Meirelles 2013, p.167) as well as Offenhuber's (2002) take on the subject. It makes sense that if travel speeds are higher, the perceived distances should be shorter and consequently the lengths of the roads as well. On the other hand, a counter argument is that having lengthier roads indeed emphasizes their perception, and this may well be the information that the reader wants to see first – speedier roads. Wu and Hung also present cartograms displaying travel times directly proportional to the length of road segments, and with such a map they are indirectly addressing the time-distance metaphor, having shorter segments for short travel durations and hence high velocities. Being an option of the cartographer, either an approach of having certain traffic values directly proportional or inversely proportional to the length of roads is legitimate, if clearly cited on the map. Wu and Hung's approach, albeit a first way to consider the lengths of a traffic network to display traffic information, has several limitations, such as being non-contiguous and hence not respecting the network's connectivity, being more prone to disrupting the recognizability of the map. Furthermore, as they notice, their cartogram is not suitable for more irregular topologies: "Cartograms should be used with caution. Non-connective linear cartograms do not suit street networks with long streets far away from each other or with irregular patterns" (Wu and Hung 2010, p.47).

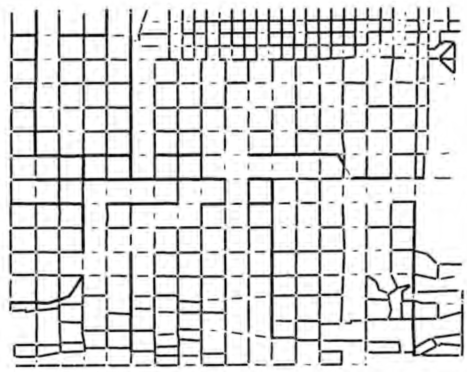

Figure 5.9 The linear non-connective cartogram by Wu and Hung (2010) displaying road segments that are proportional to traffic velocity for a section of Salt Lake City's road network.

5.2 Model for edge-based cartograms

The perspective of Keim (2004) and Tobler (2004) on cartograms implies that every cartogram is a *value-by-area*. Here, cartograms are more broadly viewed as a cartographic technique that distorts the spatial coordinates on a map according to the geographic distribution of a statistical variable.

Keim (2004) formally defines the cartogram problem for value-by-area cartograms through a set of connected polygons (a polygonal mesh) and the desired values for the area of each polygon. The objective is to generate a contiguous cartogram, and therefore, the output should also be a polygonal mesh, with ideally three constraints:

- the vertices and the edges that interconnect the vertices in the mesh should remain the same – *topology preservation*;
- the form of each polygon should remain the same – *shape preservation*;
- each polygon will be resized according to a variable – *area resizing*.

Keim then moves to prove the impossibility of this ideal solution as there are polygonal meshes and resizing parameters that are unsolvable for general cases. As Clark (1977) points out, when referring to isochronic cartograms, the time-distance transformation of a network that maintains the topological connectivity of the nodes can be done in general cases only if the transformation is done with respect to the travel time from a point within the map.

In order to implement feasible variants of the problem, the solution is to "relax" some of the constraints. Usually for contiguous cartograms, topology is considered the most important property to maintain (Keim et al. 2004, p.97) so relaxation is allowed for shape preservation and area resizing. That is to say that the cartogram problem can be solved with a certain error in the ideal solution, being when shapes are not rigorously preserved nor the area resizing accurately done.

5.2.1 Edge-based cartograms

Keim (2004) defines a cartogram as a transformation over a set of polygons, preserving their topology, and trying to preserve shape and resize areas. Here, the cartogram problem is defined based on a set of polygonal curves[39] over a finite set of points that represent a road network for traffic visualization. The intent is to depict a statistical variable through the lengths of the paths and not through the size of shapes as in value-by-area cartograms. These types of cartograms are referred to herein as value-by-length or edge-based cartograms. The typical constraints applied for the construction of value-by-area cartograms have to be redefined for edge-based cartograms. For example, the topology preservation constraint is seen through the preservation of the connections among the points in the paths, and the area resizing constraint is expressed through a length resizing constraint. Such as in value–by-area cartograms, the general case of the cartogram problem is not solvable. For example, considering a triangle of three edges {L1, L3, L3} which are resized according to a statistical variable to new lengths {L1', L2', L3'}, then if topology has to be maintained, all of the edges are constrained to sizes between zero and the sum of the other two edges' lengths. This makes it impossible for the triangle to represent every triplet of lengths (for example the lengths {4, 1, 8} cannot be represented as a triangle).

As previously discussed in the context of cartograms, map recognizability is paramount. This is assured in value-by-area cartograms with its topology and shape constraints. Edge-based cartogram paths may not be topologically connected, and thus only ensuring the preservation of shape (for example by means of preserving the angles between pairs of segments of the path) is not enough. Consider figure 5.10, with an initial arrangement of two non-connected paths. If the red path is resized while preserving its shape, it might only create new arrangements that greatly impair the map's recognizability. For instance, new intersections can be created or the relative positions of the paths to each other can be heavily altered. This problem can be avoided if an additional constraint is inserted that scales the relative distances of the paths to each other, but in a way that is related with the resizing of the paths themselves. The purpose here is to distort the map as uniformly as possible, trying to preserve the map's initial features.

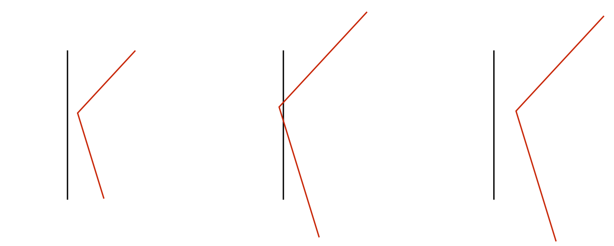

Figure 5.10 Diagram illustrating the problems of resizing a path without considering the relative distances to its neighboring paths. As showed, resizing the red path can create new intersections that impair the map's recognizability if relative distances are not considered.

Having these aspects in mind, the constraints for the cartogram problem in the context of edge-based cartograms are:

— the vertices and the edges that interconnect the vertices of the paths should remain the same, such as in value-by-area cartograms – *topology preservation*;
— the shape of each path should remain the same, by preserving, for example, the angles between its segments – *shape preservation*;
— each path or segment on a path will be resized according to a variable – *length resizing*;
— the distances of a path to its neighboring paths should be scaled proportionally – *distance scaling*.

[39] The terms polygonal curve, polygonal chain, polyline, and path are used interchangeably in this text.

These constraints are considered in the context of this work to contribute for map recognizability when distorting an edge-based cartogram to represent a statistical variable by resizing its edges while preserving topology, i.e., constructing a contiguous edge-based cartogram.

5.2.2 Model

As previously mentioned, the edge-based cartogram problem is not solvable for general cases. Nevertheless, compromises can be attained by relaxing some of the constraints except for topology. The model presented here is a mesh of springs that when simulated presents an approximated solution for the cartogram problem. A spring is an elastic device that, when compressed or stretched in relation to its rest length, exerts a force proportional to its change in length. The idea is first to replace the edges of the cartograms by springs, and by altering their rest lengths, simulate how the system would evolve in order to obtain a good compromise for a solution. Consider once again the triangle that has to be resized by lengths {4, 1, 8}. One solution is to relax the length, resizing the constraint along with the shape preservation constraint, as illustrated in figure 5.11. In this figure, by setting the rest lengths of the springs to {4, 1, 8}, and by simulating the system for a few steps, it stabilizes then the resized edges have effective lengths of approximately {≈5, ≈2, ≈7}.

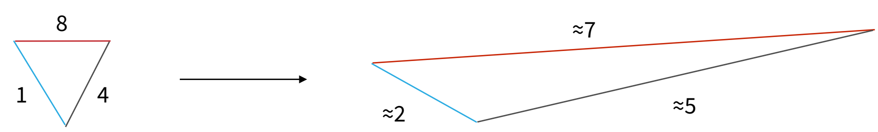

Figure 5.11 Illustration of a compromise when resizing the edges of a triangle to {4, 1, 8} while maintaining topology.

Spring-based approaches have already been used to build contiguous area cartograms (House and Kocmoud 1998), using springs to resize regions in a way that maintains shape whilst preserving topology. Due to the nature of an edge-based cartogram, this solution cannot be directly applied since it was devised for regions and polygons. Their approach does not directly simulate a mesh of interconnected springs, instead they selectively apply forces with a spring-like behavior to certain vertices of the polygons in order to meet a set of constraints. In our case, the subjacent mesh that represents the map is formed by springs, i.e., in the case of polygons, the edges would be made of springs. Furthermore, in their approach the spring constraints are dynamically adjusted to the error of the map relative to the constraints of area, size, and topology – if these errors are large, the spring constants will be increased in order to more effectively counteract them. The implementation presented here takes a different approach by empirically selecting an extensive set of spring constraints and testing them in order to find the smallest error. Moreover, the subjacent model that is described ahead is much simpler and creates a mesh of springs that is left unchanged through the simulation (except for the springs' rest lengths), adapting to different lengths of the edges and presenting an approximate solution for the edge-based cartogram problem.

The model proposed here has three types of springs, each one with a different role in adhering to each constraint in the edge-based cartogram problem. A possible arrangement for these three types of springs for a path is illustrated in figure 5.12.

— *Backbone springs* form the edges on the map. They are used to preserve topology since each pair of vertices that is connected in the original map will stay connected in the resulting cartogram. Furthermore, they play a major role in the length-resizing constraint, because by altering their rest lengths their corresponding paths will be resized accordingly if possible.
— *Inner springs* of each path in the map. Additional to the backbone springs, each path has its own skeleton of springs that tightly connects the vertices on the path. Those inner springs are used for the shape-preservation constraint since, if it is scaled by the same ratio, the path will be resized uniformly, preserving its shape and its internal angles.
— *Connective springs* connect vertices of one path to neighboring vertices of a neighboring path. These connective springs are used for the distance-preservation constraint, since they would be able to adequately scale the distances between paths and thus preserve the overall shape of the map and its recognizability.

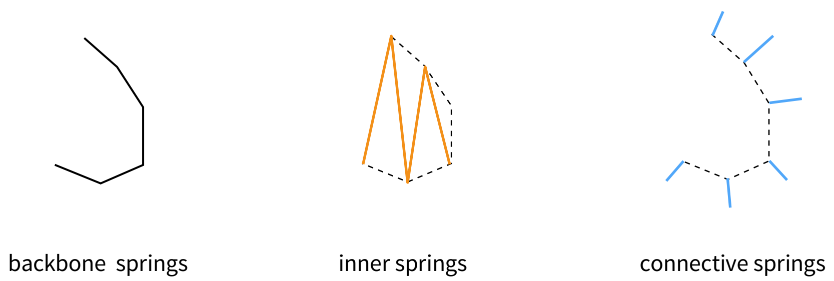

Figure 5.12 A possible arrangement of backbone, inner and connective springs for a path. Backbone springs form the original path, inner springs tightly connect it, and connective springs connect the path to other neighboring paths.

In the context of edge-based cartograms, it is assumed that the initial map is described through a set of paths. A path connects sequentially with edges of a list of vertices. Paths can share vertices, but each edge is unique in the map as it connects two different vertices only once. As suggested before, the backbone springs make the edges of the paths themselves. In order to find the inner and the connective springs, the Delaunay triangulation[40]. The Delaunay triangulation of a set of points links those points in order to divide the region into triangular tiles in such a way that no point is inside the circumcircle of any tile (see figure 5.13). This method tends to avoid thin triangles by maximizing the minimum angle between any two edges.

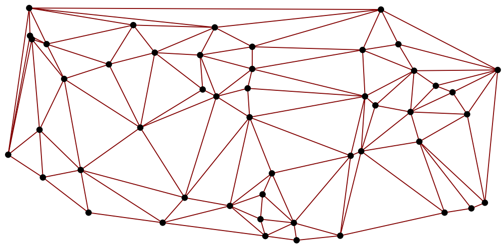

Figure 5.13 Example of a Delaunay triangulation of a set of points.

The inner springs are determined for each path by computing the Delaunay triangulation of its vertices as if they were unlinked. It is important to notice that the result of this triangulation can include edges of the path and therefore a backbone spring. With this, the inner springs of a path are the edges of the Delaunay triangulation of its vertices that are not already part of the path. For example, referring to figure 5.12, the inner springs are only the orange edges in spite of the Delaunay triangulation of that path also including the dashed edges. A similar exclusion procedure is done to extract the connective springs, but this time instead of considering each path individually, all the vertices of the map are taken into account as a set of unconnected points. By computing again the Delaunay triangulation of these points, the connective springs are the edges that are not a backbone or an inner spring of any path.

To more formally recapitulate on these aspects, consider once again the following:

— Each path describes a set a backbone springs B_i that connects the corresponding set of vertices V_i.
— If the Delaunay triangulation of a set of points A is $\mathcal{DG}(A)$, then the set of inner springs for the path i is $I_i = \mathcal{DG}(V_i) \setminus B_i$.
— If all the vertices of all the paths make the set $\mathcal{V}$, all the backbone springs of every path are the superset $\mathcal{B}$ and all the inner springs of every path are the superset $\mathcal{I}$, then the connective springs of the system are given by $\mathcal{DG}(\mathcal{V}) \setminus (\mathcal{B} \cup \mathcal{I})$.

5.2.3 Simulation

A spring is an elastic device, meaning that when compressed or stretched in relation to its rest length, it exerts a force proportional to its change in length. This coefficient of proportionality is a characteristic of the spring and can be interpreted as its strength. The system is simulated through a set of particles of unitary mass interconnected with springs. For each spring that connects a pair of particles a and b, the constraint applied is described in pseudo-code below (a and b are the particles' position vectors; vectors are in red).

```
ConstraintParticles(spring):
    delta = b - a
    dist = magnitude(delta)
    diff = 1 - spring.restLength / dist
    shift = delta * 0.5 * diff * spring.strength
    a = a - shift
    b = b + shift
```

This constraint pushes or pulls the particles proportionally to the spring's change in length and its strength. When simulating a system of many interconnected springs, it is not trivial to solve the corresponding system of differential equations. In such complex systems, using simple numerical integration techniques usually leads to instability (Jakobsen 2001). Ideally, all the constraints for each pair of particles connected with a spring should be solved at once, which involves solving a system of equations. The simulation of this system, however, proceeds indirectly through local iteration. The constraints are applied to each pair of particles connected by a spring, one after each other. Such is done for each simulation step, translating in the following pseudo-code.

```
SimulateStep:
    for i = 1 to NUM_ITERATIONS:
        for each spring:
            ConstraintParticles(spring)
```

This method is called relaxation (Press et al. 2007, p.964) and it works by consecutively satisfying several local constraints and then repeating itself for a determined number of iterations, converging to a global configuration that satisfies all the constraints at the same time. The simulation of the spring-based cartograms uses a number of 50 iterations for relaxation, which was observed as appropriate for the level of complexity of the system.

5.2.4 Exciting springs: ratios and lengths

The springs' mesh reflects a road network map. Initially, the springs extracted from this network have a rest length equal to the distance of the nodes that it connects on the map. Therefore the `diff` factor presented in the pseudo code is 0 (`diff = 1-spring.restLength/dist`). This means that the constraints are met and the map is left unaltered during the simulation. In order to create a cartogram, the spring's lengths have to be modified in order to reflect a certain variable for the corresponding edge of the map. To do this, a *ratio* of the ini-

[40] The concept behind the Delaunay triangulation was first introduced by Boris Delaunay in 1934 (Delaunay 1934) for n-dimensional spaces by using the concept of a sphere that moves between points, shrinking and expanding at will, but obeying to the condition of being empty i.e. not containing in its interior any of the points. In order to determine the Delaunay graph of a set of points the QuickHull (Barber et al. 1996) algorithm was used, which has a time complexity of *O(n log(n))*.

tial rest length of each spring is computed in order to traduce the spring's new length. The spring's length is altered accordingly, being the only type of constraint being modified – the spring's strengths remain unaltered. With this, it can be said that the springs are excited since the spring's particles will change their positions in order to accommodate the new lengths, adapting to the new constraints.

The springs' strengths are important parameters in the system simulation. They influence how quickly the springs adapt to new lengths as well as establishing how a spring's constraints can be solved at the expense of others. In this system there are only three different strengths for each type of spring: k_B for the backbone springs, k_I for the inner springs, and k_C for the connective springs. Intuitively, this enables prioritization of each of the cartogram's constraints (resizing, shape preservation, and distance preservation) through the springs' strengths in detriment of others. A delicate balance among the three forces has to be experimentally found for each specific implementation of the edge-based cartogram: a good balance would mean that the cartogram constraints are met in the best way possible.

The computed ratios vary from spring to spring, depending on the road that they pertain as well as the type of spring they are. Each road has a specific computed ratio that represents the statistical variable being shown. The objective is to scale the road uniformly, according to this ratio, so that it will shrink if the ratio is below 1 and distend if it is above 1. This means that the applied ratios are the same for any of the backbone and inner springs of the same road, resulting in a uniform scaling that tends to preserve the road's shape. The connective springs yield a slightly different strategy. Each connective springs connects two different roads. Their objective is to make space to accommodate the different roads' scalings, ideally resulting in a map that distends globally if most of the roads are scaled up and compress if the roads are downsized. Since a connective spring connects two different roads with different scaling ratios, then it is proposed that the scaling ratio of that spring is the average of the two road's scaling ratios that it connects. This approach works globally in order to make space to accommodate the new ideal lengths of the roads, yielding the following code:

```
for each simulation step:
    for each road:
        road.ratio = computeRatio(road)
        for each spring in
            road.backboneSprings and road.innerSprings:
            spring.restLength =
                spring.initialRestLength * road.ratio
    for each spring in connectiveSprings:
        ratio = 0.5 * spring.road1.ratio + spring.road2.ratio
        spring.restLength = spring.initialRestLength * ratio
    SimulateStep
```

The system can quickly adapt to new lengths of the edges. For the adopted interactions in the relaxation, the simulation displays smooth animated transitions between different sets of rest lengths. Notice that the planar orientation of the springs is implicitly determined by the positions of the pair of particles that it connects. The nature and specific way that the ratios for the springs are determined and related with the data and the statistical variable are visualized and described in the next section.

5.3 Data from Lisbon

The dataset[41] contains information on GPS traces of 1,534 vehicles in Lisbon's municipal area. The total number of GPS registries is around two million, spanning through the 31 days of October 2009. Each GPS registry has both a latitude and longitude point in time, with information about speed in kilometers per hour. Looking at the raw number of registries through time in graph 5.1, several patterns can be perceived. The decrease of business activities during the weekend as well as the smaller influx of vehicles to the city means that less registries are available during weekends. Saturdays have more reported registries than Sundays, and the weekdays follow the same general profile: morning increase of traffic; slight reduction around lunch time; a new smaller increase in late afternoon, and an abrupt decrease in the evening. This is the main story to convey from the data, portraying the rhythm of the city: when it wakes up, when it is more frenetic, and when it goes to sleep. This pattern can be better observed when superimposing every day in graph 5.2.

Graph 5.1 Number of registries throughout the days. The dataset, ranging 1-31st October 2009. Monday 5th does not follow the profile of other Mondays since it was a national holiday.

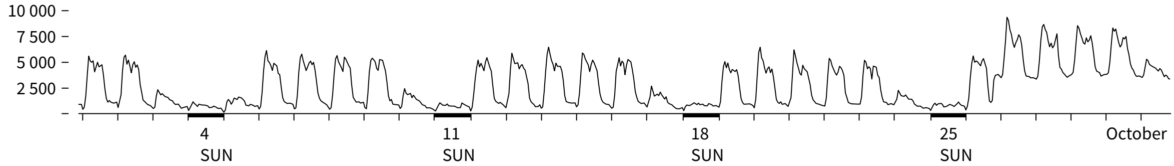

Graph 5.2 Number of registries by hour superimposed: ▪weekends, ▪weekdays before the 26th, ▪weekdays after the 26th. The days after the 26th have a higher number of GPS registries. This was discovered to be from vehicles reporting their current GPS position uninterruptedly, even at night when they are stationary.

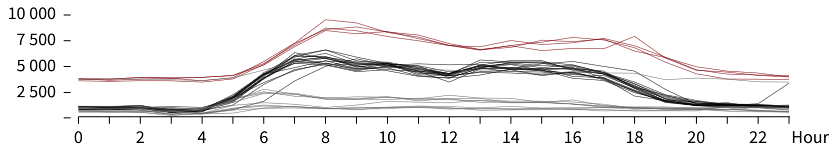

The first approach to visualize this dataset illustrates the physical world in a logic where visual complexity is present, but can be deciphered in a story. With this, a set of detailed representations of data events work together to generate a general representation of the data, enabling the portrayal of a story without losing the visual richness that can be extracted from a dataset. The physical nature of the dataset refers to 1,534 vehicles that report their GPS positions during time, hence vehicles that are moving or stationary. Graph 5.3 shows the distribution of the number of vehicles per hour across the dataset's timespan. As shown, if the dataset was to be illustrated by displaying moving vehicles in a linear fashion for the whole month of October, it would account for a repeating story where vehicles appear from minimum quantities of ≈50 to peaks of ≈three hundred. This accounts for an average of 217 vehicles per hour, that is insufficient to either have statistical relevance or to create visual richness from the dataset. A way to address this is to display every day at the same time, creating higher data density with a more visually complex portrayal, while visually enhancing patterns of data features. For instance, graph 5.4 shows the number of vehicles combined per hour independently of the day. Combining the days results in more data density, with the presence of vehicles oscillating between ≈four hundred and ≈950, with an average of 691 vehicles present per hour.

Graph 5.3 Number of vehicles throughout the days. Minimums of ≈50 vehicles are usually set at night, and peaks of ≈300 at the morning rush hour, averaging at 217 vehicles per hour. Notice that in the last week of the data, the minimum does not go below ≈two hundred, while the peaks remain coherent with the previous weeks. This is due to what was mentioned on graph 5.2, where in the last week vehicles keep registering their positions uninterruptedly, even at night.

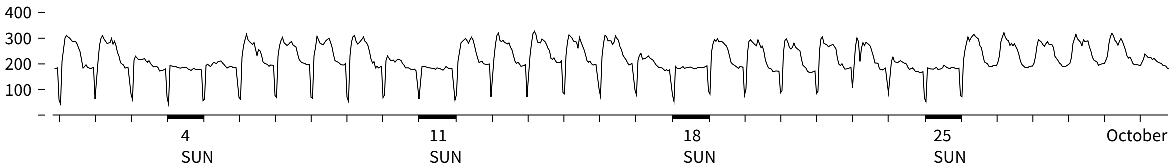

Graph 5.4 Number of vehicles by hour of the day. This aggregation permits the creation of greater data density, putting the number of vehicles in the range [≈400, ≈950] with an average of 691 vehicles per hour.

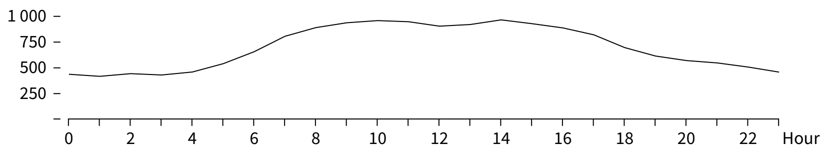

Another important feature in the dataset are the vehicles' velocities. Graph 5.5 shows the average speeds per hour across the data-span. The velocities' profiles are more irregular than the profiles of the number of registries. The last week in graph 5.5 should be observed since it presents the most regular patterns that vary between approximately zero and maximum values that are lower than the ones of the previous weeks. This is due to the fact that vehicles register their GPS position uninterruptedly, even when they are stationary, and even at night. This softens the average velocity profile for that week, having smoother transitions between night and day, and with close to zero velocities during night, showing that most of the vehicles are stationary. Again, it should be noted that the number of vehicles circulating in the city did not increase in that week, only the frequency of GPS reporting. As can be seen in graph 5.6, the weekends are the most irregular, certainly due to the lower data density on those days. When superimposing the velocities' profiles for each weekday

[41] The dataset is provided by FROTCOM (http://www.frotcom.com) a company that provides solutions for managing vehicle's fleets. The information concerning the specific types of vehicle fleets of the dataset were not disclosed, but probably refers to goods' distribution trucks given the company's profile. The data was accessed in the context of the City Motion Project – MIT Portugal.

(graph 5.7), the weekdays of the first weeks generally show a peak in velocity around 04:00 and a local minimum around 07:00, which is probably due to the morning rush hour. The profiles of the last week have lower average velocities. They do not display a peak around 04:00, and the most compelling aspect is the local minimum around 13:00, here hypothesized as due to more traffic in the city at lunchtime.

Graph 5.5 Average velocities throughout the days. The last week has a more regular profile due to greater data density.

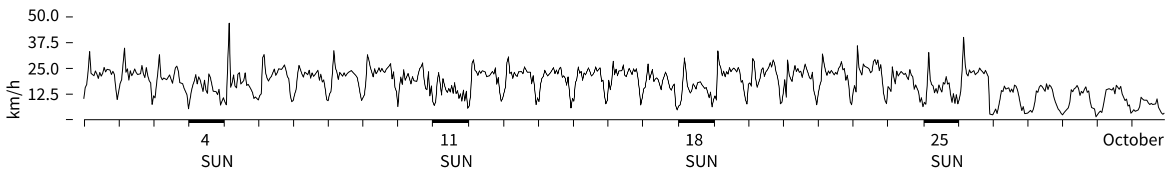

Graph 5.6 Weekends and holidays average velocities by hour of the day.

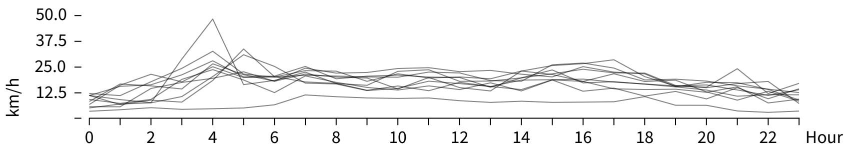

Graph 5.7 Weekdays average velocities by hour of the day: ▪weekdays after the 26th (except the holiday), ▪weekdays before the 26th.

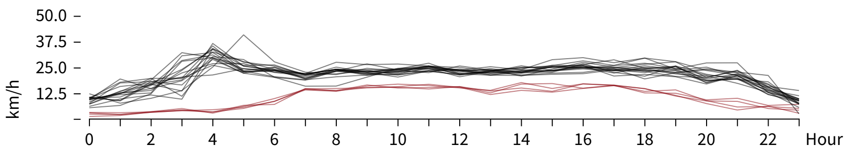

The profile in graph 5.8 is the average of all the weekday profiles (except the holiday). This average was computed in order to extract a common narrative through the dataset. The resulting profile has a local minimum for the morning rush hour, but does not display another local minimum for the afternoon rush hour when people are leaving the city. This is probably due to the aggregation time interval of one hour being too large. Nonetheless, the question as to whether the afternoon rush hour is noticeable or not in the data remains unanswered.

The next section describes the first visualization approaches to this data in order to unveil a common story in this dataset. Due to the disparities of profiles in velocities and data density between weekdays and weekends, the weekdays will be excluded from the data in order to work with a more robust and homogenous collection of traces. Such filtering left approximately 1.8 million GPS traces of the initial ≈two million.

Graph 5.8 Overall average of velocities by hour. This profile misses one key aspect of the expected behavior of the city: the afternoon rush hour. The main reason would be granularity of time aggregation and is addressed in the next section.

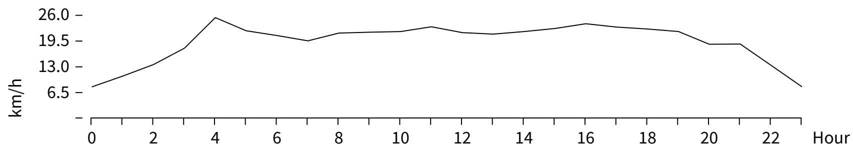

5.4 Visualization of trajectories

A key feature of Lisbon's dataset is the positions of vehicles along time. Displaying this geographical component would tell a more compelling story, depicting the traffic flow in concrete areas and arteries of the city with both quantity and vehicle speeds. A standard technique for display movement of discreet entities is the visualization of trajectories (N. Andrienko and G. Andrienko 2013), where each physical entity leaves a visual trail of its previous positions[42]. This can be done through either static or animated maps. In the first visualization approach to Lisbon's dataset, recent trajectories of the vehicles are displayed in an animated map from 0:00 to 23:59, meaning that the trajectories from different days are superimposed on the visualization canvas. This superimposition visually emphasizes the main patterns in data, circumventing the low data density problem mentioned in the previous section. The result in an effective device to tell the story of the daily rhythm of traffic in Lisbon.

5.4.1 Trajectories

The registries are displayed per time of the day, regardless of the day, resulting in an animation ranging from 0:00 to 23:59, which displays a visual superimposition that iterates one by one minute per frame. For each minute, the GPS traces of the previous 30 minutes are also displayed. This is to say that a moving visual average of the data is displayed for each minute t in the range $[t-30, t]$ minutes, increasing data density per frame and providing smoother transitions between consecutive frames. Each vehicle present at time t is represented by a white dot. The positions of the vehicle in the previous 30 minutes are also displayed by connecting them with a path. These paths display a glimpse of the vehicle's recent trajectory, and when superimposed with other vehicles, they are able to draw a global map of trajectories for traffic in Lisbon. The registries *(latitude, longitude)* were mapped using the Mercator Projection[43].

5.4.2 Metaphorical color mapping

The trajectories have a color signifying the velocity of the vehicle that ranges from red to cyan. It is well known, according to Bertin's graphic semiology (Bertin 2010, p.69), that color is not recommended for quantitative or ordinal information. Nonetheless, even if Mackinlay (1986) agrees with Bertin that color hue is one of the worst visual variables to show quantity, he presents color hue as the fourth best visual variable for ordinal information, after position, density, and color saturation. Ware (2004, p.132) notes that color should be used with caution for ordinal pseudo-color sequences, stating that spectrum sequences can be effective if the red, yellow, green, and blue regions can match significant data classes, which for this visualization would be "almost stopped", "slow traffic", "regular traffic", and "good traffic" respectively. In the scale adopted, blues are never reached as they are perceptually much darker than pure reds, yellows, greens, and cyans. The red-cyan color scale is justifiably given a metaphorical embedding in this specific context. The metaphor refers to traffic signals such as red is "an almost stopped" for velocities below 10 km/h, green is a "you are good to go" 50km/h, and for velocities above, bluish greens and cyans are used. The segments of the trajectories can have different colors, depending on the velocities of the vehicle at each previous position. Figure 5.14 illustrates this and the scale used to color the trajectories.

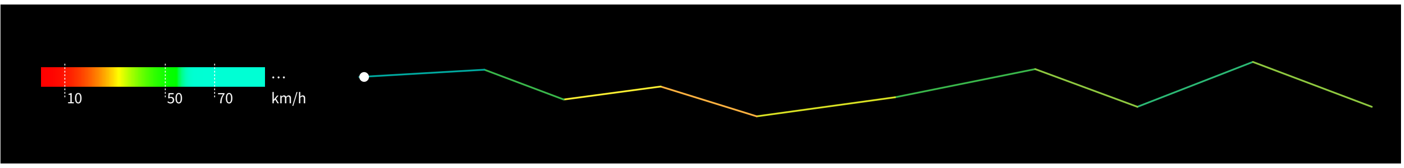

Figure 5.14 A vehicle represented by a white dot in time *t*, with a path connecting its previous positions for [*t*–30, *t*] minutes. The colors of the path correspond to the velocity of the vehicle in its previous positions. The colors are mapped according to the depicted scale. The path can have varying colors through its length.

5.4.3 Transparency

Another important aspect is that each trajectory is drawn with a certain amount of transparency, creating a more seamless composition, unveiling all the complexity of the dataset and enabling the extraction of visual estimates of velocities on clustered trajectories over main roads. As mentioned by Andrienko (2013), reducing the opacity of graphical items can contribute to decreased cluttering in the visualization artifact. This way, the trajectories' paths are not completely opaque, which permits the observation of superimposed paths to a certain extent. Indeed, if the opacity factor for the paths is constant, what can be observed is that major highways with more traffic volume, but with better traffic speeds than areas with narrow streets, get drastically emphasized. Having such emphasis for major arteries is important since much of the traffic flows through them. Nevertheless, providing an adequate emphasis for areas that cannot accommodate much traffic but that have slower flows is also important, since they can refer to denser areas of the city that can easily get congested. In order to tackle this, the opacity mapping of the paths is dependent on the velocities such as when velocities are lower, the opacity is higher. This way, problematic areas of the city in terms of congestion can be better emphasized. For the specific scale for the opacity mapping refer to figure 5.15.

---

[42] The visualization of trajectories technique applied to city traffic was popularized to broad audiences by BBC's "Britain From Above" in 2008.

[43] The Mercator projection maps the earth's sphere on a plane as if it was projected on a cylinder. It was invented by Gerardus Mercator in 1569 (Snyder 1987), and it is still the most often used map projection in the world. This projection considerably distorts the size and shapes of large objects that are closer to the poles. Given the large scale of our map and the distance of Lisbon from the North Pole, the distortions are negligible.

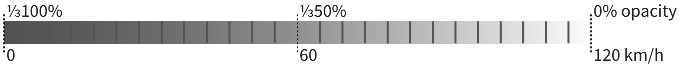

Figure 5.15 Linear opacity mapping of the path's segments according to the velocities. Slower velocities are more opaque, and higher velocities are more transparent. The maximum opacity possible is ⅓100% when the velocity is zero and 0% for the velocity of 120 km/h and beyond.

5.4.4 Results

When observing the first results of the visualization a representation of the daily rhythm of the city can be observed. The city sleeps at night with few vehicles circulating that are mainly in the periphery with varying speeds. Clusters of standing vehicles can be mainly noticed in the periphery of the city, represented by groups of white dots (see figure 5.16). As previously mentioned, the presence of such vehicles during night time is mainly a consequence of the last week of data. The city has then its arteries illuminated by numerous flowing vehicles at various speeds (see figure 5.17).

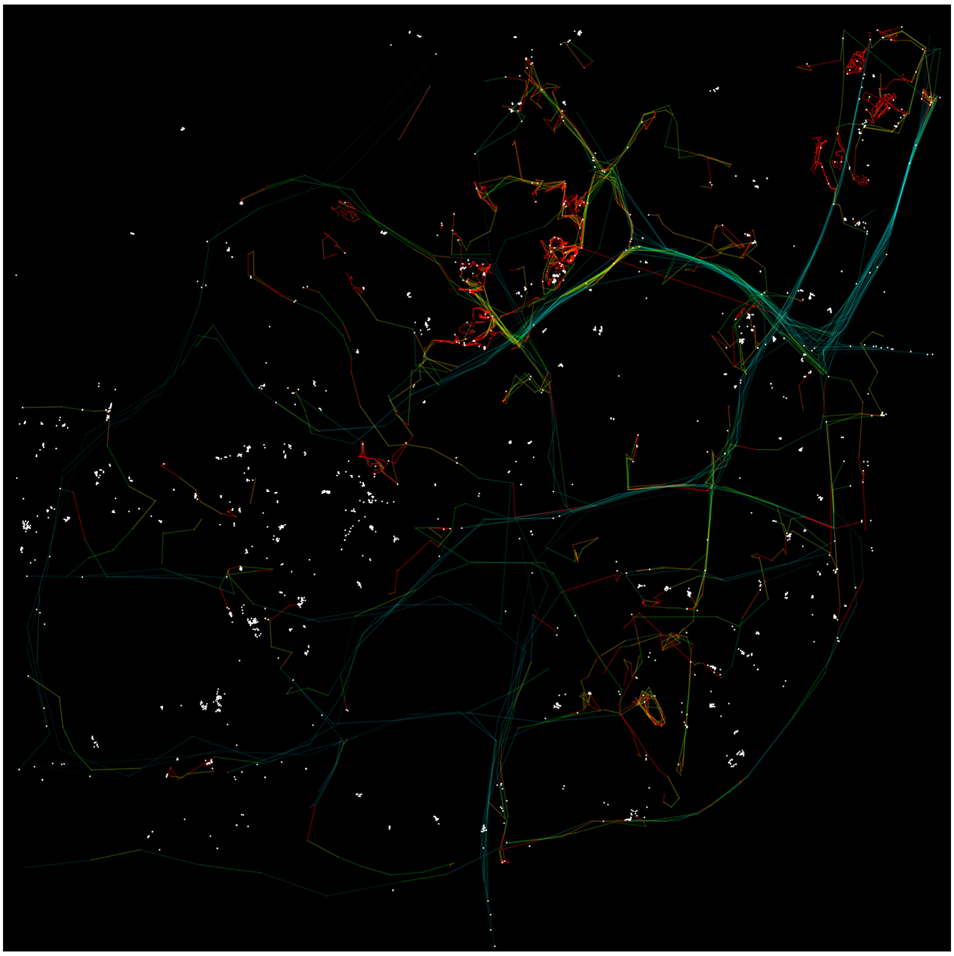

Figure 5.16 Visualization of trajectories at 00:30. The city is at sleep with few vehicles circulating. Several clusters of standing vehicles can be seen mainly on the periphery of Lisbon by identifying clusters of white dots.

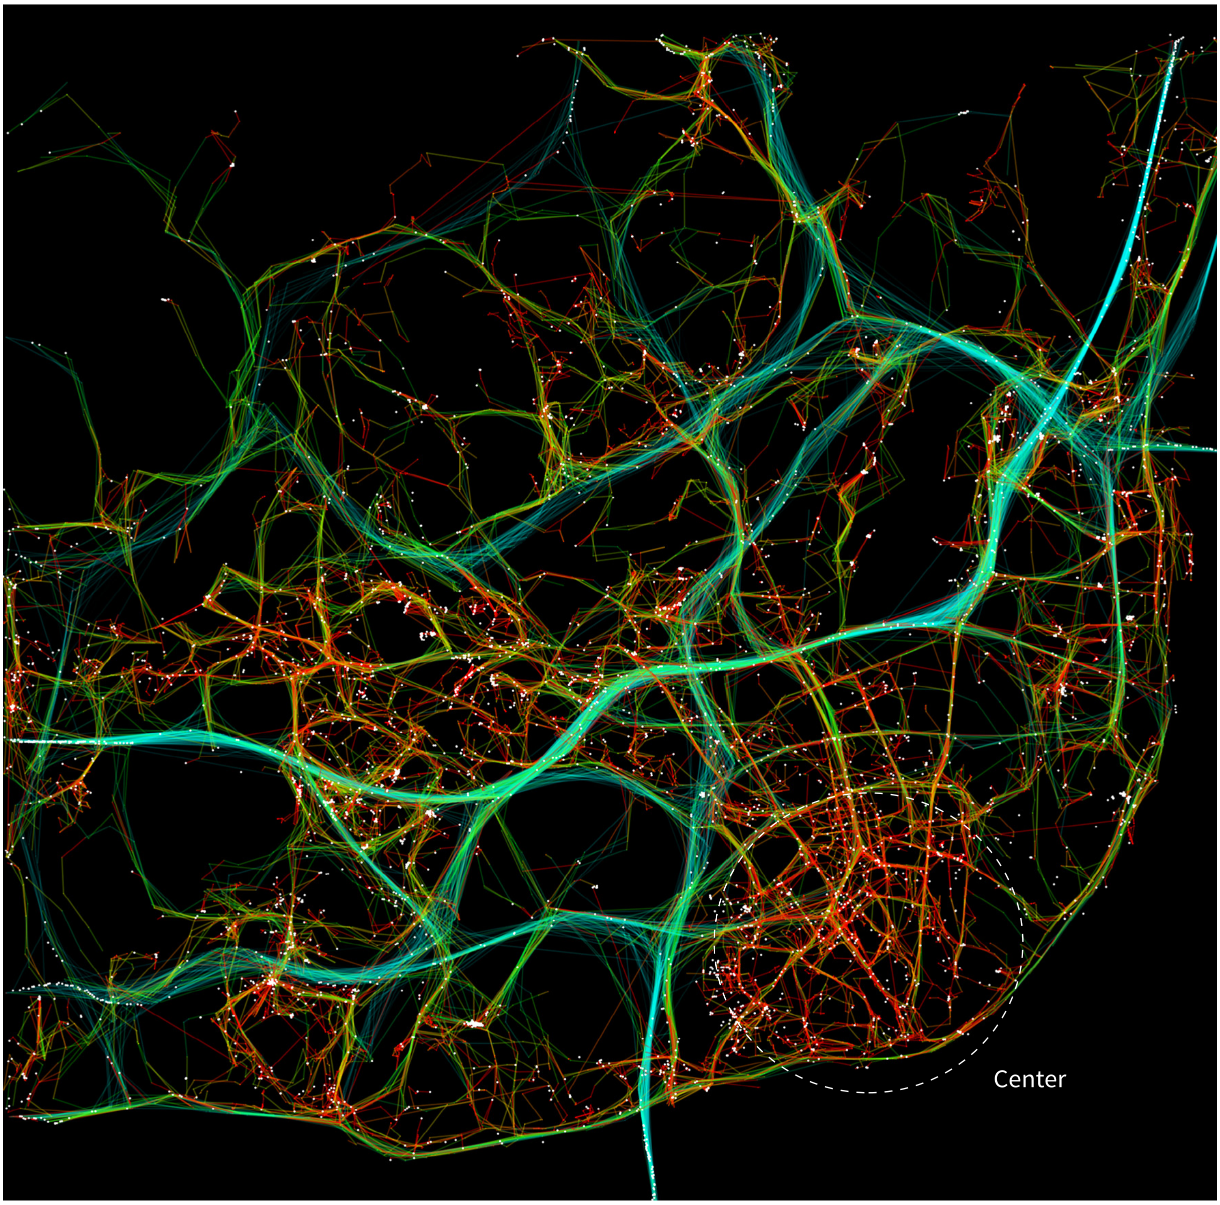

Figure 5.17 Visualization of trajectories at 11:00. The city has its arteries illuminated by numerous trajectories of vehicles at various speeds.

With this visualization, trajectories can be visually grouped, providing a glimpse of traffic volume and speed: thicker clusters of trajectories indicate higher volumes, with their average color indicating velocity. When comparing this with the road map of Lisbon (figure imgs/intro-lisbon/mapa-lisboa.svg), these thick clusters identify Lisbon's major highways, namely 2ª circular, IP7, A5 and A36. Furthermore, the result also enables us to distinguish areas of the city with different velocity profiles. For example, highways show a much more green-cyan profile than the center of the city where traffic is generally much slower.

Figure 5.18 Lisbon's road map with major routes colored, namely: A36 (also known as IC17 and CRIL), IP7 (also known as Eixo Norte-Sul), A9 (also known as CREL), EN6 (also known as Estrada Marginal), A5, A8 and 2ª circular (Segunda Circular)[44].

When observing the average velocities in each minute of the visualization, the profile on graph 5.9 is observed. It should be noticed that instead of the hourly average in graph 5.8, this is a running average of velocities in each minute t for the interval $[t{-}30, t]$ minutes. This running average takes into consideration the number of vehicles in the 30 minutes intervals and thus is a more robust measure for the general velocities' profile. In this graph the afternoon rush hour can be seen through a local minimum velocity of 37 km/h at 17:36. Another extreme to highlight is the absolute maximum at 4:31 and the absolute minimum with morning rush hour at 7:35. When observing frames from the visualization (refer to figure 5.19), the morning rush hour and the afternoon rush hour can be visualized when compared to a time in between. In fact, the main arteries such as 2ª circular have much warmer hues during the rush hours than at 11:00. Also, it is worth noticing the center of the city, which is full of vehicles at 11:00 but with much less traffic at 17:36, when vehicles are more concentrated on the main highways leaving the city.

[44] Map background (Toner) by Stamen Design, under CC BY 3.0. Data by OpenStreetMap, under ODbL.

With this visualization of trajectories, it is possible to observe the daily rhythm of the city, confirm the existence of an afternoon rush hour, as well as depict it visually. Coloring trajectories works as a way to extract a comparative notion of the average of velocities on a certain area of the city. It tells the story of the daily rhythm of the city, illuminating, depicting two rush hours, and then fading way. Further development of this work aims at conferring a more dramatic and figurative tone to this story.

Graph 5.9 Average velocity (km/h) of the vehicles in the trajectories visualization by time of the day. This is a running average because at each minute the registries of the last 30 minutes are considered for the average. Some extrema are highlighted in red: absolute maximum at 04:31 of 62 km/h; absolute minimum at 07:35 of 35 km/h; local minimum at 17:36 of 37 km/h and a local maximum at 21:09 of 53 km/h.

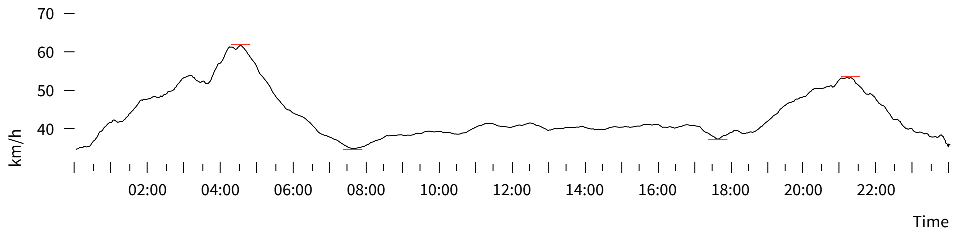

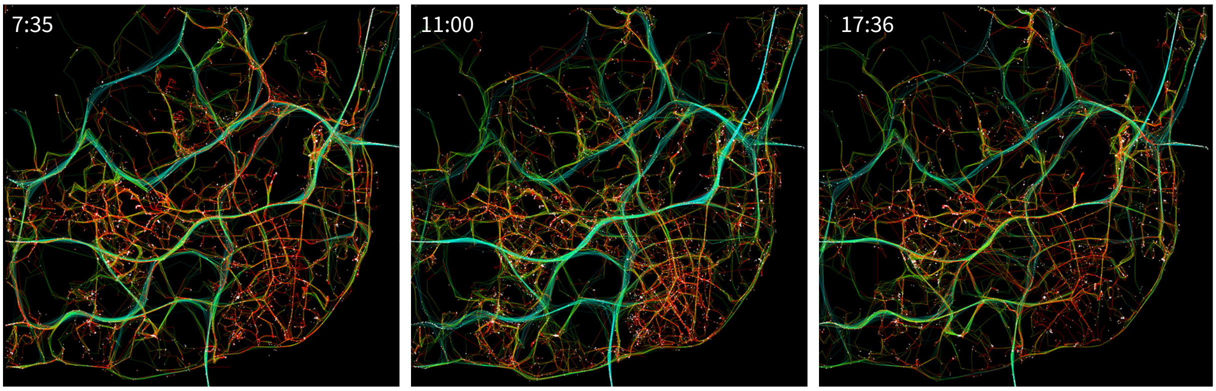

Figure 5.19 Three snapshots of the visualization at 07:35 (morning rush hour), 11:00 and 17:36 (afternoon rush hour). It is possible to observe during the rush hours the main arteries have warmer hues, where traffic is more congested. The animation can be watched in https://vimeo.com/131835197.

This visualization is perhaps the one that goes more in the direction of Manovich's concept of "direct visualization" or "visualization without reduction" (Manovich 2011). In fact, except for the time aggregation applied, the visualization tends to display the data as it would represent the real world, showing moving vehicles as moving vehicles and not reducing this or other aspects to more abstract forms. Instead of presenting a line graph with the number and average velocities of vehicles through time, which despite being more abstract would tell the same story in a more quantifiable manner, the visualization of trajectories privilege spatial variables, and display a system closer to reality – an intent which aims to connect with the viewer within the context of casual information visualization and metaphors.

5.5 Blood vessels

The visualization of trajectories the portrayed complexity of the city with its own its own routines and patterns emerging from the behavior of individual agents. This activity, complexity and life within the city inspired a more figurative metaphor: depict the city as a living organism.

The idea was implemented as a complex system that displays the city as a set of pulsing blood vessels. Each road is mapped to a pulsing vessel, with a pulsing motion proportional to the velocities on that road. If vessels are inside an organism, then the canvas is dark. Vessels are red, but vessels with slower vehicles are darker like clotted blood. Thicker vessels signify more traffic and thus more vehicles. Implanted on the vessels, are streams of white cells. These cells or particles also have a flowing motion proportional to the speed on the vessel and a density proportional to the number of vehicles in it. This visualization, adds visual cues with an inherent metaphorical value that reflects the author's intent. Although strictly non-data, these added aspects are data related, and happen to provide advantages over the visualization of trajectories.

5.5.1 On thickness

Size is systematically rated as a strong visual variable to show quantitative, ordinal, and nominal information. In fact, in terms of accuracy, it is rated second only after position (Bertin 2010, p.69; Mackinlay 1986; Cleveland and McGill 1984). According to Bertin (2010, p.66) the line implantation of size consists of varying line thickness. Line thickness will signify the number of vehicles on the corresponding road, which is a much more accurate way of depicting this information than relying on estimated visual clusters from the juxtapositions of trajectories. Although the problem of juxtapositions was addressed in the visualization of trajectories by using semi-transparent traces, the solution of varying road thickness better deals with this problem by forcing the visual aggregation of vehicles. Nonetheless, such problem was not completely eliminated since major roads that have each direction represented as two bordering roads, will sometimes incur in occlusion problems when their thicknesses expand. In order to avoid such occlusion problems, vessels are drawn with a transparency of 50%. Although undesirable in terms of accurate information portrayal, these juxtapositions are defensible when referring to roads as a living system of blood vessels that are intertwined among themselves and arguably excused when renderings are under the umbrella of casual visualization.

5.5.2 Vessel structure and pulsing motion

Each vessel's structure is built for a road[45]. A road is a set of segments, and thus the structure for the corresponding vessel is computed based on a set of perpendicular segments as illustrated in figure 5.20. The length of each of those segments is proportional to the current number of vehicles on the corresponding road, and therefore manipulates vessel thickness. The shape of a vessel is formed by uniting perpendicular segments' endpoints with Catmull-Rom[46] splines in order to attain a smooth and organic feel.

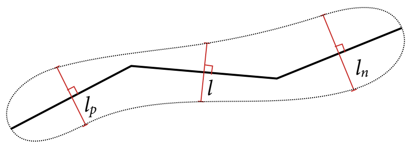

Figure 5.20 A vessel is formed by uniting the perpendicular segment's endpoints with Catmull-Rom splines.

The vessels have a pulsing motion with a rate proportional to the velocities of the corresponding road. In order to attain this, the position and angle of the perpendicular segments vary. Each perpendicular segment travels along the respective road segment and restarts when finished. Let $x \in [0,1]$ be the position of each perpendicular segment, such as when $x = 0$ it is at the beginning of the corresponding road segment, when $x = 0.5$ it is at the middle, and when $x = 1$ it is at the end. Consider the right angles showed in figure 5.20 as the natural angles of each segment l, but that are expressed relatively to the same world coordinate axis, so that the angle of l is α, the angle of the previous segment l_p is α_p and of the next segment l_n, α_n. The actual angle β that each segment l makes in relation to that same world coordinate axis depends on the position x on l (see figure 5.21). When the segment l is at the beginning of the road's segment, its angle will be the average of its natural angle and the natural angle of the previous segment l_p; when the segment l_i is at the middle then its angle is its natural angle; and when it is as the end its angle is the average of its natural angle and the natural angle of the next segment l_n. This way, the angle β is calculated by linearly interpolating through these conditions: when $x = 0$, $\beta = (\alpha_p+\alpha)/2$; when $x = 0.5$, $\beta = \alpha$; and when $x = 1$, $\beta = (\alpha_n+\alpha)/2$. The conditional linear equations that represent this interpolation is as follows:

$$\beta = \begin{cases} (\alpha - \alpha_p)x + 0.5(\alpha_p + \alpha) & 0 \le x \le 0.5 \\ (\alpha_n - \alpha)x + 0.5(3\alpha - \alpha_n) & 0.5 < x \le 1.0 \end{cases}$$

[45] The vessel's pulsing motion and the stream of cells were implemented together with António Cruz

[46] Catmull-Rom splines (Catmull and Rom 1974) are smooth parametric curves that interpolate between a set of points, and are widely used in computer graphics. This method does not require the definition of additional control points for the curves, since the original set of points also makes up the control vertices for the curve.

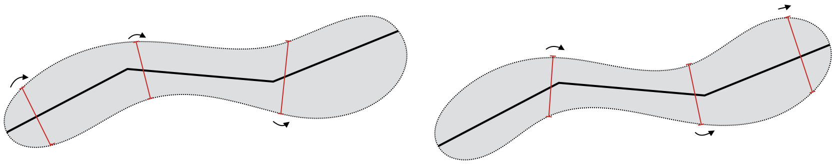

Figure 5.21 The angle of each perpendicular segment varies according to its position over the road's segment. In the left image they are at the beginning and at the right image they are at the end of their corresponding road segments.

By altering the perpendicular segments' orientation, it is possible to depict a pulsing motion. As seen on figure 5.22, varying the angles yields slightly different rendered forms, which when interpolated give the illusion of a pulsing motion. The resulting motion is subtle but sufficient to carry this specific metaphorical intent. Naturally, the rate of such motion is directly dependent on the traveling speed of the perpendicular segments, therefore dependent on how fast x varies from 0 to 1, restarting at 0. This is to say that the pulsing motion's speed can be manipulated by simply varying the iteration rate of x in a way that is proportional to the current average velocity at any given road.

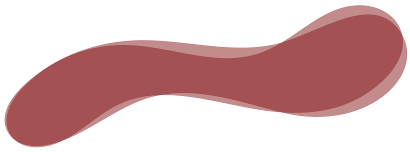

Figure 5.22 Overlay of the rendered shapes from figure 5.21. The continuous interpolation of the vessels' shapes gives the impression of a pulsing motion.

5.5.3 Velocity and cells

Beyond the subtlety of the pulsing motion, velocity is primarily portrayed, in an ordered way, by the colors of the vessels. Brighter reds signify higher average velocities and darker reds slower traffic. By keeping the same red tonality, only the brightness is manipulated, by corresponding velocities to brightness values as shown in figure 5.23.

Additionally, a "stream" of cells was added inside each vessel to emphasize the flow of vehicles (see figure 5.24). The number of cells is proportional to the number of vehicles, and their velocity is proportional to the average velocity on the road. This way, velocity is also encoded in motion in addition to color. The cells are created at the start of each vessel and travel inside the outline of the vessel from one end to the other, restarting when meeting the end point. When the number of vehicles on a road increases, the corresponding amount of cells is added at random positions, and when the velocity decreases, the corresponding amount of existing cells is randomly selected and removed. Each cell is positioned relatively to a perpendicular segment of the road, namely by storing the relative position to that segment as they were created. Therefore, when the perpendicular segments move, the cells move accordingly.

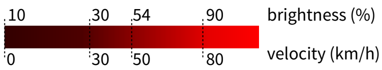

Figure 5.23 Correspondence between the brightness of the vessels and the average velocities in it.

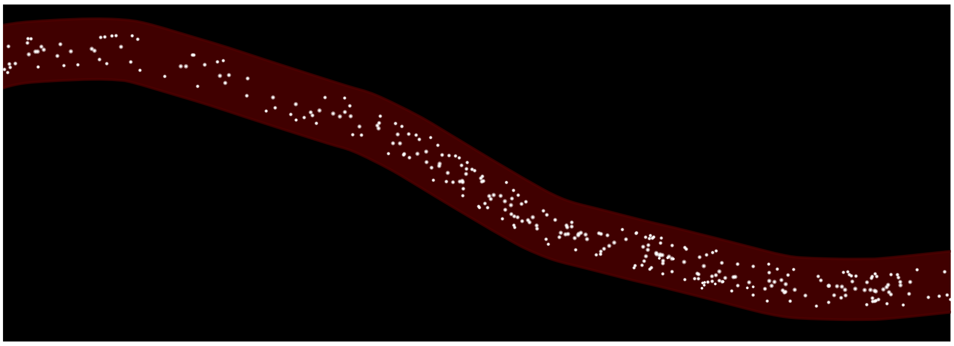

Figure 5.24 A vessel with a stream of cells implanted in it. The density of cells is proportional to the number of vehicles on the corresponding road.

5.5.4 Results

The blood vessels metaphor works particularly well for Lisbon because it is delimited by the Tagus River, which resembles a human heart or an organ. The often-chaotic disposition of roads in a city is the main factor behind the organic expression of this visualization, which can be applied to other cities as well.

The most prominent visual variable at play in this visualization is the size of the vessels, or their thickness. Independently, if they have good or bad circulation, the thickness increases with traffic volume. The implanted streams of cells bring an additional visual variable which is texture. High traffic volumes are then supported by two visual variables: increased densities in the stream's textures, and increased thicknesses of the vessels themselves.

By varying the brightness of reds to depict velocities, brighter reds standout and signify healthier arteries with good vehicle circulation, and darker reds signify low velocities like stagnant blood. It is important to consider this relation in light of the traffic volume mapping. When velocities are low, arteries can fade into the dark background since they will have darker reds, and the attention to those arteries can be lost. When velocities are low but congestion is high, meaning that there is high-traffic volume, such attention can be regained by other visual variables. Thickness alone would not bring attention to heavily congested roads since its color might be very dark. But the high density of the stream's texture in that vessel increases the perceptual vividness of the artery and brings salience to it.

It is important to consider all these visual variables in order associate each type of artery to the state of a road. For example, figure 5.25 shows the blood vessels visualization during the afternoon rush hour, where several types or arteries can be spotted, signifying different traffic conditions:

— Bright red vessels with healthy thickness and cell density means that vehicles are circulating with good velocity and congestion is non-existent.
— Dark and thin vessels with regular texture density indicates that traffic on those roads is indeed slow, but congestion is also low and not many vehicles are using those paths.
— Dark and thin vessels with very high cell density indicates slow traffic and congestion problems. They can be small but easily identifiable on the map since their textures can make them almost solid white.
— Dark and thick vessels with dense textures show where most of the traffic is, with slow velocities, but distinguishable in terms of congestion from the other thinner and denser vessels. The most prominent example of this is 2ª circular during rush hour.

An additional visual variable that reinforces brightness to display velocity is the motion of the stream of cells. The addition of such a visual cue can be seen as redundant or as reinforcing, but it adds a more tangible dimension to the artifact, as if there were in fact agents that make the system itself and cause all the described behaviors. The streams are much slower in the center of Lisbon than in other major arteries, and this can be easily spotted when the artifact is observed in motion. It is argued that such differences in the speed of motion make the visual differences in traffic speed more noticeable, and bring the artifact closer to reality with a more figurative approach, one that shows that overall traffic conditions are caused by individual agents and are not conditions without a cause.

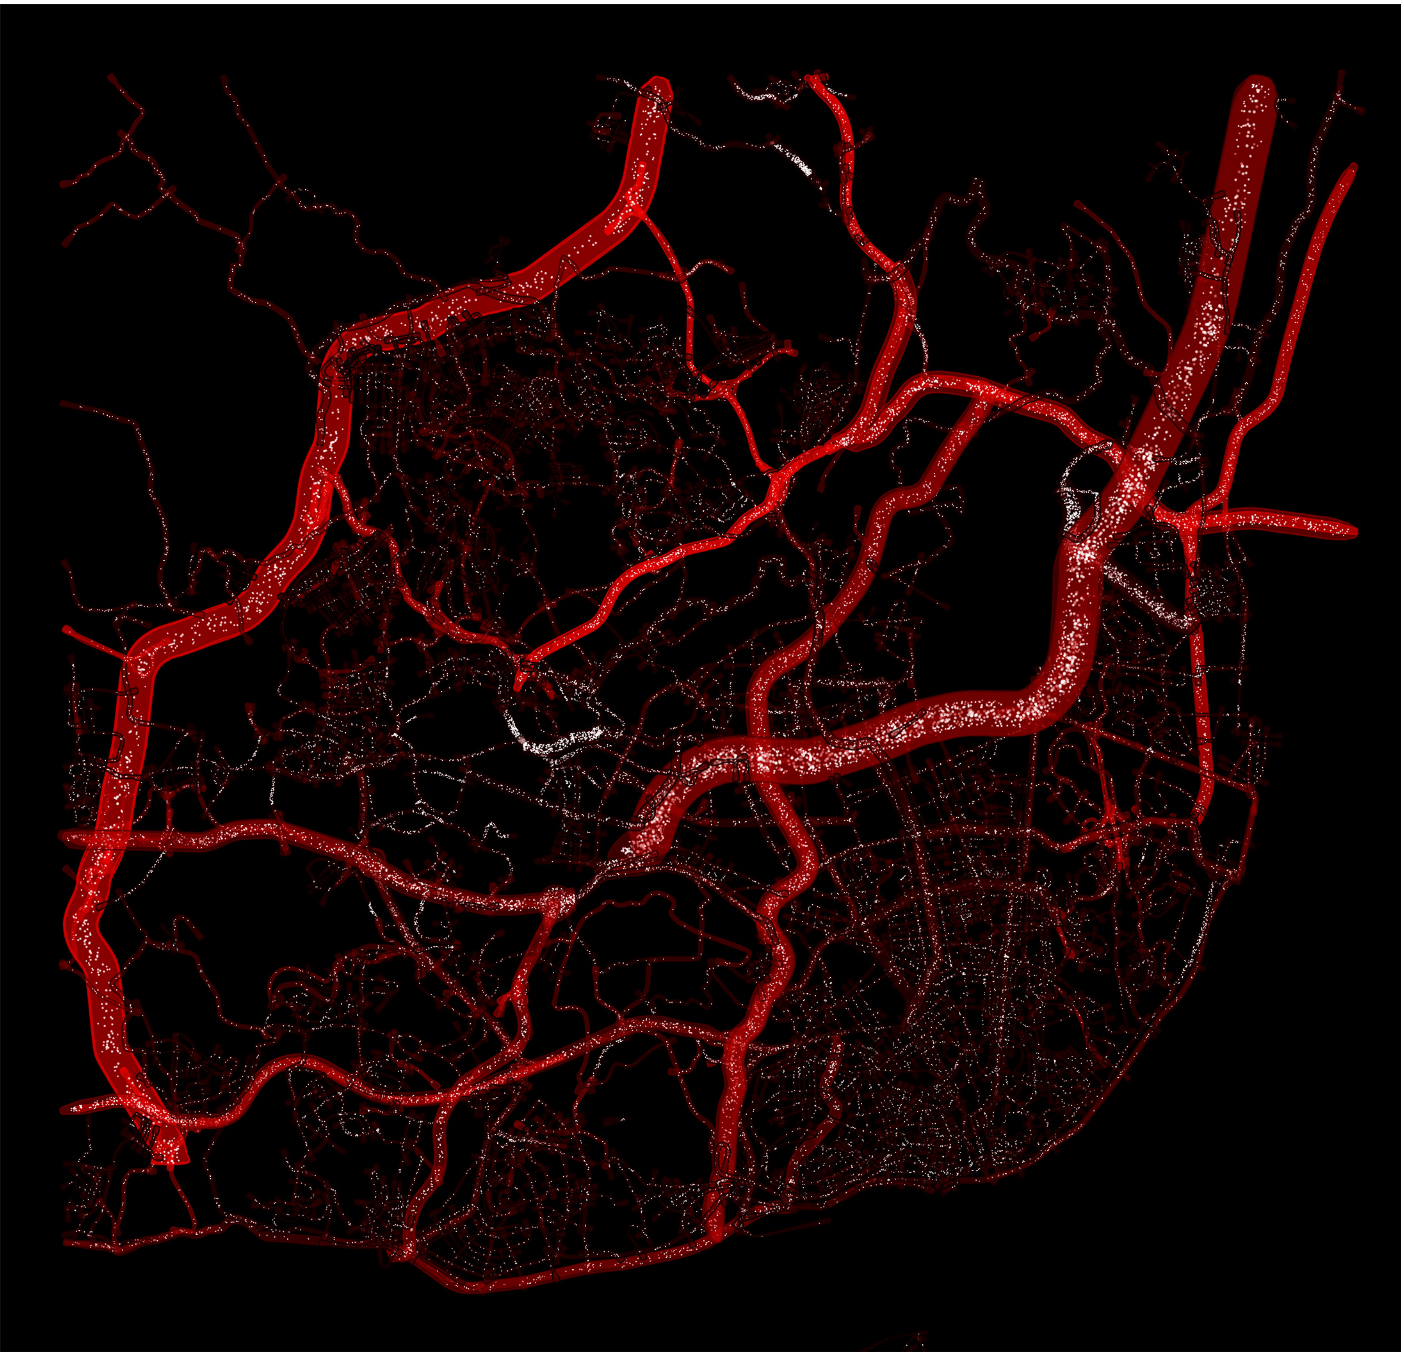

Figure 5.25 Blood vessels visualization during Lisbon's afternoon rush hour. The thickness of the vessels, the densities of the streams of cells, and the color of the vessels signify distinct traffic conditions. The animation can be visualized at https://vimeo.com/131835178

5.6 Building the cartogram

The model previously described to implement the animated cartograms consists on a set of backbone, inner, and connective springs. The backbone springs are needed to extract both the inner and connective springs – these backbone springs reflect the structure of Lisbon's road network. OpenStreetMap[47] (OSM) is used as a source to develop a suitable road network structure made of unique points and edges that connect them. OSM was queried for Lisbon's area (latitudes [38.69°, 38.84°] and longitudes [−9.28°, −9.08°]) and mapped to meters on a plane by using the Mercator projection and the WGS 84 standard equatorial radius. The OSM's format can have roads split through several data structures. In order to eliminate this inconvenient arrangement, split roads were merged if they had the same name and coincident endpoints, resulting in 4,321 roads from the initial 5,131 data structures. Furthermore, the OSM's format also does not guarantee that each point (*x,y*) is only defined once – duplicated points were also eliminated, resulting in a scheme that describes each road as a sequence of references to uniquely defined points (*x,y*). This means that different roads that intersect, can have their intersection points shared, better reflecting the connectivity of the network and the structural interdependence of the roads among one another. OSM's data also came with unnecessary resolution since points were often too close to one another. By removing points on the same road that were closer than one hundred meters, 20,315 points were removed, leaving 17,664 unique points. Given the chosen map scale of 1:60,000, this only resulted in loss of detail for distances inferior to 1.66 mm. Finally, roads that were only made up of one point were removed, leaving 4,184 roads and 17,640 unique points, of which 8,387 (48%) are shared between two or more roads. Using the cartogram model previously described and the respective Delaunay triangulations, the inner and connective springs are computed as illustrated in 5.26.

[47] http://www.openstreetmap.org

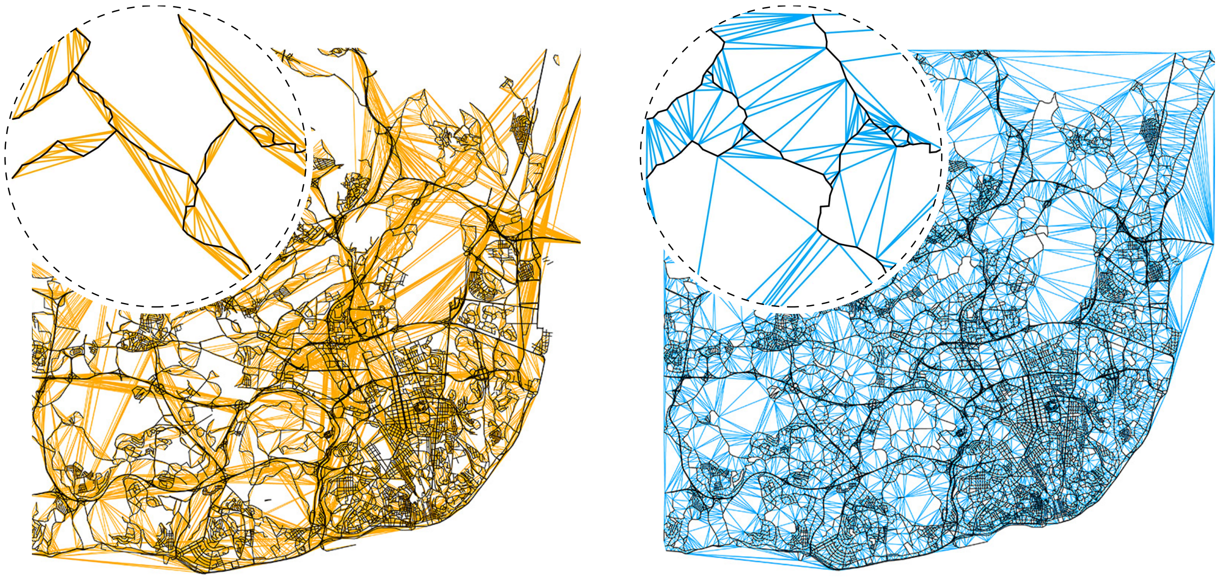

Figure 5.26 The inner springs for Lisbon's road network on the left in yellow, and the connective springs on the right in blue. Recapitulating, the inner springs preserve the inner structure of each road, together with the backbone springs that are formed by the road itself. The connective springs preserve the relative distances of the roads among themselves.

Velocities and distortion ratios

A cartogram distorts a map according to a statistical variable. In this case study, the vehicles' velocities are used to distort the map of Lisbon. The principle is that roads should compress if their vehicles' velocities are low, such as if travel times and hence distances were higher. Conversely, roads with high traffic velocities should compress as if distances were shorter. As discussed in the opening section of this chapter, this approach has already been chosen by several authors as a time-distance metaphor. The central aspect to define within this approach is what the original map reference and the undistorted lengths reflect. Here, instead of imposing an arbitrary reference such as the original lengths signifying velocities of 50 km/h, the original lengths show that traffic velocities are within that road's average. The purpose here is to display through distortions the parts of the day that deviate more from this average.

As previously discussed for the visualization of trajectories, in order to have more data density, all the weekdays of data are aggregated in one single virtual day and time is visualized in aggregations of 30 minutes [t–30, t]. In order to implement this approach, the data had to be structured in a way that translated for each road the current number of vehicles and the current average velocities of those vehicles for each time interval of 30 minutes. In order to do this, every GPS position in the data had to be intersected with the original projected map of Lisbon in order to determine the road that they pertain to for each moment in time. The dataset did not semantically provide any temporal relations between between roads and vehicles. Therefore, such relations had to be extracted using additional techniques. The global average of velocities of each road for the whole virtual day was also determined. The distortion ratio is then computed for each road at a given time interval and is shown as the following pseudo-code. The ratio reflects the current average velocity in relation to the global average velocity for a given road. If the current average is above the global average, the ratio is below 1, meaning that the lengths should be shorter and translating that velocities are higher than usual. For ratios above 1, the current average is smaller than the global, reflecting slower traffic velocities and the lengths should distend together with travel times.

```
computeRatio(road, t):
     ratio = road.globalAvgVel / road.currentAvgVel(t)
     return ratio
```

5.7 Model setup and results

The model for edge-based cartograms does not define the strength of springs, so they have to be experimentally found for this specific case. A promising balance that was initially found is $k_B = 1.20$, $k_I = 0.02$, and $k_C = 0.20$. . It was observed that this configuration made the map globally stretch in the morning and afternoon rush hours when velocities are low, and compressing the map during nighttime. The cartogram was animated from 00:00 to 23:59, iterating minute by minute, and considering, for each minute, the GPS traces occurring in a time-window of the previous 30 minutes. The results for four specific states on this initial configuration can be observed in figure 5.27. The depicted times in this figure refer to minima and maxima of the average velocities: absolute maximum of 62 km/h at 04:31, when velocities are higher and distances are most compressed; absolute minimum of 35 km/h at 07:35, referring to the morning rush hour when the city is most distended; local minimum of 37 km/h at 17:36, distending the city; and local maximum of 53 km/h at 21:09, compressing the city once again. Each road has interpolated colors according to the current distortion ratios, as blue for 0.5 and below, gray for 1.0, and orange for 1.5 and above. It is noticeable that when the city is distended the ratios are usually above 1.0 and when the city is compressed they are below 1.0.

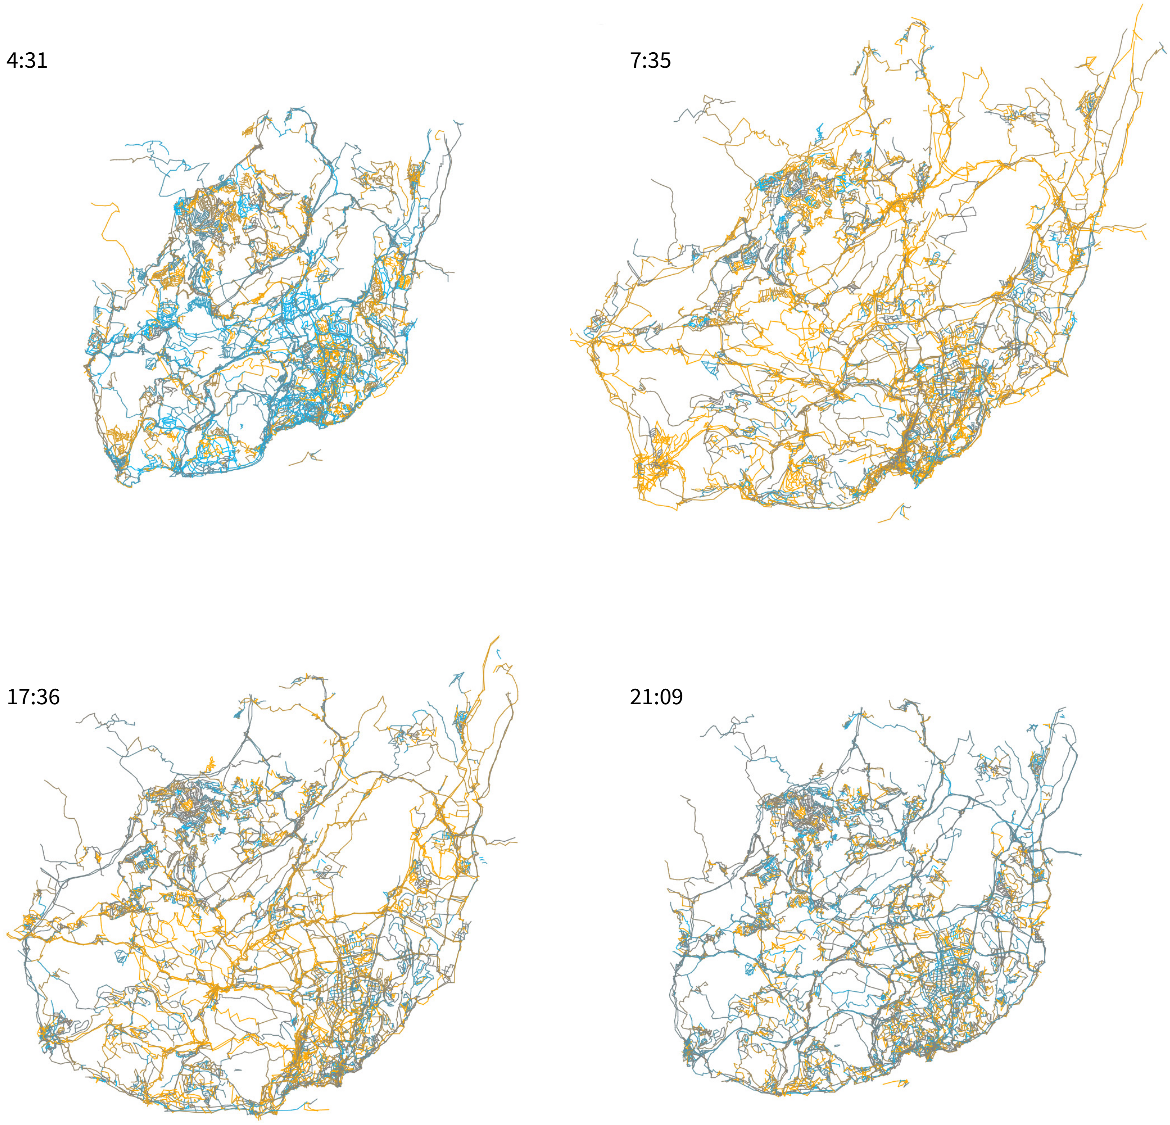

Figure 5.27 First results for the cartogram of Lisbon with $k_B = 1.20$, $k_I = 0.02$ and $k_C = 0.20$. The depicted times refer to minima and maxima of the average velocities. The current distortion ratios of each road are shown with the following color scale: 0.5 ... 1.0 ... 1.5
.

5.7.1 Cartogram errors

The general case of the cartogram problem is not solvable. Therefore, in order to preserve the map's topology, other constraints have to be relaxed, resulting in a deviation from the ideal solution and hence an error. The quality of value-area contiguous cartograms is generally asserted through a value error and a shape error (Keim et al. 2004; Henriques et al. 2009; House and Kocmoud 1998). The value error measures how the represented areas deviate from the values in the dataset – having these values accurately represented would be an ideal solution. For each of these areas, the shape error measures how much the areas were distorted in relation to their initial form. The shape error hence translates how wellthe map can be recognized.

In the case of the edge-based cartogram, the value error is given by the deviation of the road's lengths in relation to their ideal lengths according to the dataset. The formula for this length-resizing error, E_{value}, is adapted from the area error of Keim et al. (2004). Each road i has a value error e_{value} that is a normalization by the absolute difference between the desired length and the actual length of that road at a given time. The value error for the cartogram on a given time, E_{value}, is the weighted average of all the roads' errors by their desired lengths. The actual length of a road is computed by summing all the lengths of its backbone springs. In a similar way, the desired length of a road is the sum of all the applied rest lengths for the backbone springs at a given time – it should be remembered that the rest lengths applied are computed according to the ratio that translates the velocity on that road.

$$e^i_{value} = \frac{\left|\ell^i_{desired} - \ell^i_{actual}\right|}{\ell^i_{desired} + \ell^i_{actual}} \qquad E_{value} = \sum_i \left(e^i_{value} \cdot \frac{\ell^i_{desired}}{\sum_i \ell^i_{desired}} \right)$$

In order to measure the degree of distortion of the map and assert the recognizability of the map, two additional error measures were devised. The distance error, E_{dist}, measures the degree of distortion of distances among roads in relation to the original map, and the shape error, E_{shape}, measures how much the shape of those roads was changed in relation to the original map. Considering the distance error, each connective spring i has a distance error e_{dist} given by a normalized value of the absolute difference between its actual and original lengths. The original length of a connective spring is given by its length in the original undistorted map. The distance error for the cartogram E_{dist} at a certain time is given by the average of the connective springs' errors, weighted by the original lengths of each spring in order to emphasize longer distances.

$$e^i_{dist} = \frac{\left|\ell^i_{original} - \ell^i_{actual}\right|}{\ell^i_{original} + \ell^i_{actual}} \qquad E_{dist} = \sum_i \left(e^i_{dist} \cdot \frac{\ell^i_{original}}{\sum_i \ell^i_{original}} \right)$$

The shape error takes a different approach. Each road, formed by a set of segments, has a set of n internal angles between 0 and π radians. Across the simulation these angles vary from their initial state as in the original map. The shape error e_{shape} of a road j at a given time is the sum of the absolute differences between the original and the actual internal angles, normalized by $n \cdot \pi$. The shape preservation error E_{shape} for the cartogram at a given time is the average of the shape errors of the roads, weighted by their original lengths in order to give more emphasis to longer roads.

$$e^{i}_{shape} = \frac{\sum_{k}^{n} \left| \alpha^{ik}_{actual} - \alpha^{ik}_{original} \right|}{n \cdot \pi} \qquad E_{shape} = \sum_{i} \left(e^{i}_{shape} \cdot \frac{\ell^{i}_{original}}{\sum_{i} \ell^{i}_{original}} \right)$$

5.7.2 Experimental setup

The experiments presented herein explore how the conjugation of k_B, k_I, and k_C can lead to the best errors when using the edge-based model cartogram for Lisbon. Each triplet of (k_B, k_I, k_C) is used to simulate the animated cartogram during 24 hours, iterating one minute per frame with 50 simulation steps per frame.

The original undistorted map corresponds to having all the springs' strengths nullified, i.e., $k_B = 0.00$, $k_I = 0.00$ and $k_C = 0.00$. Accordingly, with the definitions of distance and shape errors, E_{dist} and E_{shape} are zero since the angles and distances between roads are preserved. Nevertheless, across the simulation, the value error E_{value} has a minimum of 0.086, a maximum of 0.208, and an average of 0.112. The purpose is to find solutions with better value errors than for the case where the strengths are null. The cases that have value errors higher than the null case are not solutions at all, since the velocity data would be better represented by the undistorted map. The configuration previously displayed of $k_B = 1.20$, $k_I = 0.02$ and $k_C = 0.20$, shows average value errors approximately five times lower than for the null case. Naturally the improvement of the value error comes at the expense of increased errors for shapes and distances (see table 5.1).

Table 5.1 Minimum, maximum and average of the errors for ($k_B = 1.20$, $k_I = 0.02$, $k_C = 0.20$).

	min	*max*	*avg* ($\overline{E}$)
E_{value}	0.018	0.031	0.021
E_{shape}	0.178	0.193	0.191
E_{dist}	0.063	0.134	0.100

In the first experiments with the springs' strengths, it was noticed that the backbone springs had always to be the strongest in order to have smaller errors than for the null case. The best values for k_B seemed to be around 1.00, with the results definitely worsening for values above 1.50. As for the connective and inner springs, the former generally had to be stronger than the latter, with k_C ranging between 0.05 and 0.20, and k_I ranging between 0.01 and 0.05, seemingly resulting in a low error for the shape and distances. Having this initial orientation, a series of tests were conducted, experimenting with any conjugation with the values (0.00, 0.01, 0.02, 0.05, 0.10, 0.20, 0.50, 1.00, 1.10, 1.20, 1.50, 2.00) for k_B, k_C, and k_I, accounting for a total of $12 \times 12 \times 12 = 1728$ tests.

5.7.3 Springs with zero strength

Having strengths of zero tested in the experiments is important to assert on the effectiveness of each type of spring in controlling each type of error. The backbone springs affect the representation of information the most. When considering the 1728 tests, the experiments that have the backbone springs with zero strength, more than doubles the value error on average. In the same way, the connective springs exert control over the distance error, with errors doubling when the connective springs have no strength. The relation of the inner springs with the shape error is shier. When the inner springs have zero strength, the results on the shape error tend to be 30% higher on average. This is probably due to the fact that either the backbone or the connective springs also contribute to preserve the roads' shapes. The comparison of the errors for springs with zero strength can be seen in graphs 5.10, 5.11, and 5.12.

Graph 5.10 Comparison of E_{value} for null and non null k_B. Average E_{value} through the 144 experiments with $k_B = 0.00$ versus the remaining 1584 experiments with $k_B \neq 0.00$, sorted by value error. The backbone springs are necessary to improve the value error.

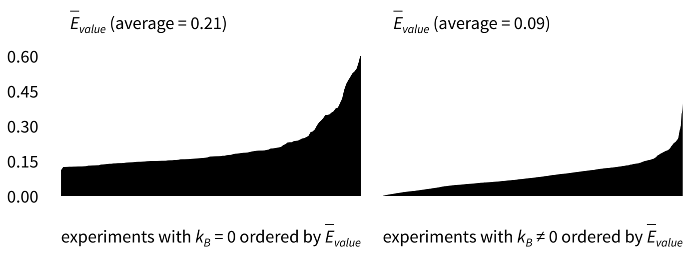

Graph 5.11 Comparison of E_{dist} for null and non null k_C. Average E_{dist} through the 144 experiments with $k_C = 0.00$ versus the remaining 1584 experiments with $k_C \neq 0.00$, sorted by distance error. The connective springs are necessary to improve the distance error.

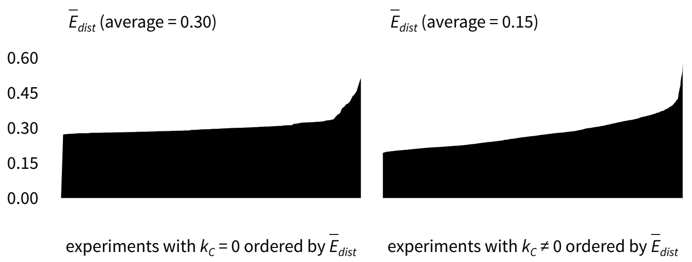

Graph 5.12 Comparison of E_{shape} for null and non null k_I. Average E_{shape} through the 144 experiments with $k_I = 0.00$ versus the remaining 1584 experiments with $k_I \neq 0.00$, sorted by shape error. The inner springs are necessary to improve the shape error.

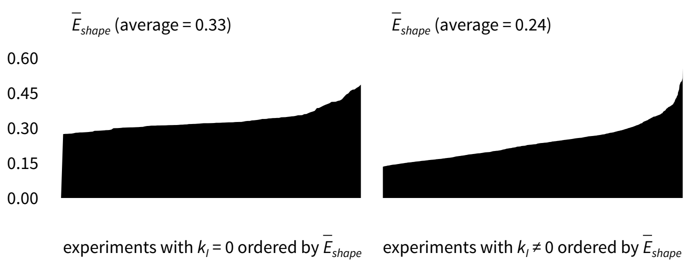

5.7.4 Role of each type of spring

Each type of spring was designed to control a specific type of error: backbone springs for the value error, inner springs for the shape error, and connective springs for the distance error. The efficacy of each type of spring in controlling the errors can be asserted by looking at how their strengths relate with the errors. When analyzing the results, it is important to bear in mind that very stiff springs with strengths of 2.00 worsen the errors due to two things:

— they lead to instability in the system as the particles in the simulation are shifted with distances that are larger than ideal in order to meet their constraints;

— they do not enable less-strong competitor springs to enforce their constraints and thus control the error they were designed for.

Each experiment results in an animated cartogram that has varying errors per frame. This analysis starts by looking at the average errors of each experiment, $\overline{E}_{value}$, $\overline{E}_{shape}$, $\overline{E}_{distance}$. When considering all the 1728 experiments, the overall results for the errors can be seen on table 5.2. When comparing with the results for ($k_B = 1.20$, $k_I = 0.02$, $k_C = 0.20$) on table 5.1, it shows that there is the potential to have better solutions.

Table 5.2 Descriptive statistics of the errors. Minimum, maximum, average, and standard deviation for $\overline{E}_{value}$, $\overline{E}_{shape}$, $\overline{E}_{distance}$, considering the 1728 experiments. *Values that exclude the special case ($k_B = 0.00$, $k_I = 0.00$, $k_C = 0.00$) where both the shape and distance errors are zero.

	min	*max*	*avg*	*std dev*
$\overline{E}_{value}$	0.002	0.614	0.100	0.072
$\overline{E}_{shape}$	0.134*	0.556	0.246	0.078
$\overline{E}_{dist}$	0.096*	0.514	0.152	0.056

Graph 5.13 shows that the backbone springs can control the value error. The value error is generally higher when $k_B < 0.50$, with the higher errors being registered when $k_B = 0.00$. The best results for $\overline{E}_{value}$ are between $k_B = 0.50$ and $k_B = 1.50$. However, stiff backbone springs ($k_B = 2.00$) do not allow the system to make compromises for the length-resizing constraint, leading to general increased value errors. When the other springs are most stiff (both k_I and k_C are 2.00, marked in red), the backbone springs are blocked from scaling properly given their surrounding forces, causing higher value errors. Backbone springs also control the shape error since their strength is as important as the strength of the inner springs, which ensures the proportional resize of a road and hence preserves its internal angles. The shape error improves as k_B increases, stabilizing on $k_B = 0.50$ onwards until $k_B = 2.00$, when the springs are stiffer and the error increases. Similarly with the value error, the shape error has the worst

results for $k_I = 2.00$ and $k_C = 2.00$ (marked in red), regardless of k_B. The distance error is generally unaffected by the backbone springs, except for when $k_B = 0.00$ and $k_B = 2.00$. Nevertheless, marked in blue, is a significant group of experiments in which $\overline{E}_{dist}$ is substantially higher. These highlighted experiments have the connective springs nullified ($k_C = 0.00$), which are the springs specifically designed to control the distance error.

Graph 5.13 Relation between backbone springs and errors. $\overline{E}_{value}$, $\overline{E}_{shape}$ and, $\overline{E}_{dist}$ of the experiments versus k_B (backbone spring strength) with a fitted polynomial trendline of order four. Experiments in which both k_I and k_C are 2.00, are marked in red. On the rightmost graph, for k_B versus $\overline{E}_{dist}$, the experiments highlighted in blue have $k_C = 0.00$. The backbone springs are able to control both the value and shape errors. The backbone strengths do not generally affect the distance error.

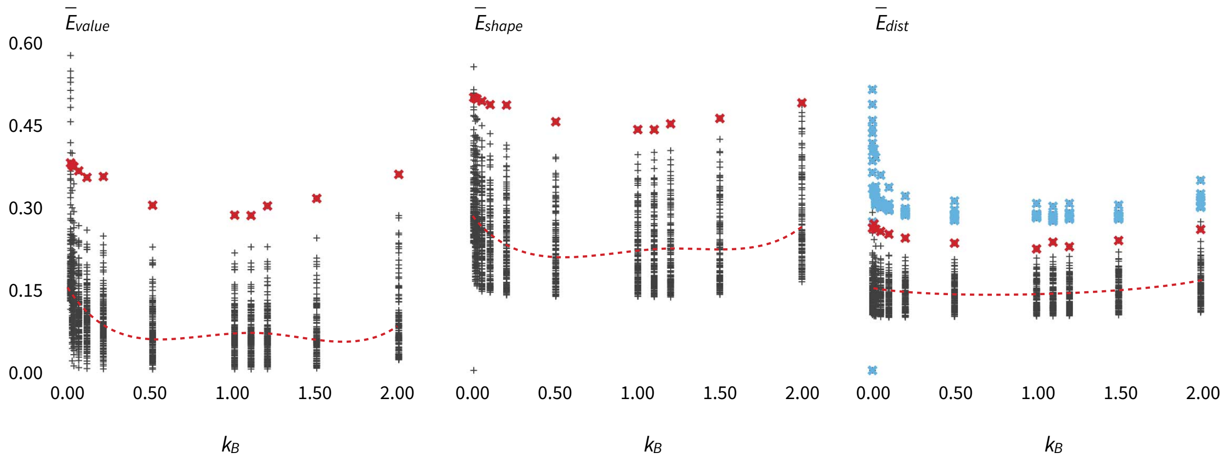

Graph 5.14 shows that inner springs worsen both the value and distance errors, since forcing to maintain the roads' inner structure does not allow the other constraints to efficaciously take place. Nevertheless, the inner springs contribute to improve the shape error, but only for strength values between 0.01 and 0.50. The fact that inner springs fail to maintain the shape error for when their strengths are above 0.50 can be explained through the E_{shape} dependence on the backbone springs – when inner springs are too strong when compared with the backbone springs, the later can fail to adequately adjust the roads lengths in order to ensure the preservation of the road's inner angles, and thus minimize the shape error. Furthermore, the shape error is sensitive to both stiff backbone and connective springs. The highlighted experiments in orange on graph 5.14, in k_I versus $\overline{E}_{shape}$, both have the connective and backbone springs with strengths of 2.00. This demonstrates that stiff k_C or k_B increase the shape error. Considering the distance error in graph 5.14, the experiments with $k_C = 0.00$ (marked in blue) have substantially higher errors, being the worst case scenario for when k_B is also 0.00, marked in red. This demonstrates that even if the connective springs are strong, they cannot be efficacious in maintaining the distances between roads if the backbone springs are not present.

Graph 5.14 Relation between inner springs and errors. $\overline{E}_{value}$, $\overline{E}_{shape}$, and $\overline{E}_{dist}$ of the experiments versus k_I (inner-spring strength) with a fitted polynomial trendline of order four. Experiments in which both k_C and k_B are 0.00, are marked in red. For k_I versus $\overline{E}_{shape}$, the highlighted experiments in orange have both $k_B = 2.00$ and $k_C = 2.00$. On the rightmost graph, for k_I versus $\overline{E}_{dist}$, the experiments highlighted in blue have $k_C = 0.00$. The inner springs are able to control the shape error, but as their strengths increase so do the value and distance errors.

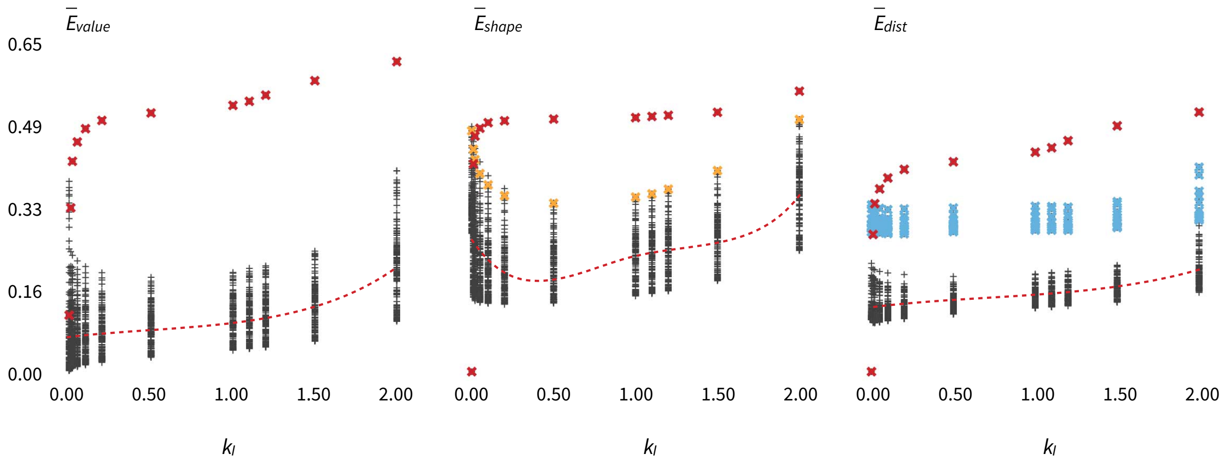

As seen in graph 5.15, as the connective strengths increase, so do the value and shape errors. The distance error can be controlled with the connective springs, with best results between $k_C = 0.01$ and $k_C = 0.20$. On graphs 5.13 and 5.14, the $\overline{E}_{dist}$ results show that stiff backbone and inner springs of strength 2.00 increase the distance error. This can be better visualized on the $\overline{E}_{dist}$ graph in figure 5.15, in which the experiments highlighted in blue have either $k_I = 2.00$ or $k_B = 2.00$. Except for when $k_C = 0.00$, stiff inner and backbone springs are the solely responsible for maximizing the distance error, regardless of k_C.

Graph 5.15 Relation between connective springs and errors. $\overline{E}_{value}$, $\overline{E}_{shape}$, and $\overline{E}_{dist}$ of the experiments versus k_C (connective spring strength) with a fitted polynomial trendline of order four. Experiments in which both k_I and k_B are 0.00, are marked in red. On the rightmost graph, for k_C versus $\overline{E}_{dist}$, the experiments highlighted in blue have either $k_I = 2.00$ or $k_B = 2.00$. The connective springs can control the distance error, but as their strengths increase, so do the value and shape errors.

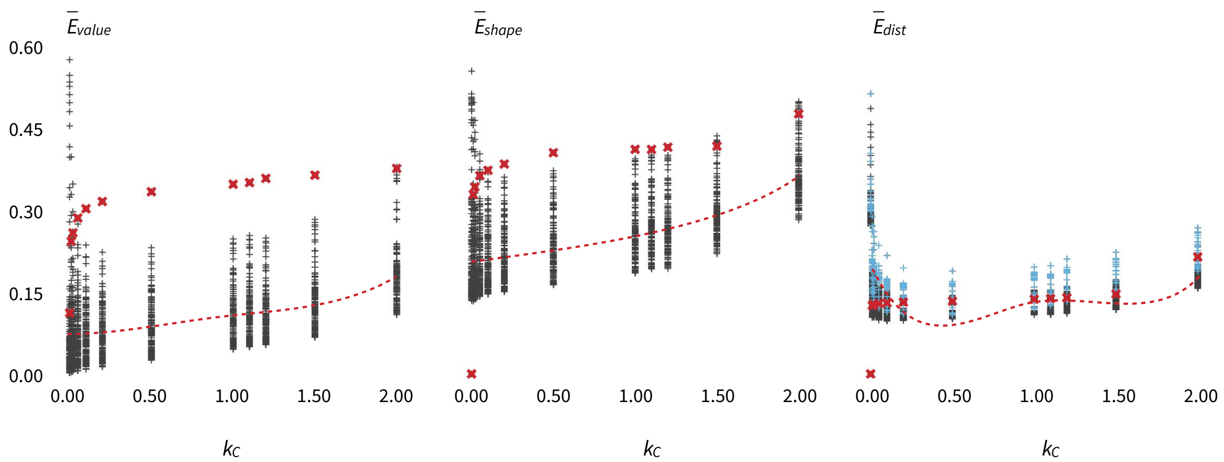

To summarize, this first analysis of strength versus error leads to the following conclusions.

- — Backbone springs provide best results for the value and shape errors when $k_B \in [0.50, 1.50]$.
- — Backbone springs do not influence the distance error except when $k_B = 2.00$.
- — Inner springs are able to control the shape error when $k_I \in [0.01, 0.50]$.
- — Inner springs worsen both the value and distance errors.
- — Connective springs are able to control the distance error when $k_C \in [0.01, 0.20]$.
- — Connective springs worsen both the value and shape errors.

5.7.5 Quality of the solutions

Another way to interpret the results it to lay out the experiments by their value, shape and distance *maximum* errors. The scatterplot on graph 5.16 shows the experiments with non null springs by their maximum errors: value error in x, shape error in y, and color for the distance error. Small values for distance error are red and shown in bigger sizes. The experiments that have k_B, k_I, or k_C as zero are plotted separately on graph 5.17 since they are inherently bad solutions and do not follow the general distribution. The pareto frontier for experiments with non-null spring strengths is plotted in graph 5.18. As suggested in the previous analysis, the value and shape errors seem to be related since they generally increase together. Therefore it is easy to identify experiments that minimize both the maximum errors for value and shape. The maximum distance errors, although, do not follow this trend. The best maximum distance errors do not minimize the value and shape errors, and therefore a compromise has to be met in order to choose the best solutions. The Spearman's correlation coefficients among the three errors are described in table 5.3.

Table 5.3 Spearman's correlation coefficients among the three types of errors in experiments with non-null strengths.

	ρ
$\max(E_{value}), \max(E_{shape})$	0.73
$\max(E_{value}), \max(E_{dist})$	0.59
$\max(E_{shape}), \max(E_{dist})$	0.20

Graph 5.16 Scatterplot of experiments with non-null spring strengths. Maximum value error on *x*, maximum shape error on *y*, and maximum distance error by color. The best maximum distance errors are not the best maximum value and shape errors. Nevertheless, several candidates for the best solution can be identified when x ∈ [0.0, 0.1] and y ∈ [0.1, 0.2].

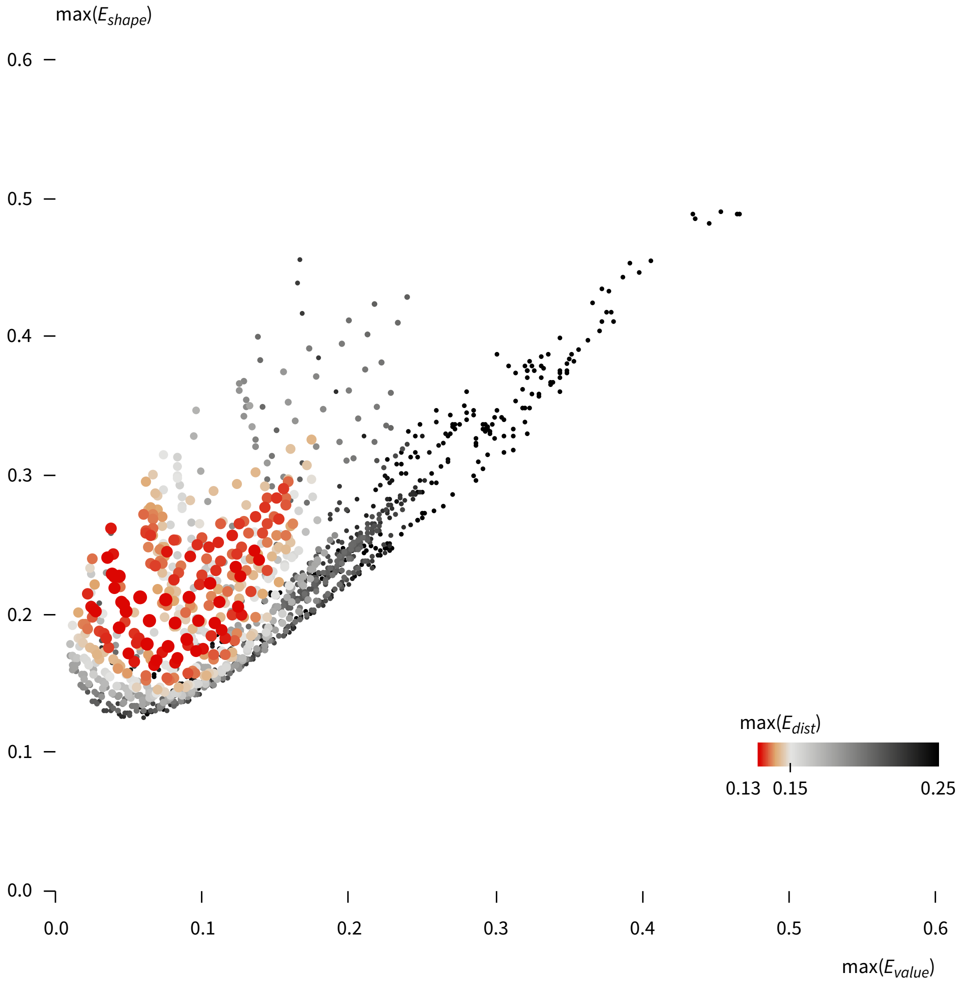

Graph 5.17 Scatterplot of experiments with $k_B = 0.00$, $k_I = 0.00$, and $k_C = 0.00$. Maximum value error on x, maximum shape error on y, and maximum distance error by color as in graph 5.16. These cases constitute bad results since at least one error component is left uncontrolled.

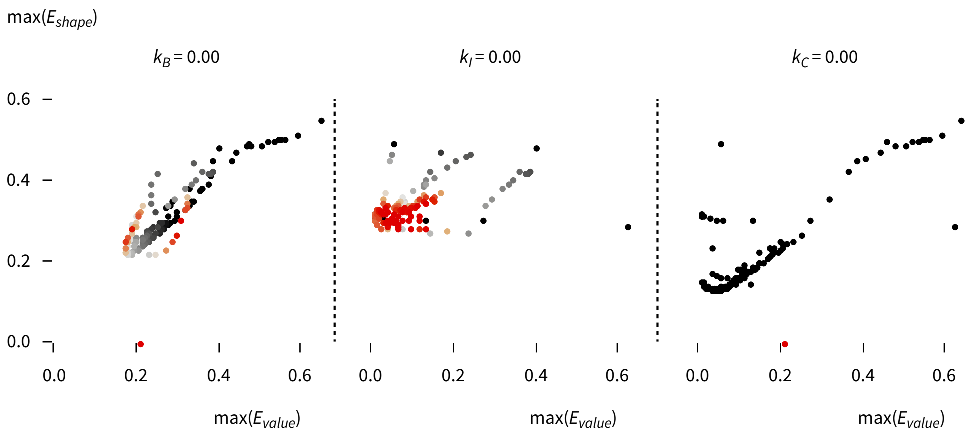

Graph 5.18 Pareto frontier of experiments with non-null spring strengths. Maximum value error on x, maximum shape error on y, and maximum distance error by color.

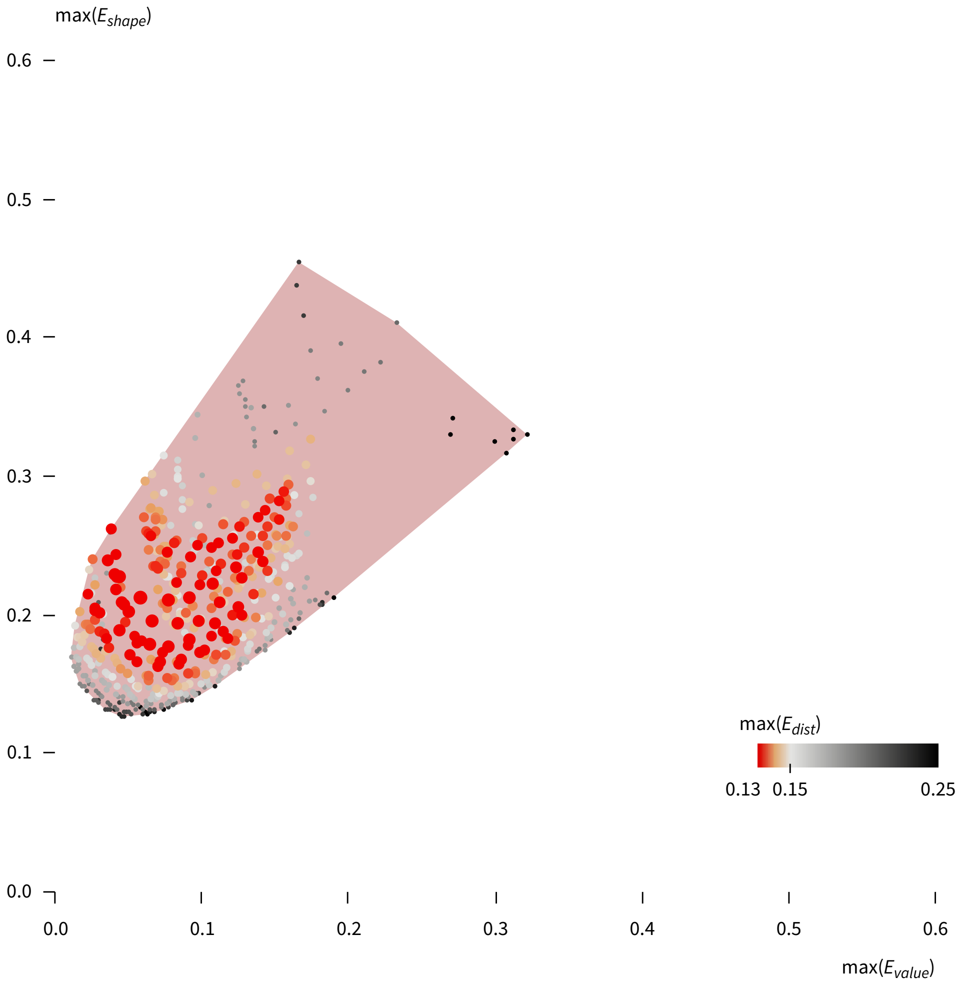

In order to determine the best solutions, a quality measure that attributes the same importance to the value, shape, and distance errors of each solution was devised. For each experiment, its quality is given by the sum of the maximums E_{value}, E_{shape}, and E_{dist}, normalized by their respective absolute maximums and minimums through the whole set of experiments. The smallest values of Q look to minimize the maximum errors of the solutions. The formula for the quality measure Q of one experiment i the universe of n total experiments, where $i \in n$, is as follows:

$$Q = \sum_{e \in \{\text{value, shape, distance}\}} \frac{\max E_e - \min_i(\max E_e)_i}{\max_i(\max E_e)_i - \min_i(\max E_e)_i}$$

The best solution from the considered settings, is $k_B = 0.50$, $k_I = 0.01$ and $k_C = 0.05$, and the second best is $k_B = 1.00$, $k_I = 0.02$, and $k_C = 0.05$. When considering the first experimental solution of $k_B = 1.20$, $k_I = 0.02$, and $k_C = 0.20$, it comes in 45th position out of the 1728 experiments.

As described, each experiment shows an evolving city during 24 hours, compressing and distending due to traffic velocities. This is the most *dramatic* effect of the animated edge-based cartogram, and it varies from experiment to experiment. Therefore, it can be said that some solutions are more dramatic than others even if they have well controlled errors. This variation in size of the city across each simulation was estimated by computing the minimum and maximum areas of the city's bounding box. The most dramatic moment for this dataset is around the morning's rush hour where the maximum area variation is expected. This variation occurs abruptly, and exposes the viewer to the most noticeable contrast in the city's size. Other metrics, such as the standard deviation of the area, would measure instead the *dynamism* of the city during simulation time, and not necessarily how dramatic it was. The area difference between the maximum and minimum gives an estimate of the dramatic effect that can be seen on table 5.4. The best solution according to Q accounts for an area variation of 93% and the second best, 82%. The area variation of the initial experimental solution ($k_B = 1.20$, $k_I = 0.02$, $k_C = 0.20$) is only 72%.

Table 5.4 Hundred best conjugations of backbone, inner, and connective spring strengths. These parameterizations have the smallest values of Q, and have a corresponding area variation (Q ▪ *Area variation*).

		k_B							
k_C	k_I	0.05	0.1	0.2	0.5	1	1.1	1.2	1.5
0.01	0.01		.63 ▪ .88	.62 ▪ .85	.61 ▪ .90	.63 ▪ .83	.62 ▪ .87	.63 ▪ .86	.64 ▪ .84
	0.02			.64 ▪ .94	.63 ▪ .93	.63 ▪ .90	.63 ▪ .90	.63 ▪ .90	
0.02	0.01	.63 ▪ .84	.62 ▪ .92	.61 ▪ .87	.60 ▪ .87	.61 ▪ .86	.61 ▪ .88	.62 ▪ .85	.63 ▪ .89
	0.02		.64 ▪ .93	.62 ▪ .89	.61 ▪ .86	.61 ▪ .84	.62 ▪ .85	.62 ▪ .84	.63 ▪ .83
	0.05					.64 ▪ 1.01	.64 ▪ 1.00	.64 ▪ .99	
0.05	0.01	.64 ▪ .88	.61 ▪ .91	.60 ▪ .89	.60 ▪ .93	.61 ▪ .86	.61 ▪ .86	.61 ▪ .90	.63 ▪ .86
	0.02		.62 ▪ .91	.61 ▪ .89	.60 ▪ .88	.60 ▪ .82	.61 ▪ .81	.61 ▪ .83	.62 ▪ .78
	0.05			.64 ▪ 1.01	.62 ▪ .98	.62 ▪ .96	.62 ▪ .92	.62 ▪ .93	.63 ▪ .90
	0.1					.64 ▪ 1.06	.64 ▪ 1.06	.64 ▪ 1.06	
0.1	0.01		.63 ▪ .93	.61 ▪ .93	.61 ▪ .92	.62 ▪ .83	.62 ▪ .87	.62 ▪ .86	.63 ▪ .86
	0.02		.63 ▪ .89	.62 ▪ .88	.60 ▪ .86	.60 ▪ .81	.61 ▪ .79	.61 ▪ .80	.62 ▪ .75
	0.05			.63 ▪ 1.01	.62 ▪ .98	.61 ▪ .92	.61 ▪ .91	.61 ▪ .90	.62 ▪ .86
	0.1					.63 ▪ 1.01	.63 ▪ 1.00	.63 ▪ 1.00	.64 ▪ .97
0.2	0.01			.64 ▪ .88	.63 ▪ .85	.63 ▪ .78	.64 ▪ .82	.64 ▪ .82	
	0.02			.64 ▪ .86	.62 ▪ .83	.62 ▪ .76	.62 ▪ .77	.62 ▪ .72	.63 ▪ .68
	0.05				.62 ▪ .94	.61 ▪ .92	.61 ▪ .90	.62 ▪ .86	.62 ▪ .83
	0.1				.64 ▪ 1.05	.63 ▪ 1.01	.63 ▪ 1.01	.63 ▪ .95	.63 ▪ .94
0.5	0.05						.64 ▪ .77		

Q
.60 – .64

Area variation
▪ .68 ▪ .80 ▪ .90 ▪ 1.00 ▪ 1.06

Naturally, each error affects Q since they are directly included in the definition of Q. Such correlation can be observed on graph 5.19. Another characterization of the cartogram simulation that does not follow such correlation is area variation. In fact, experiments can have different area variations for more than 100%, while yielding a similar Q. What can be asserted according to 5.19 is that a low Q, or a good solution, limits the amount of area variation. It should also be noticed that area variation has comfortable values around 100% for smaller Qs. This means that looking for good solutions does not end up in static cartograms with no area variation.

Graph 5.19 Scatterplots of quality versus error and area variation. In red, to 100 best solutions, or with the smallest Q. Errors and Q are strongly correlated. Looking at the relation between Q and area variation, having good solutions imposes limitations on the amount of area variation of the cartogram.

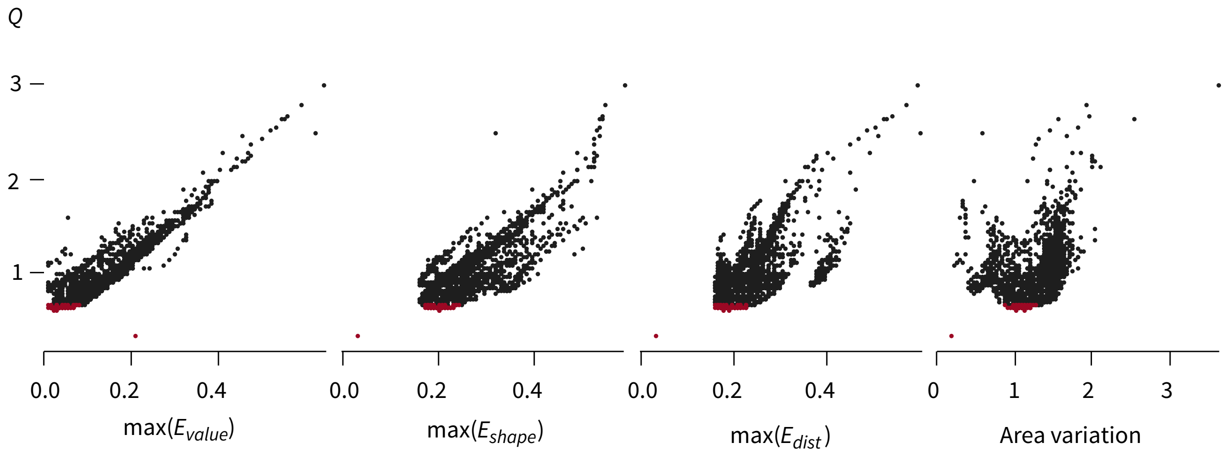

Table 5.4 shows the top one hundred best solutions according to the quality measure as well as their respective backbone, inner, and connective spring strengths. Although not being the fittest according to Q, other solutions can display area variations as high as 106% and be on the top 100 experiments with the best errors. Table 5.4 can be used in order to choose a system configuration that provides solutions with low errors while enhancing the dramatic effect of the cartogram.

Figure 5.28 shows five generated cartograms for Lisbon at 07:35 for different configurations. The time of 07:35 is the peak of the morning rush hour, when velocities are lower and the city is expected to be most distended. This maximum expansion translates into higher errors, since the city is forced to mutate its shape, having a maximum deviation from its original form. In this case, the positions of the system's particles are changed more abruptly, in quantity and magnitude, leading to more chances of having instability in the system. Therefore one can argue that this particular state of the animation for the cartogram is a good testing criteria in order to visually assert how well the solution can solve its more problematic case. For example, when visually observing the best one hundred results according to Q, there were no false positives, meaning that there was not a solution that notoriously presents an unrecognizable Lisbon. For the best case according to Q ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$), the city is both more distended along its main roads as well as more compressed in particular areas such as center, when compared with the initial experimental solution ($k_B = 1.20$, $k_I = 0.02$, $k_C = 0.20$) on figure 5.27. The case of $k_B = 1.10$, $k_I = 0.05$ and $k_C = 0.20$ on figure 5.28 is another good solution, being 22nd best overall. In this solution, the shape of the city is better preserved, naturally at the expense of a slightly increased value error – in the best case $\overline{E}_{value} = 0.020$ and $\overline{E}_{shape} = 0.178$, and in this case, $\overline{E}_{value} = 0.026$ and $\overline{E}_{shape} = 0.169$.

The bottom row of figure 5.28 shows three bad solutions. For $k_B = 1.00$, $k_I = 0.01$,and $k_C = 0.00$, the absence of the connective springs cause the roads to behave more independently, rendering more fine-grained road structures unrecognizable and stopping the city to compressing and distending uniformly. This configuration is the 623rd best solution according with Q. The configuration that has every spring strength maxed out ($k_B = 2.00$, $k_I = 2.00$ and $k_C = 2.00$) is the 19th worst result out of the 1728 experiments, causing instability in the system and rendering the city unrecognizable. The very worst result according to Q has stiff inner springs and the remaining springs nullified ($k_B = 0.00$, $k_I = 2.00$, $k_C = 0.00$). The simulation of this configuration renders the city structure completely unrecognizable, since the only springs responsible for intertwining different roads are the backbone and connective springs.

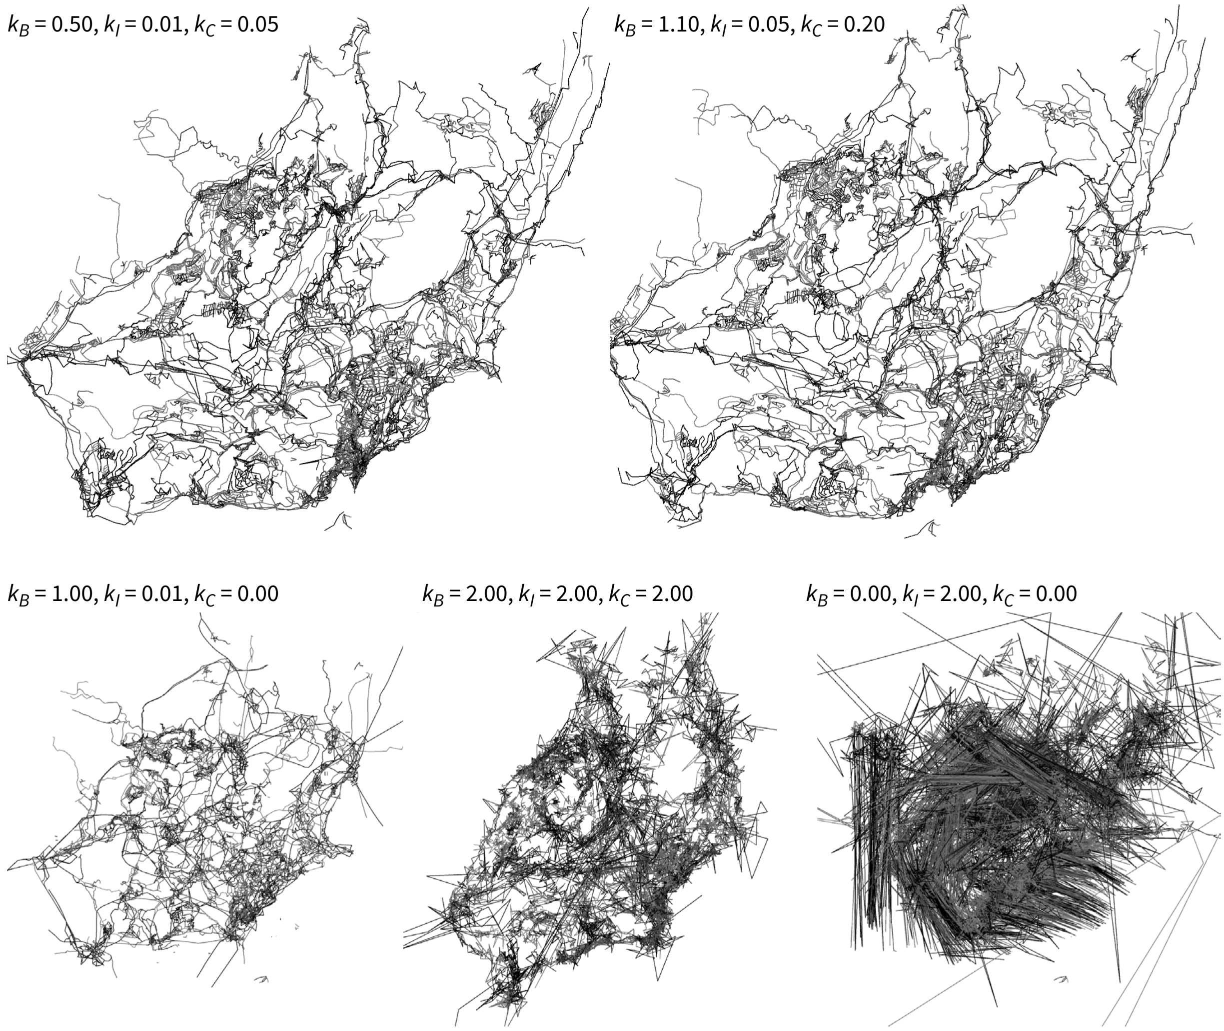

Figure 5.28 Five generated cartograms for Lisbon at 7:35 for different spring-strength conjugations. Top left: best solution; top right: 22nd best solution; bottom left: 623rd best solution; bottom center: 19th worst solution; bottom right: the worst solution.

5.8 Over the cartogram

The edge-based cartogram that was devised to depict the city's traffic models the spatial arrangement of the city and thus modifies its underlying structural metaphor, which is the first step to an approach to semantic figurative metaphors. The cartogram is dramatic in its nature, but this contrast with the familiar map of the city should not be abstracted from the data that provides such a counter-vision. With this, the most direct visualization of Lisbon's traffic data, the visualization of trajectories, can be overlaid on the cartogram, displaying the data it uses to distort the map, and providing a much more detailed and intricate portrait of the city.

In order to do this, each vehicle trace is expressed in relation to the closest backbone spring at any given moment. Each trace is at a given distance from its closest backbone spring, and the projection of that trace on the spring gives a one-dimensional position over that same segment. These relative positions of the traces to the closest backbone spring are expressed in a similar way to the perpendicular segments that travel over a road for the blood vessels implementation – in this case a distance to the segment and a one-dimensional position $x \in [0, 1]$ on the spring. The relative positions are computed before the simulation starts for every trace at any given time, relatively to the static backbone springs (i.e., the segments of Lisbon's road network). When the backbone springs change in position and size, the traces are mapped according to their relative position to the springs, and thus the trajectories of the vehicles follow the distorted map of Lisbon. The results of this implementation can be observed in figure 5.29 for the configuration ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$). The continuous animation of the visualization of trajectories over the cartogram is highly complex, presenting variable behaviors depending on the city's region as well as emergent ones, so that the city compresses and expands as a whole. This approach adds a redundancy in data representation that is used to emphasize aspects in data. For example, regions with slow traffic are not only displayed in orange and red, but also compress as a result of the slow velocities, therefore translating velocities into both space and color. A detailed version of this duality can be observed in figure 5.30, where the main arteries of the city can be seen expanded and represented with yellow and orange tones.

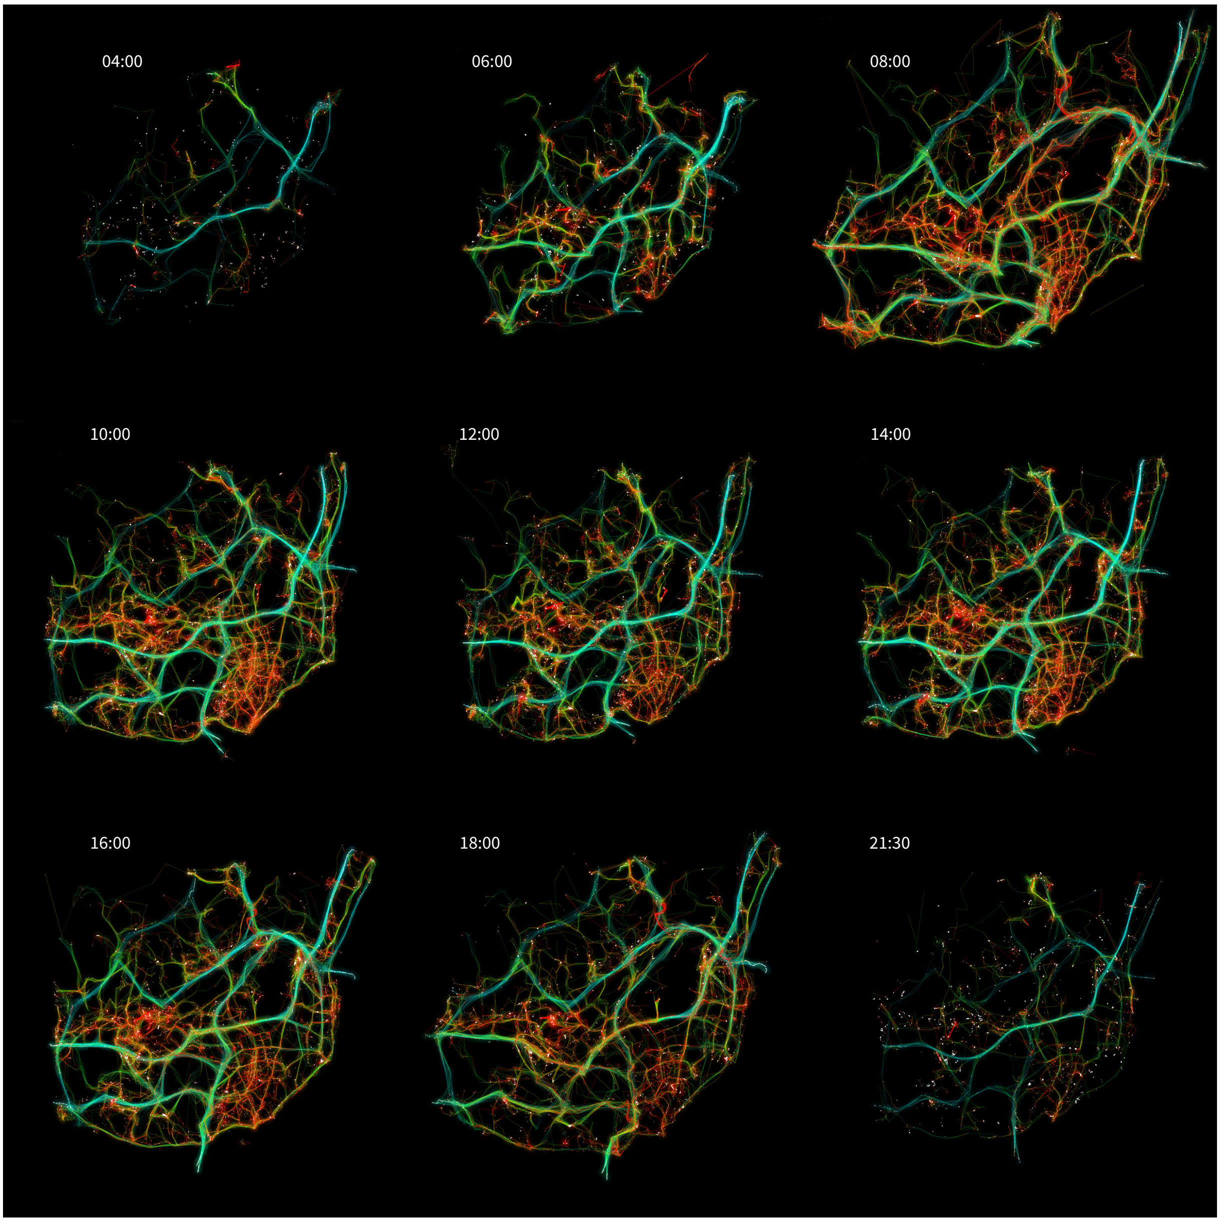

Figure 5.29 Snapshots of the visualization of trajectories applied over the cartogram for several times of the day for the configuration ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$). Overlaying the trajectories on the cartogram shows a city that expands and compresses with a complex representation originating such behavior. The animation can be watched at https://vimeo.com/131835139.

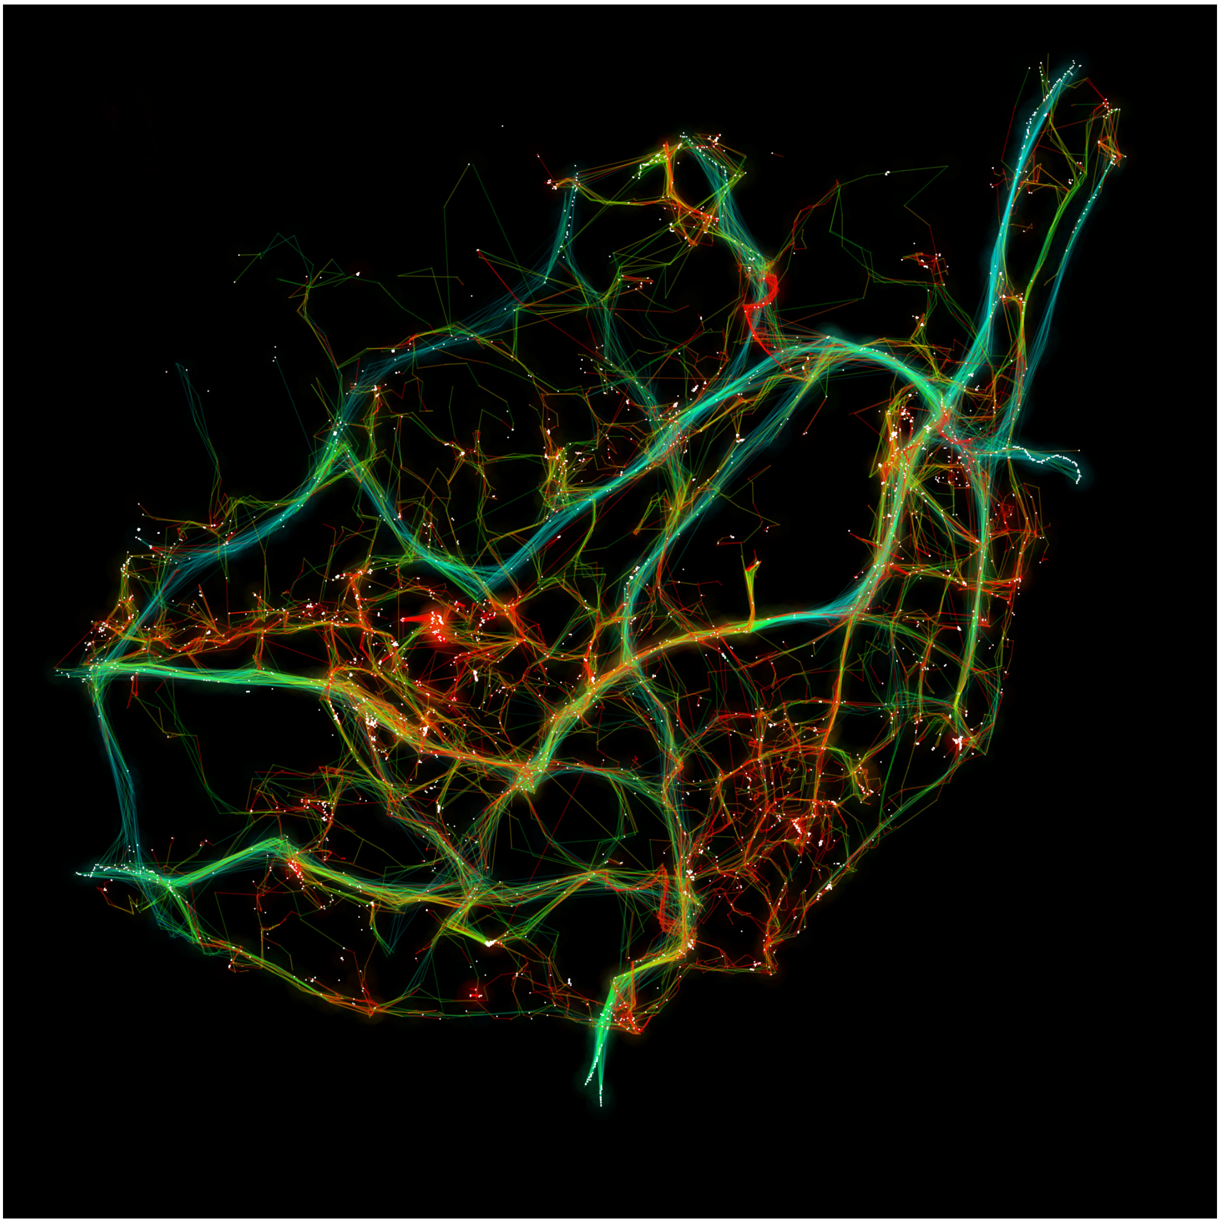

Figure 5.30 Snapshot of the visualization of trajectories applied over the cartogram at 18:00 for the configuration ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$). The animation can be watched at https://vimeo.com/131835139.

The blood vessels visualization was as well applied to the cartogram. Such application required the backbone springs of the cartogram to act as the road segments of the vessels. The results of this implementation can be observed in figures 5.31 and 5.32, for the configuration ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$). As previously noted, the blood vessels visualization works particularly well for Lisbon since its geographical arrangement resembles a heart or an organ. The application of the blood vessels to the cartogram brings a strong metaphorical component, because it not only uses the time-distance metaphor in a data-related manner, but also offers a new language when portraying traffic. The result is a figurative portrayal that has semantic meaning tightly coupled with the data – the city is a living organism, but it behaves within the data, with varying thicknesses, cell densities and speeds, pulsing motions, and overall expansions and contractions.

Until now, the cartogram has explored the modification of the structural metaphor while the vessels visualization brought new visual cues that embody the figurative component of the artifact. Having these two steps in the same visualization results in a rendering of a semantic figurative metaphor at its fullest. The motion of the vessels and the streams of the cells, together with the expansions and contractions of the cartogram, add visual behaviors that are unrelated to the data and are perhaps too distracting, potentially having reached the limits of eccentricity when working with figurative metaphors in visualization. The application of blood vessels on a dynamic cartogram is graphically strong and is figuratively close to its metaphorical domain. Nonetheless, it also distances itself from the elegance of simplicity, where the authorial intent superimposes the functionality of data portrayal. Although permissible in the context of semantic figurative metaphors, there is a point where the beauty of complexity might turn into visual confusion, and such should be avoided and well framed within specific authorial intents.

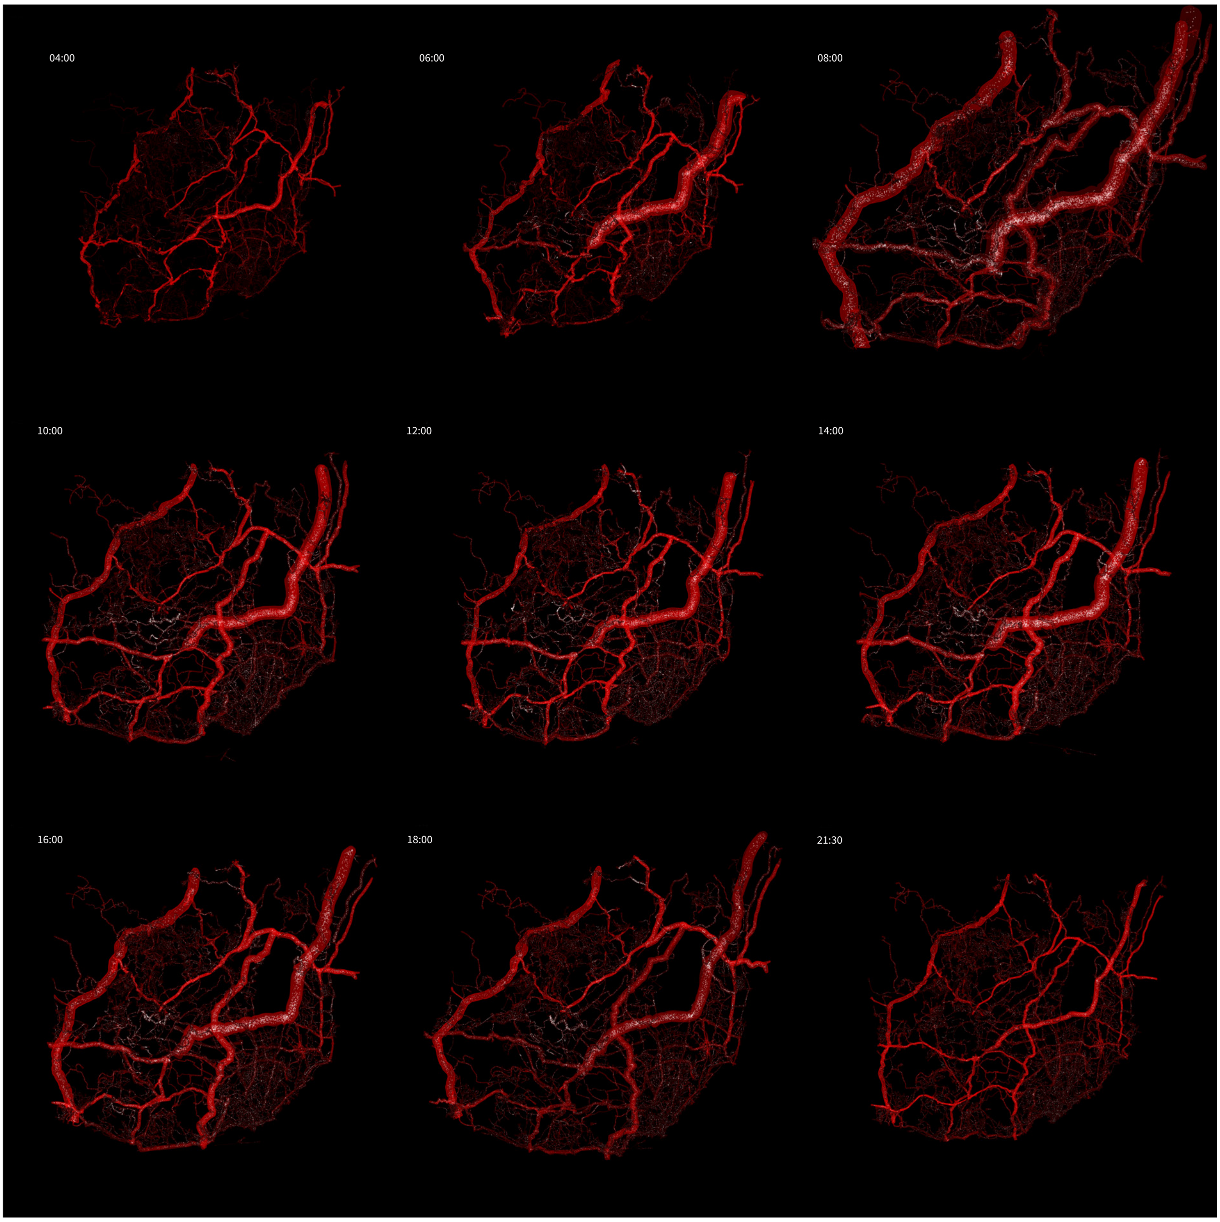

Figure 5.31 Snapshots of the blood vessels applied to the cartogram for several times of the day for the configuration ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$). Merging the two visualization emphasizes the metaphor of the city as a living organism. The animation can be watched at https://vimeo.com/88842273

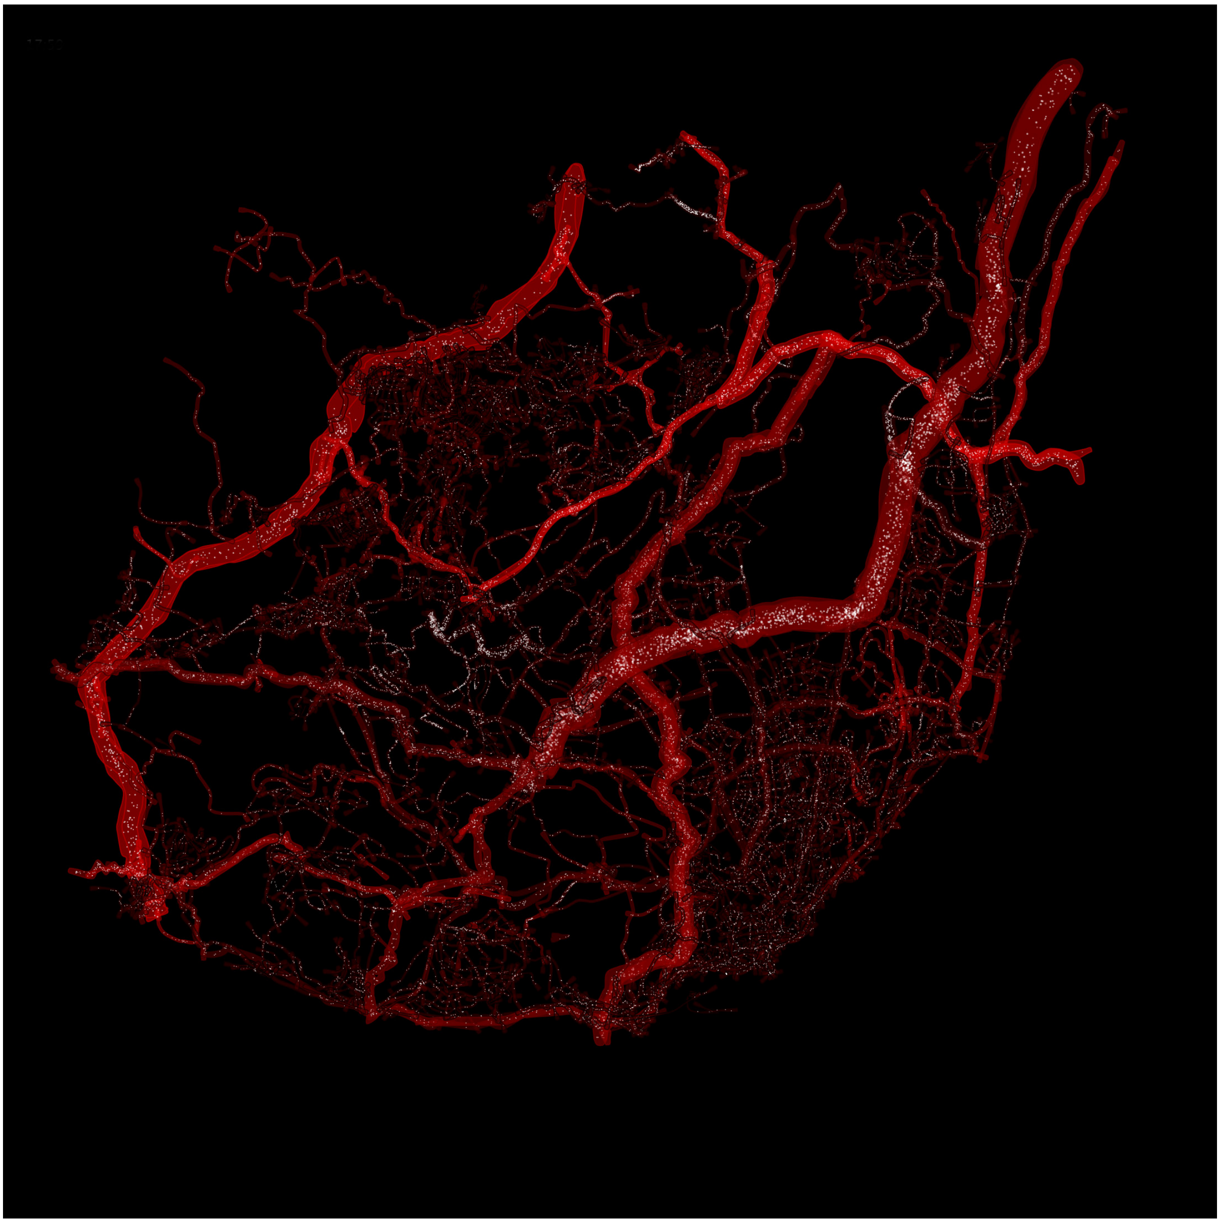

Figure 5.32 Snapshot of the blood vessels applied over the cartogram at 18:00 for the configuration ($k_B = 0.50$, $k_I = 0.01$, $k_C = 0.05$). The animation can be watched at https://vimeo.com/88842273

5.9 Testing with London

5.9.1 Data

This section presents results on the application of the cartogram model and the figurative visualization of blood vessels to another city and dataset[48]. The dataset reports vehicles' GPS traces and their speeds for a courier service in the city of London for every day of July 2010. The data is significantly denser than for Lisbon, having a total of about 12 million GPS registries for one month, when Lisbon's dataset had about two million registries. Nonetheless, this does not mean that London's data is more representative of its road traffic than Lisbon's. For example, London's dataset accounts for only 141 vehicles, contrasting with Lisbon's 1,534 vehicles. It was noticed that London's dataset indeed reports a much higher time rate of GPS traces for each vehicle than Lisbon's.

Graph 5.20 shows the number of vehicles with active registries throughout each day. The weekends can be clearly spotted as the number of vehicles abruptly diminishes along with the regular usage of the company's vehicle fleet at around 110 vehicles per day. It should be noted that the number of active vehicles approaches zero in the evenings. When displaying the average velocities throughout the days as in graph 5.21, no discernible pattern can be easily noticed. This can be explained by the low number of vehicles that are active in the dataset during evenings (see graph 5.22), which consequentially express neither a regular nor a robust profile for London's traffic behaviors.

Graph 5.20 Number of active vehicles throughout each day of July 2010. The weekends can be clearly spotted and the company's vehicle fleet has a usage of around 110 vehicles per day.

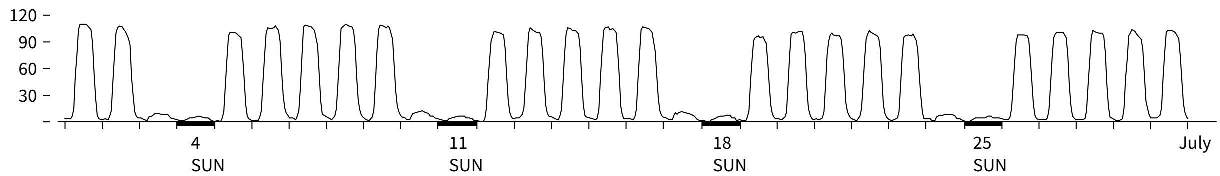

Graph 5.21 Average velocities throughout each day of July 2010.

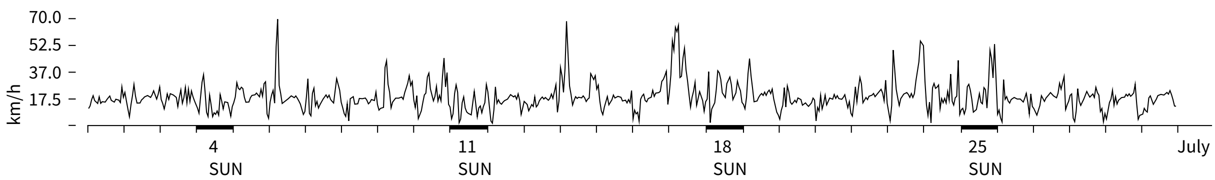

48 London data was provided by eCourier – www.ecourier.co.uk

Graph 5.22 Number of active vehicles by time of the day, juxtaposed for every day of July 2010. It can be observed that during the evenings the number of vehicles is not significant enough to extract robust measures for the global traffic behavior of the city during those periods. During these periods the number of circulating vehicles can be as low as three.

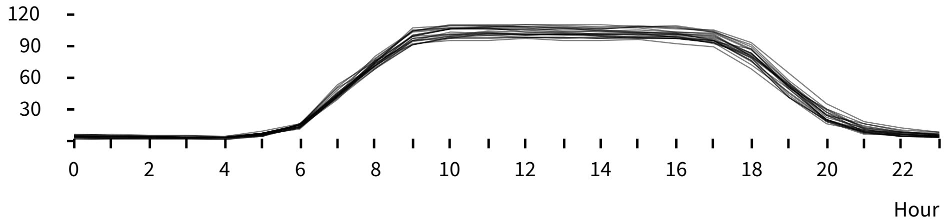

When applying the running average of velocities to this data as it was for the visualization of trajectories of Lisbon, the result is the profile in graph 5.23. London's traffic velocities are generally lower than for Lisbon. This profile does not display as clearly as Lisbon the local minima of velocities that may express its rush hours. This may be due to fewer vehicles traveling through the city, conferring a less robust estimate of global traffic conditions, or it may be due to London's global traffic having its own signature, as any other city. An irregular profile can be observed in the evenings, until velocities consistently drop from their peak of ≈30 km/h at 06:00 to less than 20 km/h after 09:00. The irregularity of the evening's patterns is due to the lack of vehicle registries for these times of the day. There is a global minimum around 09:00, but as much as this tends to fit with morning rush hour behavior, it should hardly be asserted as such based on the data. In fact, when the velocities decrease, the number of active vehicles increases, starting their day in an already busy hour, but missing to register with a considerable number of vehicles the profile that leads to it. With most of the vehicles in the city, traffic speed remains considerably stable throughout the day, but after 18:00 it increases to around 25 km/h for four hours. After 22:00, the irregular pattern due to the lack of GPS registries, returns.

Graph 5.23 Running average of velocities for the last 30 minutes of each time of the day. The velocities decrease at the start of the day and rise when the day ends.

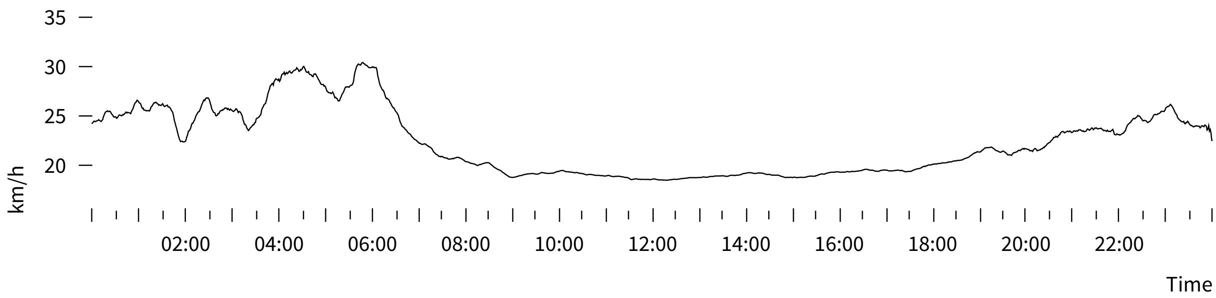

5.9.2 Testing the cartogram

Even if this dataset is not the best indicator of London's traffic conditions, with abrupt variations that are mostly due to the lack of vehicles circulating during certain hours of the day, it remains a good test object for the edge-based cartogram model. The abrupt variations in the evening translates into what seems to be randomly scattered roads in London compressing and distending. Using this data enables the testing of the model's stability and performance in terms of errors, as well as to observe how the model deals with less robust or absent data. The data was aggregated as it was for Lisbon, condensing July 2010 into a single virtual day, with trails of 30 minutes, having with this aggregation the profile of velocities seen in the running averages in graph 5.23.

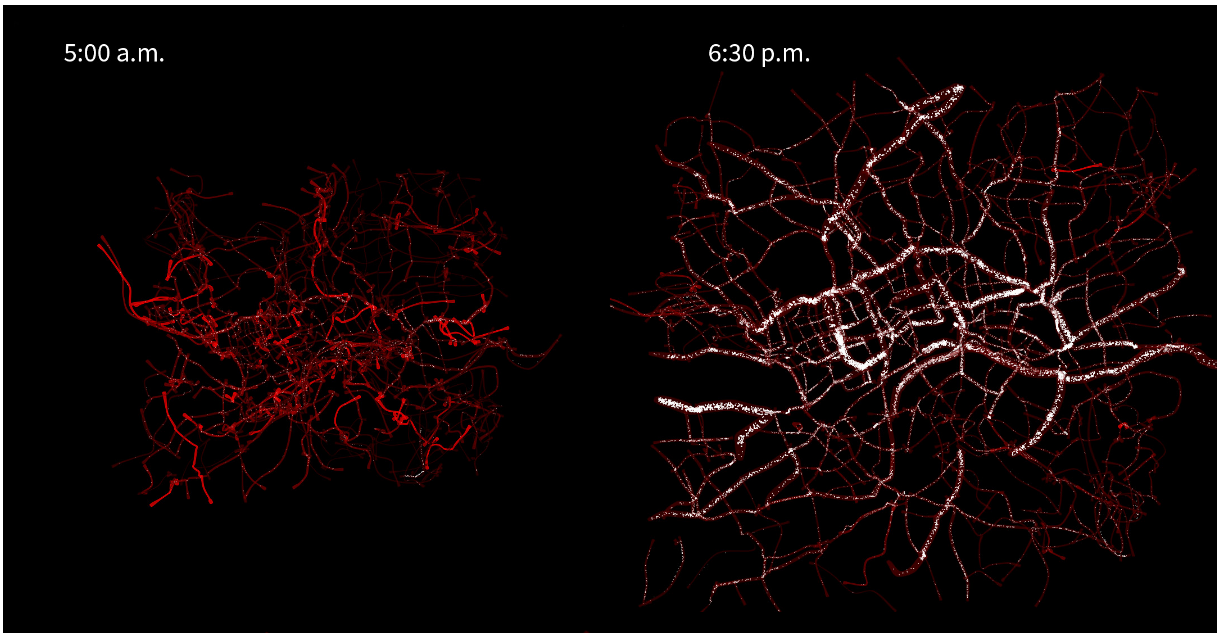

Figure 5.33 Two snapshots of the blood vessels on the cartogram for London for the configuration ($k_B = 1.00$, $k_I = 0.10$, $k_C = 0.02$). The animation can be watched at https://vimeo.com/135803140.

Furthermore, despite the data's questionable significance for the reality of London traffic, it shows a city that has higher velocities in the evening, followed by a two-hour decrease in the average velocities that tend to stay stable through the day. The velocities increase again in the evening, mainly after 18:00, but through a more stretched period of four hours. With this, the city is expected to start compressed, greatly distend during the morning, and return slowly to its more compressed state during the night. This can be observed in figure 5.33 with the blood vessels visualization applied over the cartogram. Figure 5.34 shows a snapshot for the same configuration of 5.33 in the afternoon, one of the busiest times as reported by the data. Several major roads in figure 5.34 can be observed as congested, mainly in the city center. Since the velocities are generally smaller than for Lisbon, and the number of vehicles per road differs as well, the color scale of the blood vessels and the density of cells had to be adjusted in order to appropriately reflect them – e.g. maintaining the original color scale would render vessels almost invisible, considering how dark they were; the number of vehicles, although globally fewer than in Lisbon, were more concentrated in certain roads, making undesirably thick vessels that occluded their neighbors and rendered an extremely dense stream of cells that, unless appropriately parametrized, remove the patterned appearance of the streams, becoming opaque white. The blood vessels visualization does not resemble an individual organ anymore. Nevertheless, London's has an organic disposition of roads and hence the blood vessels visualization is able to maintain its metaphor of road traffic as a living organism, with the system depicting more an intricate tissue that expands and compress, with clotting problems in its center.

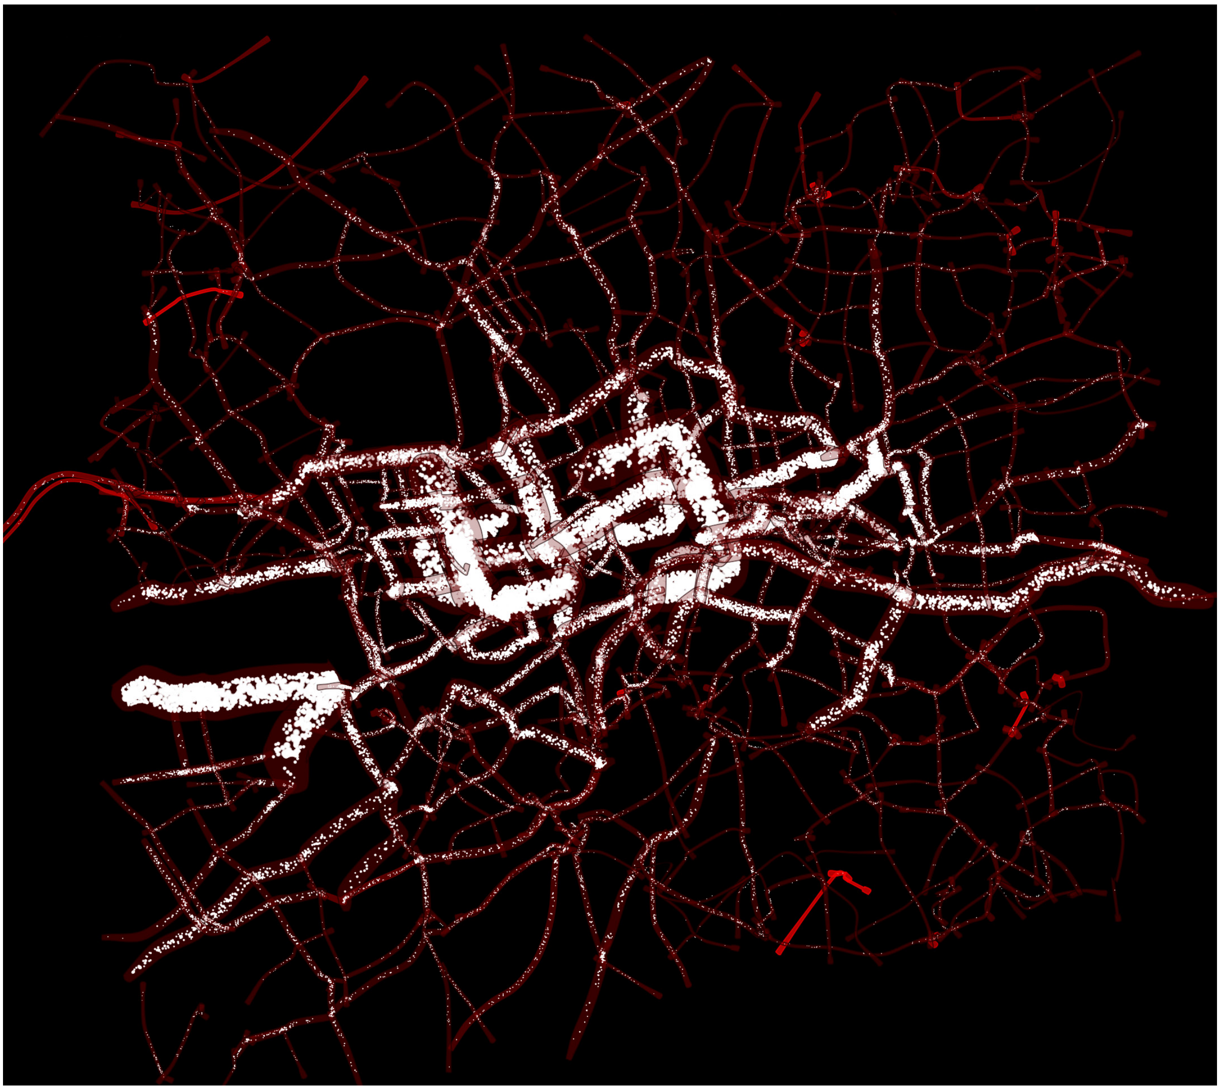

Figure 5.34 Snapshot of the blood vessels on the cartogram for London for the configuration ($k_B = 1.00$, $k_I = 0.10$, $k_C = 0.02$). The animation can be watched at https://vimeo.com/135803140.

In order to test the robustness of the cartogram model in handling other datasets while also finding the best conjugations of spring strengths, the same test of experiments was conducted, but this time for London's dataset. The cartogram was simulated 1728 times for every conjugation of the values in the set (0.00, 0.01, 0.02, 0.05, 0.10, 0.20, 0.50, 1.00, 1.10, 1.20, 1.50, 2.00) for k_B, k_C and k_I. Initial results for the errors of these experiments are reported in table 5.5. The errors are generally well controlled, with averages and standard deviations similar to Lisbon's (Lisbon's results can be compared through 5.2). The average Ēvalue for London is 0.120, higher than Lisbon's 0.100, and significantly more dispersed with a standard deviation of 0.087 to Lisbon's 0.072. Even higher, is the average Ēshape at 0.297 to Lisbon's 0.246. On the other hand, the average of $\overline{E}_{dist}$ is significantly lower at 0.134, when Lisbon's has 0.152. This difference contributes to maintaining the quality of London's solutions in a similar way to Lisbon's – the top solutions have similar values of Q

Table 5.5 Descriptive statistics of the errors. Minimum, maximum, average, and standard deviation for $\overline{E}_{value}$, $\overline{E}_{shape}$, $\overline{E}_{distance}$, considering the 1728 experiments for London. *Values that exclude the special case ($k_B = 0.00$, $k_I = 0.00$, $k_C = 0.00$) where both the shape and distance errors are zero.

	min	*max*	*avg*	*std dev*
$\overline{E}_{value}$	0.004	0.582	0.120	0.087
$\overline{E}_{shape}$	0.154*	0.560	0.297	0.081
$\overline{E}_{dist}$	0.082*	0.457	0.134	0.046

Every experiment for London except the ones with null springs' strengths is presented in graph 5.24 by their maximum E_{value}, E_{shape}, and $E_{distance}$. The resulting distribution can be compared with graph 5.16 for Lisbon. The pareto frontier for experiments with non-null spring strengths is plotted in graph 5.25. The Spearman's correlation coefficients among the three errors are described in table 5.6. It can be observed that the shape errors are usually higher for London's cartograms – the smallest maximum E_{shape} for an experiment in graph 5.24 is 0.17 while in Lisbon it is 0.13. As expected from table 5.5, the distance errors are lower than Lisbon's, having graph 5.24 a max($E_{distance}$) of 0.11 while Lisbon has 0.13. It is also worth noting that in the case of London smaller distance errors are also more frequent – that can be observed by the number of red circles in graphs 5.24 and 5.16. There are significantly more experiments with max($E_{distance}$) in the range [0.11, 0.13] in London, than in the range [0.13, 0.15] in Lisbon.

Table 5.6 Spearman's correlation coefficients among the three types of errors in experiments with non-null strengths for London.

	ρ
max(E_{value}), max(E_{shape})	0.82
max(E_{value}), max(E_{dist})	0.46
max(E_{shape}), max(E_{dist})	0.28

Another characteristic that distinguishes these results from Lisbon's are the different variations among the three types of errors for both scenarios. Just like in Lisbon, as the value errors diminish so do the shape errors, for most cases. Notably the profile of the best distance errors for London, diverge from Lisbon's case. In Lisbon, as experiments have smaller value and shape errors, the possibility of having small distance errors increases as well, until a certain point where the distance errors start being higher again. On a different profile, the best distance errors for London can be found while observing increasing value errors but decreasing shape errors. The best distance errors in London are less related with good value errors. They are more dispersed, reporting even good values for distance errors when the errors are bad for the value and shape errors, being that some of these bad shape errors are some of the worst. This is due to the abrupt changes in velocities of certain roads caused by fewer vehicles during evenings. Those roads have to drastically and momentarily expand and compress in order to accommodate their new lengths, but being restrained by other connected roads that have no registered velocities and hence do not follow their behavior. These constraints and the need to accommodate abrupt variations cause the value errors to rise when backbone springs are resized. The stretching factor of the connective springs also weights the surrounding roads that they connect, and if these other roads for the most part do not change their lengths, neither will their associated connective springs. The connective springs are strongly associated with the distance error, and if they tend to sustain their lengths, the distance errors tend to be maintained in low values. The abrupt variations caused, for example, by as few as three vehicles on three different and unrelated roads in London, probably causes the dispersion of good distance errors in solutions with some of the worst value and shape errors. Another reason for this dispersion takes the different topological properties of London into account. The mapped portion of London represented here has more roads than Lisbon, forming a more complex mesh of springs. This makes the situation of areas of the city not scaling appropriately more frequent. In a more complex mesh, more connective springs can connect a high number of segments that, unless excited with similar velocities, will not result in the connective springs' scaling and hence regions cannot be resized in order to accommodate the new roads' lengths with minimal errors. These more complex topological arrangements, together with the sparse and abrupt data variations, make the constraints harder to solve and jointly contribute to high value errors and the dispersion of good distance errors to bad overall solutions.

Graph 5.24 Scatterplot of experiments with non-null spring strengths. Maximum value error on x, maximum shape error on y, and maximum distance error by color. The maximum value and shape errors generally increase together. The best maximum distance errors are scattered through good and bad maximum shape and value errors. Several candidates for the best solution can be identified in $x \in [0.05, 0.15]$ and $y \in [0.20, 0.25]$.

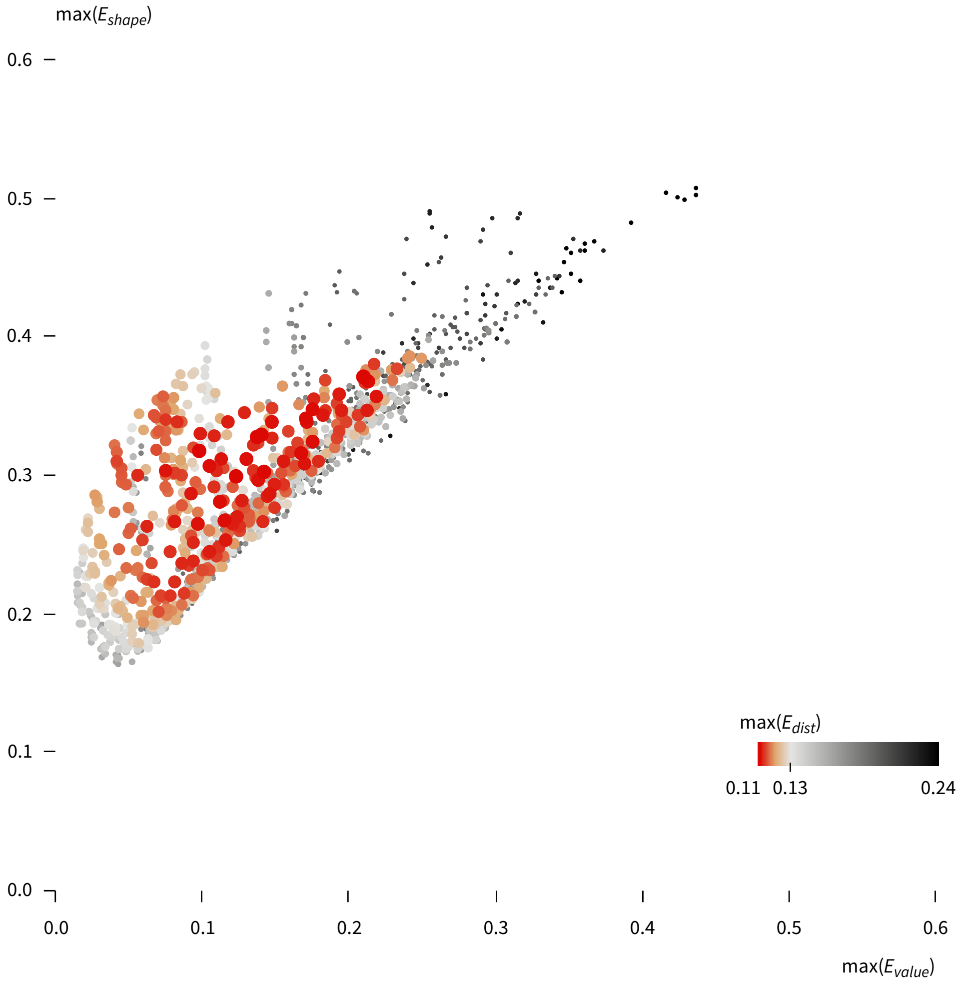

Graph 5.25 Pareto frontier of experiments with non-null spring strengths for London. Maximum value error on x, maximum shape error on y, and maximum distance error by color.

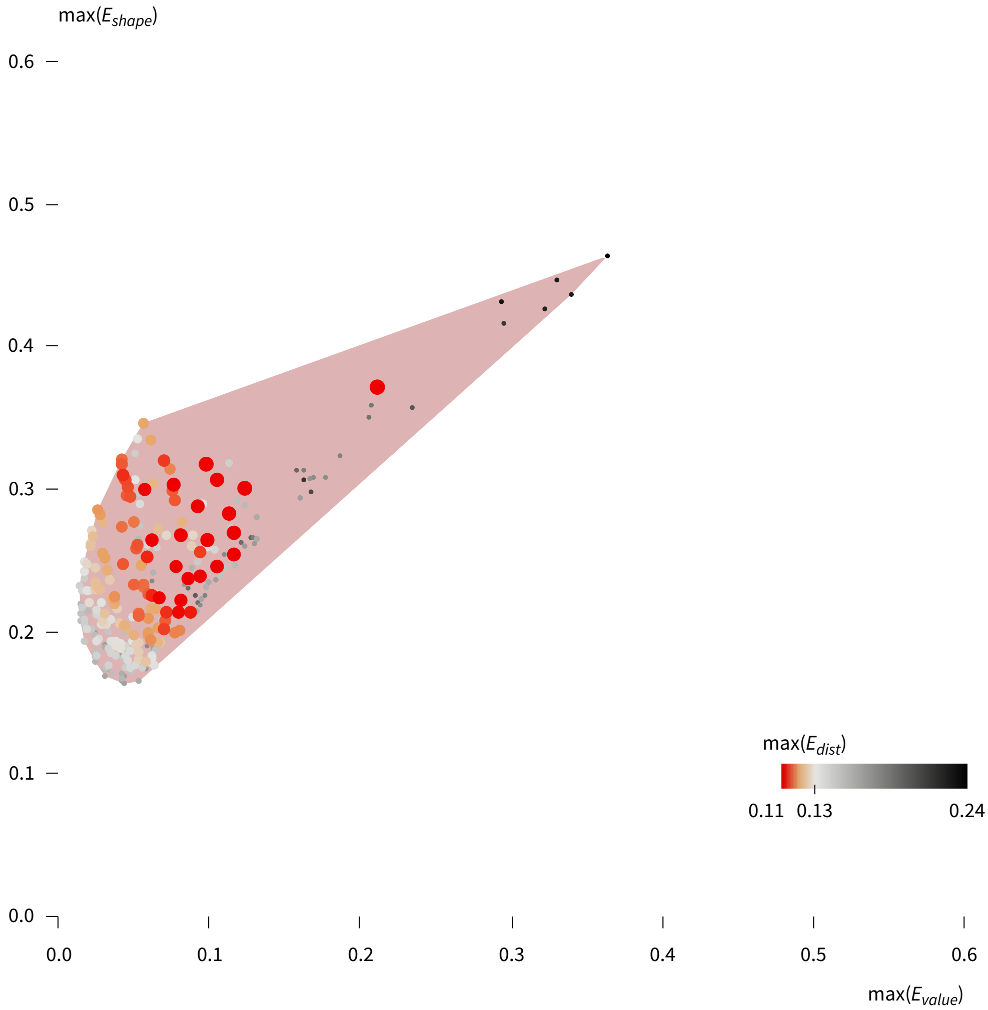

Following the same methodology as Lisbon's cartograms, a value Q for the quality of the solutions and their area variation was computed. The results for the top one hundred solutions with the best Q for London are in table 5.7. When comparing with the results for the Q of Lisbon's experiments in table 5.4, it can be concluded that according to Q the solutions for the London cartogram are similar but worse than Lisbon's: $Q \in [0.60, 0.64]$ for Lisbon to $Q \in [0.64, 0.69]$ for London. Naturally, due to differences in data and topology, the best parameterizations for the springs' strengths also differ. London also tends to have good solutions for stiff backbone springs ($k_B > 0.5$), but contrary to the Lisbon case, it tends to have some of its best solutions for stiffer inner springs and looser connective springs than the ranges of strengths that accounted for Lisbon's best solutions. In more interconnected cities like London, the number of connective springs will be higher and hence they do not need to be as stiff to produce the same effects

The area variation is substantially lower than Lisbon's. While the ranges of area variation for Lisbon's best results were between 68% and 106%, London shows variations only from 18% to 86%. Its best solution is ($k_B = 1.00$, $k_I = 0.05$, $k_C = 0.02$) according to Q and has an area variation of 72%. Other solutions exist with Q values close to the best Q but with a varying set of area variations, such as 30%, 38% and 40%. The lower area variation for London is due to the low variation in London's traffic velocities, as observed in this specific dataset. The running average of velocities throughout the day does not incur on variations that are much higher than 10 km/h, while Lisbon's dataset depicts variations in average velocities of about 30 km/h.

Table 5.7 The one hundred best conjugations of backbone, inner and connective springs strengths. These parameterizations have the smallest values of Q, and have a corresponding area variation (*Q* ▪ *Area variation*).

		k_B							
k_C	k_I	0.05	0.1	0.2	0.5	1	1.1	1.2	1.5
0.01	0.01			.68 ▪ .32	.67 ▪ .37	.68 ▪ .49	.69 ▪ .31		
	0.02	.68 ▪ .30	.65 ▪ .29	.66 ▪ .44	.64 ▪ .40	.65 ▪ .51	.65 ▪ .38	.69 ▪ .31	
	0.05		.69 ▪ .35	.67 ▪ .45	.65 ▪ .52	.66 ▪ .55	.66 ▪ .41	.68 ▪ .43	
	0.1				.67 ▪ .56	.67 ▪ .66	.66 ▪ .44		
	0.2					.68 ▪ .86	.68 ▪ .33		
0.02	0.01		.69 ▪ .41	.68 ▪ .31	.68 ▪ .22	.68 ▪ .70			
	0.02	.69 ▪ .25	.65 ▪ .48	.65 ▪ .34	.66 ▪ .26	.65 ▪ .65	.67 ▪ .30	.67 ▪ .37	
	0.05		.67 ▪ .53	.66 ▪ .39	.66 ▪ .27	.64 ▪ .72	.66 ▪ .41	.66 ▪ .22	.68 ▪ .40
	0.1			.69 ▪ .41	.67 ▪ .29	.65 ▪ .75	.66 ▪ .28	.66 ▪ .32	
	0.2					.66 ▪ .83		.68 ▪ .26	
0.05	0.01				.69 ▪ .26				
	0.02		.68 ▪ .26	.67 ▪ .39	.67 ▪ .18	.69 ▪ .33	.67 ▪ .32		
	0.05			.65 ▪ .45	.65 ▪ .27	.66 ▪ .33	.64 ▪ .30	.65 ▪ .39	.68 ▪ .47
	0.1			.68 ▪ .49	.66 ▪ .23	.65 ▪ .41	.64 ▪ .38	.66 ▪ .45	
	0.2				.67 ▪ .27	.67 ▪ .46	.66 ▪ .41	.68 ▪ .35	
0.1	0.02			.68 ▪ .26	.69 ▪ .34			.69 ▪ .49	
	0.05			.67 ▪ .31	.67 ▪ .29	.66 ▪ .30	.66 ▪ .31	.66 ▪ .40	
	0.1			.68 ▪ .35	.66 ▪ .34	.66 ▪ .28	.66 ▪ .29	.66 ▪ .47	.68 ▪ .44
	0.2				.69 ▪ .34	.68 ▪ .24	.67 ▪ .35	.66 ▪ .64	.68 ▪ .50
0.2	0.05			.69 ▪ .30	.69 ▪ .28	.69 ▪ .36	.68 ▪ .41	.69 ▪ .46	
	0.1				.68 ▪ .38	.67 ▪ .36	.67 ▪ .43	.67 ▪ .32	
	0.2					.67 ▪ .49	.67 ▪ .38	.67 ▪ .44	

Q
.64 ▬ .69

Area variation
▪ .18 ▪ .80 ▪ .40 ▪ .60 ▪ .86

5.10 Conclusions

This chapter presented several applications of metaphorical mappings for traffic visualization in the context of semantic figurative metaphors. The geographic nature of the data was maintained, and thus in order to distort the structural model of the visualizations, the cartogram technique was used. Cartograms are "dramatic" because they contrast familiar maps in unfamiliar representations of space and hence can "shock" the audience. The spatial dimensions of the city were distorted in order to provide traffic information. If traffic is congested and velocities are low, the map distends as if distances were perceptually farther. When traffic velocities are high, the map compresses in the affected areas. This employs a time-distance metaphor for the structural model.

The datasets used consisted of GPS traces of several vehicles with velocity information. Since vehicles circulate on the road network, and the road network is being affected, a new model to build cartograms was devised. This model, in contrast with typical area-value cartograms, is an edge-based cartogram, which distorts the lengths of the edges according to the data associated with those edges, instead of distorting areas with area-based data. The cartogram model is contiguous as it preserves topology, and it is dynamic as it can interpolate and continuously adapt to new data states, resulting in cartograms that can be animated in simulation-time. This contrasts with most of cartogram applications in the literature, which only aim at providing a static version of the cartogram and take higher computing power (Keim et al. 2004).

The devised cartogram model is based on a mesh of simulated springs, and accommodates three types of springs, each one with its own function:

- Backbone springs that form the roads in the network and are used to preserve topology.
- Inner springs for each road in the network that tightly connects the vertices in the road and are used to preserve the overall shape of the roads and their internal angles.
- Connective springs connect vertices of one road to the neighboring vertices of a neighboring road, playing a major role in maintaining the distances among roads and thus preserving the overall shape of the map as well as its recognizability.

As with any contiguous cartogram, the edge-based cartogram shows several errors when representing information: the value error that accounts for the difference between the expected and the effective length of each road segment; the shape error that accounts for deviations in the overall shape of each road; and the distance error that accounts for the position of the roads in relation to one another. The cartogram model is parameterized with different strengths for each type of spring, and depending on this configuration, on the the road network and on the the data itself, it can have varying results for the errors and the quality of the solutions. The model was tested for two datasets: one for Lisbon and another for London, resulting in a total of 1728 tested configurations each. The model can provide good results, with minimal value errors, shape and distance errors that render the city recognizable. The average value error for all roads can be as low as 0.002 among the 1728 tests for Lisbon, with an average of 0.100 and a standard deviation of 0.072. The London tests had a low of 0.004, an average of 0.120, and a standard deviation of 0.087 for the same error. Although the results for the errors of an edge-based cartogram are not directly comparable with results of value-area cartograms, other studies indicate that these values for the value error are low and satisfactory (Keim et al. 2004), since a similar way to account for the value error was utilized.

Regarding the role of each type of spring in affecting the errors, the backbone springs are able to control both the value and shape errors, but a correlation with the distance error could not be observed. The inner springs can control the shape error but not the value and distance errors. The connective springs control the distance error, but not the value and shape errors. The best solutions for each case were determined by using a quality measure that weights the value, shape, and distance errors of each solution. Generally, good solutions could be observed with stiff backbone springs ($k_B \geq 0.5$) and looser inner and connective springs ($k_I \leq 0.2$ and $k_C < 0.2$). Although good solutions can be found when $k_C > k_I$, this is not a general rule and exceptions can be found for the top hundred solutions with the best quality that have $k_C < k_I$.

The conjugations of strength values that work better for each edge-cartogram application varies with the topology of the city for factors such as the size and segmentation of the mesh of springs. Naturally this varies with the data, its density, and magnitude of values. This meant that the optimal configurations for Lisbon and London are different. For example, it was found that the data itself affected the patterns of errors in the two cities. In Lisbon's case, having small value and shape errors increases the possibility of having small distance errors as well, until a certain point. This could not be found in London's case, where the best distance errors are dispersed through bad value errors and some of the worst shape errors. Nonetheless, even if the datasets and topology of the springs provided different error profiles for each case (most notably, London has generally higher value and shape errors but has lower distance errors), the overall quality of the solutions as provided by the devised metric is similar.

The simulated model depicts an animated day of virtual traffic for each city. For Lisbon, the city starts compressed in the evening, abruptly distends during the morning rush hour, and compresses again to a stable arrangement that depicts the city's usual traffic conditions. The city distends again during the afternoon's rush hour, and from there it progressively shrinks during evening. London is also presented as a compressed city during the evenings and distends when the vehicles in the data start circulating more frequently, compressing again when the vehicles stop reporting their locations during the late afternoon. Although London's dataset does not allow inference into its rush hours, even if they were noticeable, the cartogram model was able to successfully show the traffic conditions as reported by the dataset. The animated cartograms show local behaviors with individual roads compressing or distending in a certain area, but also emerging global behaviors such as the city compressing or distending as a whole.

Two separate visualizations were developed for Lisbon's dataset: the visualization of trajectories and the blood vessels visualization. The visualization of trajectories uses a metaphorical color encoding on par with the concept of "visualization without reduction" (Manovich 2011), as it tends to display the data as it would represent the real world, showing moving vehicles as moving vehicles, displaying a system closer to reality, which intends to connect with the viewers in a context of casual visualization. Inspired by the aesthetics of motion of the cartogram, the blood vessels visualization was created to portray the city as a living system of pulsing blood vessels. The blood vessels visualization is less abstract and more figurative, depicting Lisbon as a set of blood vessels through which traffic flows. Having thicker vessels for roads with more traffic is a much stronger way to draw attention to them than relying on the overlay of semi-transparent GPS traces. For Lisbon, the blood vessels metaphor works particularly well, since it resembles an organ or a human heart.

Until this point, the cartogram model explored the modification of the structural metaphor while the other visualizations brought visual cues to be implemented on top of the cartogram. Therefore, both models were applied over the cartogram in order to advance the implementation of semantic figurative metaphors for these cases. The visualization of trajectories on the cartogram provides an additional layer of information as seen causing such distortions. The application of the blood vessels on the cartogram is graphically strong and is figuratively close to its metaphorical domain – that is why it is a semantic figurative metaphor at its fullest. Nonetheless, it is also a less elegant artifact in a minimal sense, where the authorial intent superimposes the functionality of data portrayal. The motion of the vessels and the cells, together with the expansions and contractions, add visual behaviors that are unrelated to the data, and are perhaps too distracting. With this experiment, the limits of eccentricity in semantic figurative metaphors for visualization have perhaps been reached. Although permissible in this context, there is a point when the beauty of complexity might turn into visual confusion, and such should be avoided and well framed within specific authorial intents.

The blood vessels visualization was also applied to London's cartogram, in order to better depict how the cartogram model can work for other cities, but also as way to explore the suitability of this metaphor for other traffic data and topographical arrangements. The blood vessels for London show a much more congested city on a few roads with slower velocities than Lisbon's, which is directly comparable for both artifacts. In relation to how well the metaphor suits London, it is now farther from an individual organ, and closer to an organic tissue of vessels with clotting problems. Nonetheless, the metaphorical intent of having a city figuratively portrayed as a living organism was not lost.

6

Corporate politicians and satire

Nowadays in Portugal there is a significant mistrust of the country's political class. Voting turnouts have been around 60% in legislative elections since 2009 and suspicions about political patronage and corruption are constantly present in the media. Lobbying in Portugal is illegal and unravels in a parallel stage, being widely unpopular and cast in a harsh light by the media. The high availability of public information regarding this matter has raised awareness that fuels discontentment. Within this context, the work described in this chapter intends to make this invisible matter more visible, tailoring a provocative visual form that would engage viewers through a serious but satirical portrayal of the issue.

The design and results of an interactive visualization that displays the relations between Portuguese politicians with government positions and companies is presented in this chapter. The aim of this visualization is to create awareness of these relations and, by using a highly figurative approach, captivate a broad audience. The *Ecosystem of Corporate Politicians*, or in its Portuguese variant *Um Ecossistema Político-Empresarial*, as it was called, depicts a set of organisms that chase and frenetically jump between companies.

6.1 Related work

The high availability of public information on the relations between high-ranking government members and the corporate world has raised awareness on this matter all over the world. Several platforms have been developed in order to make this information reach a broad and global audience. Behind the recent "Poderopedia" (2012), is a team of journalists, developers, designers, and collaborators interested in promoting transparency, accountability, and democracy in Chile, Colombia, and Venezuela. "Poderopedia" is an index and database of the major actors in businesses and politics, but makes sparse use of visualization for the mapping of those relations (see figure 6.1). Although lacking in visual power and communication, it is a valuable collaborative platform with a rich dataset of hundreds of individuals and companies and thus a vast index for data journalism. This will increase even more as they plan to release an API to access their content.

Figure 6.1 network map of Rafael Guilisasti in *Poderopedia*[49]. Rafael Guilisasti is placed at the center and radially connected with other companies and individuals – relations based on family ties, friendship, companionship, or sharing of administrative positions. The platform does not facilitate the browsing of individuals or companies based on their relevance in the global network, neither does it try to map a global network. The only type of mapping observed is based on a central figure (individual or company), and does not involve any visual variable to distinguish the type of relation established.

"Retórica" (rhetoric), from 2013, is a visualization of the thematic emphasis of 14,000 speeches in the Brazilian House of Representatives. The thematic and statistical analysis of the speeches resulted in 70 main themes that deputies are associated with. The themes are represented as packed circles, which in turn have a set of packed bubbles inside that represent the associated deputies (see figure 6.1). The layout resembles a circular treemap, but the positions of the bubbles are random with the exception of the relation of containment of a deputy inside a theme.

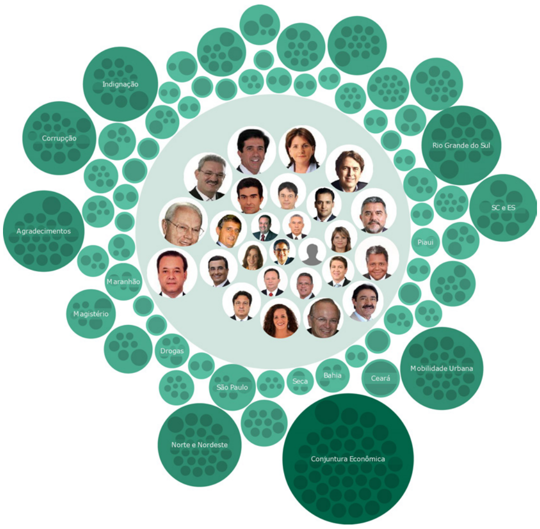

Figure 6.2 *Retórica*[50] visualization with the expanded "health" thematic. The photographs of the deputies, which are the main focus of this theme, are presented as packed circles in the center. Statistical analyses of the themes are done by Davi Moreira and Manoel Galdino. Visualization by Luis Carli.

Arguably, one of the most popular examples of visualization for transparency in businesses and politics is the project "They Rule" that started in 2001. Backed by the LittleSis[51] database, Josh On created the *They Rule* project which consists of an interactive web application that maps the connections and affiliations of US corporations and their board members. Companies and individuals are nodes in a network that can be manipulated by the user at will, adding or removing elements of a network and building, sharing, and editing networks of corporate relations. The networks have a fluid layout, enabling the users to rearrange edges and nodes – when dragging, all dependent relations are adjusted in their locations as well. More can be read on the project's website, which unveils the project's motivation:

> A few companies control much of the economy and oligopolies exert control in nearly every sector of the economy. The people who head up these companies swap on and off the boards from one company to another, and in and out of government committees and positions. These people run the most powerful institutions on the planet, and we have almost no say in who they are. This is not a conspiracy, they are proud to rule, yet these connections of power are not always visible to the public eye.[52]
>
> —The "They Rule" project

The "They Ryle" project has a well-defined and distinctive identity. It is mostly deprived from color as it would not add to the visualization semantics, opting instead for a sober tone. It is a highly figurative visualization, since nodes that represent the boards' companies are drawn as actual tables of emp-

[49] http://www.poderopedia.org

ty chairs, but connected through individuals (see 6.3). The individual's depiction is figurative to the point where one can distinguish between man and woman. As observed, the shape of the individuals, changes according to the amount of chairs they take. In fact, more chairs imply fatter silhouettes, arguably translating into a message of insatiable greed. This is an important aspect, since it is a metaphorical visual cue that was added on top of a node-link diagram in order to communicate the author's message and tone.

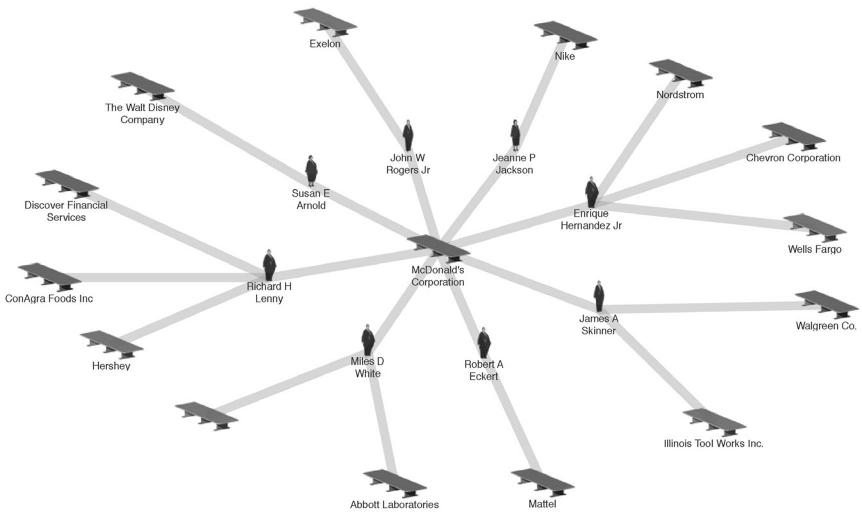

Figure 6.3 McDonald's map created by user "sh02" on 14 November 2011. The McDonald's Corporation is placed at the center and interconnected with several board members from other companies' government bodies. The map can be seen at http://theyrule.net/.

This chapter focuses on the relations between members of Portuguese governments and their involvement with companies. Nowadays in the Portuguese sphere, such a theme does not lack infrastructure. For example, "Tretas"[53] is a collaborative wiki that aggregates dossiers about important matters in politics and economics in Portugal as well as extensive curricula of Portuguese politicians. "Tretas" has a subversive tone, trying to expose for the most part, unclear and shady relations of prominent Portuguese politicians. In the context of the work described in this chapter, it is an important reference since it was used to extend the dataset.

[50] http://retoricaparlamentar.com
[51] LittleSis is a "free database of who-knows-who at the heights of business and government" (http://littlesis.org)
[52] http://theyrule.net/drupal/about
[53] The wiki was initiated by Hélder Guerreiro and José Lopes around 2008. http://tretas.org

6.2 Data

In 2010 Costa et al. published the book *Os Donos de Portugal* (*The Owners of Portugal*) that among a historical narration of the most influential corporate families in Portugal, contained a list of politicians' résumés that held high governmental positions and later moved to influential Portuguese companies (Costa et al. 2010). For each politician, the list describes each position held in government and in companies, together with information on their political party affiliation. The collected résumés mainly cover a period from 1950 to 2010, and focus on ministers and secretaries of state in strategic sectors (e.g. finance, economic, and public service). The résumé total is 115, describing each position in government and companies, with the corresponding time period and political affiliation. The companies are public or private, often corporate groups, and politicians are usually part of the administrative board of such companies.

Although not adding any new politician, the first step to augment this dataset extended it temporally to the end of 2013. In order to do so, each of the 115 politicians was investigated to check if their positions were still held in 2013. If not, a search was held for the new positions since 2010 and incorporated accordingly. The history of party affiliations was also verified. The results of these searches were not always conclusive, as the investigation was strictly based on publicly available information, such as major companies' governing bodies, companies' public reports, online public résumés, press, and news sites. It is therefore possibly to lose track of some politicians if they moved to smaller companies whose governing-body information could not be found, missed an update to their résumés, or simply were not documented by the press.

An additional study of Bianchi and Viana (2012) reported on political connections in listed companies in Portugal, considering that there is a direct political connection when managers, administrators, or major stakeholders of companies are connected with political parties. They can be politicians, public employees, or people close to political power. The politicians covered in this study were also added and merged with the dataset, conducing the appropriate searches in order to gather the maximum amount of information for each résumé. Finally, two other politicians deemed relevant with connections to businesses were added – the then Portuguese Prime Minister, Pedro Passos Coelho, and the then President of the Portuguese Republic, Aníbal Cavaco Silva. The dataset resulted in a total of 130 politicians, spanning from 1950 to 2013, accounting for a total of 906 corporate positions of every politician.

The data was manually transcribed[54] from PDFS and online pages to the JSON[55] format and aggregated by politician and company. The aggregation of politicians and companies had to deal with the normalization of politicians' names. The aggregation of companies itself resulted sometimes in merges of small companies into relevant corporate groups that were verified as being in full control of those companies. The data was made publicly available[56].

The augmented dataset resulted in 354 different companies that were frequented by an average of 2.09 politicians, with a median of one politician and a standard deviation of 3.26. Table 6.1 shows the top ten companies that had the most politicians.

Table 6.1 Top ten companies with the most politicians.

Company or group	*Politicians*
CGD	26
Champalimaud	23
Banco de Portugal	23
PT	23
BCP	20
Mello	18
EDP	18
Galp	16
BES	13
BPI	11

Politicians have an average of 6.97 positions in companies each, with a median of five positions and a standard deviation of 6.27. The average of single companies frequented by a politician is 5.69, with a median of four companies and a standard deviation of 5.58. It is important to notice that the average amount of companies is not much lower than the average amount of positions, implying that every position of a politician is usually scattered among different companies and is not reserved to one. For example, consider the table 6.2 that shows the top ten politicians by number of positions and their corresponding number of different companies.

Table 6.2 Top ten companies by positions and different companies.

Politician	*Positions*	*Companies*
Ângelo Correia	55	51
Luís Todo Bom	25	20
Daniel Proença de Carvalho	21	20
António Couto dos Santos	19	16
Fernando Faria de Oliveira	18	14
Álvaro Barreto	18	15
António Nogueira Leite	16	10
Artur Santos Silva	16	12
Carlos Tavares	15	11
António Almeida	15	13

[54] The transcriptions, searches and investigation that enabled the extension of the dataset were done jointly with Cátia Costa. http://cdv.dei.uc.pt/authors/catia-costa%E2%80%A9/

[55] JSON is a lightweight data-interchange format, described in the standard ECMA-404. http://www.json.org

[56] The dataset was made publicly available through a Creative Commons License (Attribution 4.0 International): http://creativecommons.org/licenses/by/4.0/. The politicians résumés can be accessed here: http://pmcruz.com/eco/js/worms/data/worms_12jan.json; and the file that provides the companies' aggregations can accessed here: http://pmcruz.com/eco/js/worms/data/grouping.json.

By investigating which was the ruling party during a politician's governmental position, it is possible to infer the politician's party affiliation. In some punctual cases, a politician had different party affiliations. In this case the affiliation was determined to be the most recent. The dataset showed that 58% of politicians were from the PSD, 34% were from the PS and 8% were from CDS[57]. The remaining 13% could not be determined, either because they are independent or because such information is concealed.

Regarding the distribution of the amount of politicians' positions in companies over time, it can be observed that information is clearly skewed towards more recent years (that is perhaps due to the increase of information availability in recent times), having its peak in 2010 when the study for *Os Donos de Portugal* was concluded (graph 6.1). Nevertheless the amount of contributed information for 2010-2013 in this work is not negligible. Although the gathered data is only a sample of the real dimension of traffic between companies and the government in Portugal, it is illustrative of this subject.

Graph 6.1 Number of positions in companies held by politicians throughout the years.

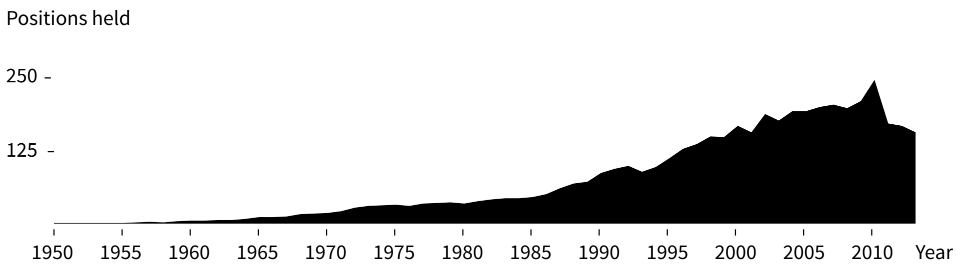

[57] PSD – Social Democratic Party; PS – Socialist Party; CDS – Democratic and Social Center / People's Party (conservatives)

6.3 Visualization

This visualization is based on the study of Costa et al. (2010), who already provided a first visualization approach that shows companies organized in a radial layout and connected through straight edges. The weight of those edges is proportional to the number of politicians that the companies shared, as observed in figure 6.4. The visualization developed in this work intended to involve the viewers more, providing an exploratory and interactive visualization with a clear graphical language complemented by subtleties that give room to the viewers to make their own cultural interpretations. The visualization is called *An Ecosystem of Corporate Politicians* (in Portuguese *Um Ecossistema Político-Empresarial*), being bilingual in Portuguese and English. The main target audience is Portuguese people interested in having a broad view of the political-corporate panorama and investigating specific relations. Every aspect of the interface is bilingual, except the ones that directly derive from the data such as positions' descriptions, which are only provided in Portuguese. In fact, translating 906 positions and 354 companies' names in English was not adequate in terms of effort given that the main target audience are Portuguese. The name "ecosystem" is a metaphor for the visualization model that was built: it represents living organisms (i.e., politicians) that dynamically jump between companies. Using such a figurative metaphor is intended to better communicate a general message and captivate the audience. This section describes approaches in building this visualization, such as message, behavior, form, interaction, and technological aspects.

6.3.1 Message

Determining the shape of the data (see graph 6.1 in the previous section) is important to justify narrative aspects in the visualization. As noticed before, politicians usually jump between different companies instead of remaining in the same one or in a few. That is a fact emphasized by the visualization. A strict chronology was not adopted since the evolution of number of positions held by politicians through time is considered biased, due to the recent increase of publicly available information. In fact, the radical increase of the number of positions depicted in the dataset in recent times would not add more to the narrative than this. The message to convey is that politicians jump frenetically between moving companies and have a hectic and incessant business life. Therefore each politician starts their business activity at the same time, in simulation time. In order to complement this message, one can extract sub-narratives from the jumps of each politician and the jumps among companies that share the same politicians.

Figure 6.4 A network of companies in a radial layout as depicted in Costa et al. (2010).

6.3.2 Structural metaphor and layout

The inherent metaphor of the dataset is a network of companies connected by politicians. Although it is true that there is not any semantic aspect to the data that could directly dictate the shape of this network, usually network-related metrics can be extracted to create a layout for the network. One popular method is the force-directed layout (Fruchterman and Reingold 1991), that positions nodes on a graph in a way that minimizes the number of crossing edges as well as the nodes' juxtapositions. This is done by establishing forces between nodes that can be related to a network metric for the corresponding edge. The system is then simulated until it reaches a state of minimum energy. Naturally, different results of force-directed layouts can arise from the same set of edges and nodes, depending on the initial configuration of the nodes at the beginning of the simulation.

The chosen structural metaphor for this visualization is a radial network (Bertin 2010, p.51). This is an authorial choice that has its limitations and advantages. In fact, the force-directed could be useful if the objective of this work was to visualize measures of the nodes' centrality in the network, or to extract possible clusters of companies that employ many politicians. Instead, a radial layout has the nodes distributed on the circumference as a more geometric configuration than an apparently disorganized and scattered set of nodes that often results from force-directed layouts. This is not in detriment to these layouts, which present a spatial arrangement based on semantic data features and have became a popular tool in visualization. The radial layout has a more organized geometry and is arguably less abstract than the resulting arrangements of force-directed layouts. This sense of organization and less abstraction is much in line with the objective to bring visualization to a more figurative level, as well as the fact that it contributes to confer a stronger identity and enhance the recognizability of the artifact. Having nodes distributed in a circular fashion alludes to a stronger sense of cohesion, as if the politicians that are jumping between these nodes are sharing the same playground. Nonetheless, the artifact also transmits a relative sense of disorganization, but this time on the politicians' layer. The radial layout generally increases the relative distances among nodes and avoids having identifiable clusters of nodes. This way, politicians that travel between nodes will often cross with each other. This results in a more complex artifact, designed to transmit a sense of hectic and incessant business activity by the politicians.

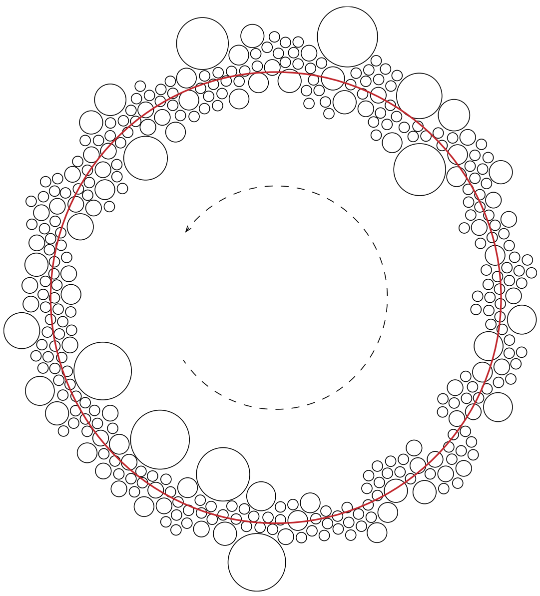

Figure 6.5 The company circles are packed in a circular belt in the center of the canvas. The circles collide with each other and have a constant rotational speed. Companies are positioned randomly over the circumference and therefore lay themselves out differently each time the simulation is started.

However, a pure geometric radial layout is sterile. It lacks an organic tone that is deemed appropriate to have a more elaborate visual appeal. Instead of having the nodes evenly distributed over a circumference, they form a belt of non-overlapping, packed bubbles (see figure 6.5). They are constantly moving, although slowly, aligned with a circumference, and organize themselves by means of collisions to avoid juxtapositions. They move in order to add dynamism to the visualization, while forcing the politicians to chase and catch those companies. Since companies have an area proportional to the total number of different politicians that had a position there, the compromise between packed bubbles distributed on the circumference and a strict radial layout, enables the economization of space due to the varying sizes of the circles. Bigger circles represent companies that had more politicians and are therefore more interconnected with other companies, emphasizing their importance in the network through size. Implementing a belt of circularly packed bubbles instead of a plain radian layout is a modification done so that the structural metaphor can accommodate the intent of the visualization's author.

6.3.3 The behavior of politicians

A politician is, before any other metaphorical considerations, a particle. Each politician has a sequence of positions in companies that they have attained. In order to do so, they visit the company corresponding to each corporate position. Therefore the simulation of a politician's movement uses a particle that has to travel in between a set of companies. Although the order of visitation is chronologically sorted, the travel times are only dependent on the particle's behavior and are not chronologically synchronized with the movement of other politicians. When iterating over the sequence of positions, the politician has the following behavior:

- Travels to the respective company following a specific trajectory.
- When a company is reached, the politician starts an encircling movement lasting an amount of time proportional to the chronological duration of the corresponding position.
- When the previous movement is over, the politician promptly travels to another company or repeats the encircling behavior if the next position refers to the same company.
- When the sequence is completed, it repeats itself. This way the system will keep running without stopping or resetting as a whole.

A politician's trajectory is influenced by two forces: one that attracts them to the company that they are traveling to, and a second that is perpendicular to the previous and always points outwards to the center of the canvas. Moreover, the traveling speed is inversely proportional to the distance of the target, diminishing as it approximates the targeted company. For example, if a politician has a set of three companies to visit, the followed trajectories are depicted in figure 6.6.

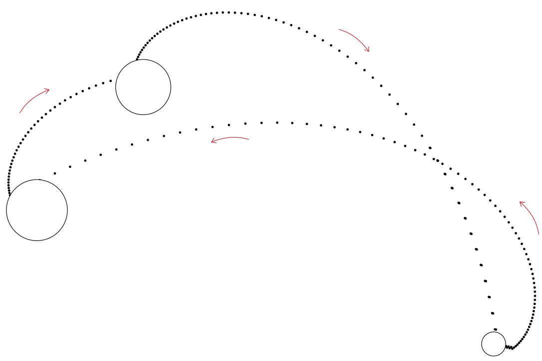

Figure 6.6 The trajectory of a politician's particle as it travels through a sequence of three companies.

The particle has the trajectory of an arc that approximates the circumference of the canvas and avoids concentrating trajectories in the center. Arc diagrams (Wattenberg 2002) and arc-based networks (Polisciuc et al. 2015) are a common technique to draw edges in networks since they contribute to declutter edges, and better distinguish one edge from another. The specific form of the arc derives from the forces in effect, but an additional modification to this arc-based radial layout is that those trajectories are not drawn explicitly but instead partially, depicted through the movement of the politician's particle and its trail. In other words, one can argue that the inherent ink part of the visualization model was erased in order to enact the metaphor of jumping politicians, depicting the connections through motion. This results in a movement that rather than crossing the central circle, entangles the circumference, intersecting its trajectories that enter and leave the circle without crossing the center (see figure 6.7).

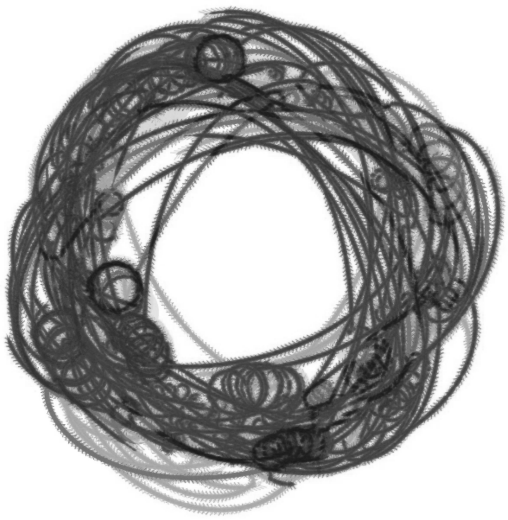

Figure 6.7 Rendering of multiple politicians' trajectories juxtaposed. The arcs do not pass through the center of the canvas and entangle the main circumference of the visualization, through the outside and the inside. Several small circumferences can be observed as the result of the politicians encircling companies for a certain period of time.

6.3.4 The politician's body

> (…) to laugh and laugh without end, mouth wide open until you can see the gullet, at the scores of grotesqueries that swarm all around us like ants in a sugar bowl.[58]
>
> —*Maria* in Pinheiro (1885, p.1)

This is how Rafael Bordalo Pinheiro opens his satirical periodical *pontos nos ii* in 1885. Pinheiro was a prominent Portuguese illustrator and caricaturist. In a later periodical called *A Parodia, Pinheiro* often illustrated political themes resorting to animals (see figure 6.8).

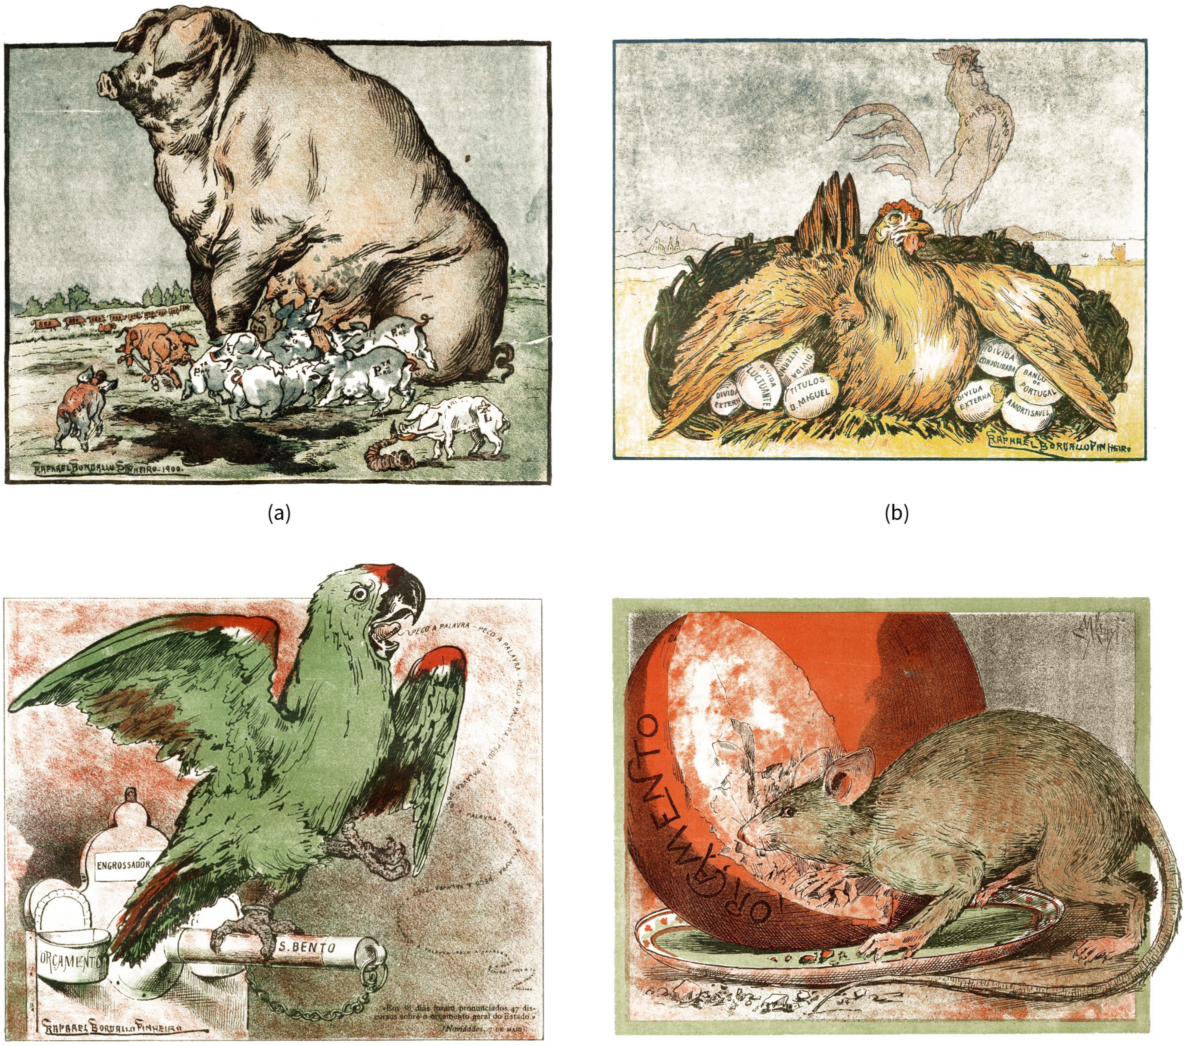

(a) (b) (c) (d)

Figure 6.8 Satirical illustrations of Pinheiro resorting to animals. (a) “Politics: the Big Pig” (Pinheiro 1900c); (b) “Economy: the Broody Hen” (Pinheiro 1900b); (c) “Parliamentary Rhetoric: the big Parrot” (Pinheiro 1900d); (d) “Bureaucracy: the big Rat” (Pinheiro 1900a).

The *Ecosystem of Corporate Politicians* also serves as a similar pictorial device. Transforming the image of politicians to animals serves the purpose of critique and satire, conferring tones of ridicule to the portrayal of this universe, and using an imaginary world to depict the theme with the strength of irony. Satire and humor can also be found in the more recent infographics of Nigel Holmes, although they are so what toned down. Holmes points this out when talking about his work.

[58] Translated from the Portuguese by Luís Coimbra “(…) rir, rir sem descanço, de bocca escancarada até mostrar o cavername, de todos os mil grotescos que por ahi fervilham como formigas n’um assucareiro.”

> I found that the best way to connect, was making people smile. Not funny, not laugh, but a smile of recognition of something. I'm trying to be friendly, to be approachable, to tap into other people's emotions, to use humor in the sense of good humor, of good feelings, of good feelings and emotions [sic]. (…) Emotions are important.
>
> —Transcription of Holmes (2015)

Using a figurative approach to depict politicians can then be a humorous and emotional way to depict information while carrying an authorial intent. It should be noticed, however, that the approach that the ecosystem takes in this direction is much subtler, and, although figurative as well, is closer to the abstract language of visualization than the more realistic illustrative approach of Pinheiro. The implementation of this is a delicate balance between soughing the satire of Pinheiro, but balancing it with the more minimalist and abstract style of Holmes.

To seek satire, the rendering of a politician took the form of an animal that could be part of a "plague", while making sense in terms of graphical language and in behavior in the context of the visualization, i.e., companies being circles and politicians having a jumping behavior. Nonetheless, the rendering could not be realistic to the point that eliminated its light, humorous effect or to the point that would distract the viewers from the seriousness of the theme. The chosen form for each politician is one of a generic, more abstract, insect, or mollusk type. The way it was rendered gives room to many interpretations, be it of an ant, a moth, a wasp, a slug, or a cockroach. The rendering is abstract to the point that makes it dependent on the viewer's own interpretations, and arguably may not even be perceived as insects nor mollusks. The emotional response of the viewers derives from their own exercise and may result in a broad set of perceptions (from the repulsive to the humorous) and more importantly, the immediate perception of the ironic tone of the visualization. From here onwards the term "insects" is chosen for the sake of simplicity and is used interchangeably with "politicians".

Choosing insects is also justified by the devised behavior that was sought for this dataset – frenetic, jumping politicians. Rendering politicians as insects represents both the modification of the visualization's structural model since edges are not drawn, while it is the most important visual cue that reflects the authorial intent of the semantic figurative metaphor. The *anatomy of the insect* politician consists of a head, a body or tail, a pair of antennas, and three pairs of legs or spikes. The silhouette of the head and body is drawn using a sine function that is mirrored on the x axis. The rest of the elements are drawn in relation to this axis (see figure 6.9).

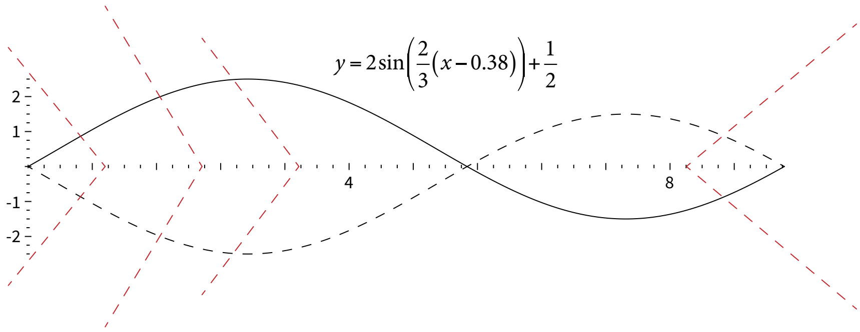

Figure 6.9 The politician is drawn using a sine wave. The factor 0.38 is an approximation to one of the solutions of 2·sin(⅔x)+½ = 0. The head is the wave with the shortest amplitude at the right and the body is the wave at the left. The legs and antennas are drawn as depicted in red.

In order to reinforce the satirical intent of this visual cue, and much in line with the work in "They Rule", the shape of the insect is related to the number of different companies that it has to visit. By varying certain parameters of the sine function, the head and the body can be made more or less distinguishable, translating into a rounder or slimmer silhouette. Politicians that have to attend more companies are depicted as *fat* and the ones that have fewer companies to attend are made *slimmer*. This characteristic of the insect also adds to their identity, enabling, for example, one to identify the politicians that have more companies in their résumés. Figure 6.10 can better illustrate this relation.

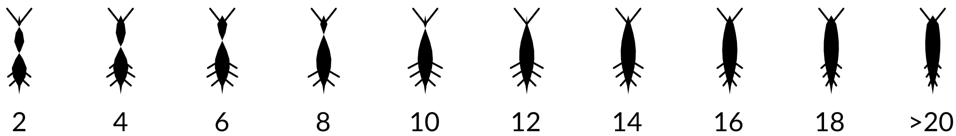

Figure 6.10 The silhouette of the insect varies in relation to the number of different companies that it has to visit.

The politicians' trajectories are arcs. Their bodies bend in line with their curved trajectories. For this, each politician's particle leaves a trail that consists of its last ten points in space. This trail is then used to map the body of a politician using the particle as the first point, and allowing to bend the politician's body on its trajectory. The jumping behavior implies a higher velocity when leaving a company than when landing the next one. This variation of velocity causes the stretching of its ten last positions, making the insect longer or shorter according to the velocity. This enables the addition of more variation to the bodies' length, emphasizing the politician when speeding up by drawing a larger portion of the traveled trajectory. This dynamism in curve and length of the politician's body conveys an organic expression that is adequate to this animalistic metaphor, and more appropriate to the invertebrate nature of mollusks. Figure 6.11 shows how politicians can have their shapes' deformed in this way.

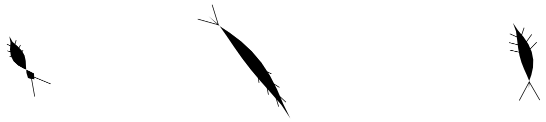

Figure 6.11 Example of three different politicians curved on different paths and at different speeds, conveying organic tones to the representation.

The politician is colored based on its latest political party affiliation, according to their customary colors in Portugal: pink for socialists (PS), orange for social democrats (PSD), blue for conservatives (CDS) and gray for the remaining. The original colors of the parties were toned down in order to have the same perceptual brightness on a white canvas. Nonetheless, those remain vivid colors that confer a sense of humor and playfulness to the visualization. A visual overview of the politicians swarming over the packed companies in a circular layout can be observed in figure 6.12.

Figure 6.12 Visual overview of the *Ecosystem of Corporate Politicians*, showing the initial state of the visualization, with politicians colored by party affiliation, and jumping from company to company.

6.3.5 Opening sequence

The visualization is an interactive web application, implemented on a single webpage. The user interacts with the visualization, pulling new contents, but without changing the page. The initial state of the visualization and the unique entrance point that every user has to go through, shows the title of the visualization and a brief description: "Interactive visualization of the relations between members of the Portuguese government and companies." The animated legend identifies each component of the visualization, depicting a company in the center being encircled by a politician (see figure 6.13). This allows the user to gasp the theme and the metaphorical encoding of the visualization when exploring. This initial screen is presented when all the data is loaded while all the graphical elements and structures for the visualization are automatically configured. After this process, the user can continue using an "explore" button that is made available.

Figure 6.13 The initial legend that is presented to the user when the visualization is opened.

After the users read the description and the legend they can proceed to explore the full visualization. The visualization has an opening sequence that cannot be skipped, as the circles are created in a radial layout, and the politicians appear sequentially from the top, swarming to visit their first companies. It is a sequence with metaphorical intent, as it does not show politicians already in their positions, but instead shows a flocking mass of insects that all at once run to meet the mass that makes the companies (see figure 6.14). While this sequence plays out, an instructional phrase appears: "Explore the ecosystem by clicking on companies and chasing the politicians". The term "chase" translates into an aspect of ludic and humorous interaction, which is explained in the next section.

6.3.6 Interaction

Interaction is an important aspect of the visualization since it contributes to the engagement of the user and adds an exploratory function to the visualization. When hovering over a company its size increases, and the name of the company is displayed. This change in size forces the surrounding companies to re-arrange in order to accommodate the new size. When the user clicks on a company, the company travels to the center of the visualization and all the companies that do not share politicians with it disappear, together with the politicians that are not related with that company (see figure 6.15). Therefore, clicking on a company places an emphasis centered on that company and the other companies that share the same politicians, displaying the sphere of influence of the company. This creates a narrower visualization and enables the extraction of specific anecdotes. For these more specific visions of the system, additional textual information is displayed. When a company is selected, a list of every politician that has held a position there appears. The user can browse through that list, and when hovering over a certain name the respective insect is highlight in the visualization. This makes the visualization more clear and tangible, directly associating the politicians name with their new form and explaining the display of the adopted metaphor. Once a company is clicked on, it stays visually marked in order to give notice to the user that such company was already explored.

Figure 6.14 The opening sequence of the visualization, as the politicians run from the top in order to encircle their first companies.

The viewer can also select a politician, either by clicking on the names list or by clicking directly on the insect. When a politician is selected, only the companies that it frequented are shown, enabling the isolation of its jumping behavior. The politician view displays two types of textual information. The first is a list of every corporate position of the selected politician, displaying the position's description as well as the timeframes. The second is a non-interactive list of every position held in government and other political activities such as party affiliations. This list too has the position's descriptions and the corresponding timeframes (see figure 6.16).

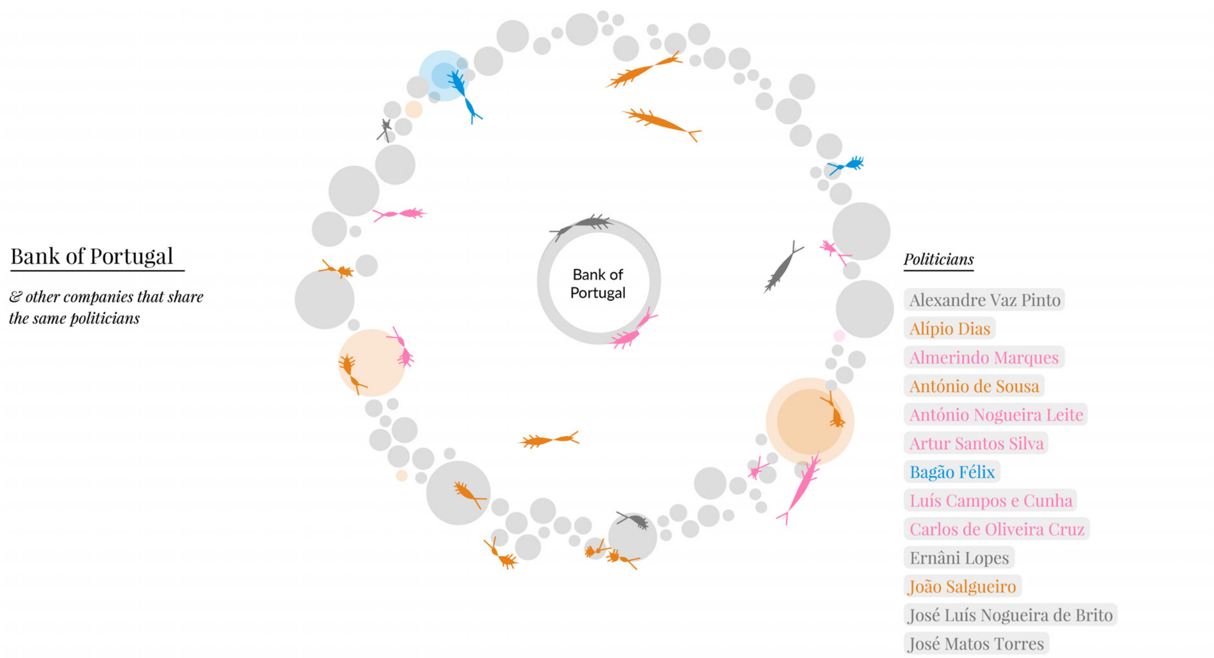

Figure 6.15 The sphere of influence of The Bank of Portugal, displaying companies that share the same politicians.

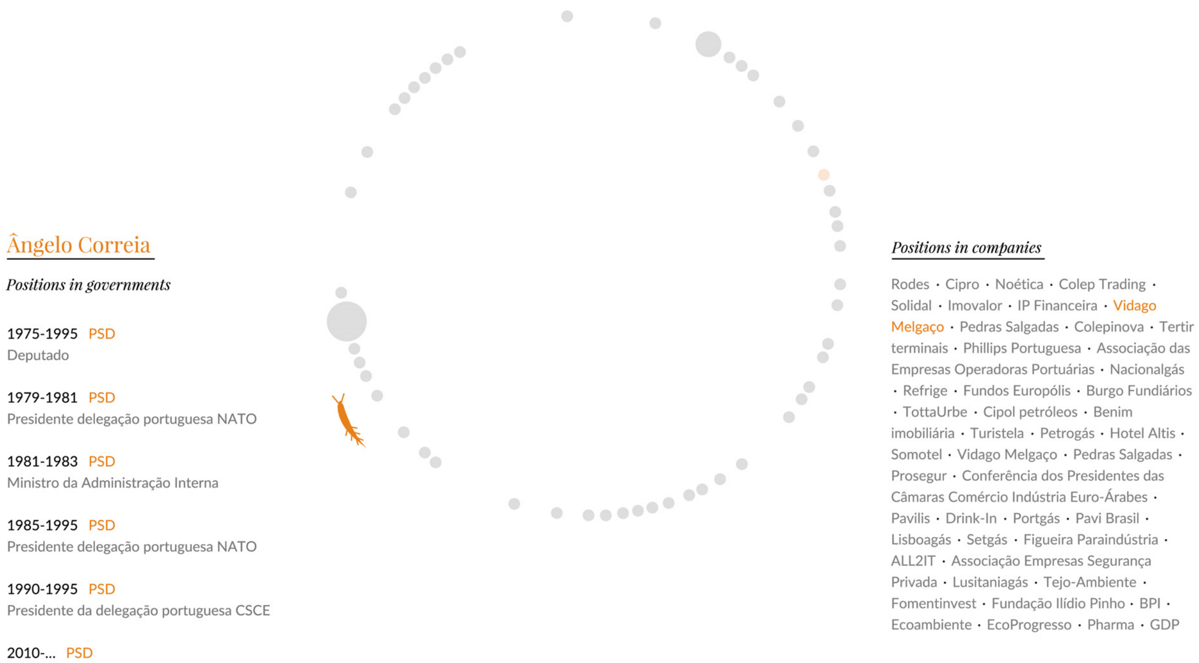

Figure 6.16 The sphere of influence of Ângelo Correia, the politician with the most positions in companies. A list of his corporate positions is displayed on the right that shows the respective position and timeframe when hovered over. On the left, a list of every government position that he had along with his party is shown.

Clicking on a politician, if not through the names list, can be challenging. The politicians have a frenetic jumping behavior that makes them hard to catch. Catching and chasing politicians is a ludic and metaphorical intent of the visualization, one that takes a satirical tone by inciting the user to interfere with the politician's busy life. In order to attain this, a mechanism was implemented that attracts the closest politician to the mouse pointer. When an insect is in the sphere of influence of the mouse pointer, it becomes stuck and encircles it repeatedly. The user can now move the mouse pointer slowly in order to keep attracting the insect and observe how it follows the mouse through sudden inversions in movement and abrupt spirals (see figure 6.17). This is to say that in fact the user can play with the insect or move the mouse away and the insect will resume its busy activity, jumping to the company that it was pursuing before being halted. The visualization presents facts in the dataset. Given the expected role of the user to absorb information and extract conclusions, it is of satire and irony to hint to the user the power of intervention in the visualization.

Additionally, other classic forms of interaction were implemented. The user can browse to the next or the previous state in the visualization. For example, if the users have a company selected and then click on a certain politician, after browsing through the politician's information they can go back to the company they were exploring. The users can also type in the names of companies and politicians in a search field and jump directly to them.

Figure 6.17 Three politicians caught by the mouse pointer.

6.3.7 Technology

As a way to reach a large audience, the visualization is web based. The ideas of openness and transparency that relate with the exposition of corporate-political relationships culminated in the use of standard technologies, such as Javascript, HTML, and CSS. More specifically the visualization was limited to run with the latest web technologies, such as ECMA5 compliant Javacript, CSS 3 and the CANVAS2D element of HTML5. Other standard recent technologies such as WEBGL could also be used to implement this visualization and enhance performance, but they are not as widely used as the CANVAS2D that, for example, can more easily be run on a mobile device.[59] In fact, today modern browsers are using the GPU to accelerate the rendering of the CANVAS2D elements. The library sketch.js[60] was used for the usual routines of drawing and animating on a CANVAS2D element, and the Coffee Physics library[61] from the same author was used to implement the companies' collisions.

When using physics simulation in a browser, performance is paramount due to the large number of bodies in the system (354 companies and 130 politicians). For numerical integration the Improved Euler method provided by the Coffee Physics Library was used. The Improved Euler is a second-order Runge-Kutta method for integration that is not as computationally intensive as a fourth-order Runge-Kutta and is more accurate than the Euler method (Süli and Mayers 2003, pp.325-329).

6.4 Dissemination and engagement

The *Ecosystem of Corporate Politicians* has the objective of reaching a wide and vast audience while raising awareness of the relations among Portuguese companies and politicians. In this concrete case of an interactive visualization specifically designed for the web, the dissemination of the visualization can be measured and characterized through web analytics. Raising awareness passes through both a broad dissemination and an effective understanding of the visualization. Although this section makes no direct claims on how well the users understood the visualization, it is suggested there was a strong and positive reaction. Engagement with visualizations can be estimated through several metrics, such as the number of user interactions with the visualization and the time spent on it (Mahyar et al. 2015).

6.4.1 Dissemination and user visits

The Ecosystem was made available on http://pmcruz.com/eco on December 23, 2013. The traffic was monitored using Google Analytics (GA) and until September 25, 2015, during the course of roughly 21 months, 145,822 users and 175,514 sessions have been reported. About 83% of the users were classified as new visitors and 17% as returning. The majority of traffic (88%) originated from Portugal, as it was its main target audience. Most of the traffic was generated in the first three months of launching (see graph 6.2). Between December 23, 2013 and March 31, 2014 it reported 112,196 users (77% of the total users) of which 84% were new visitors. For this period, the average user visits per day was 1,275.

Graph 6.2 Number of users' visits for the visualization (December 22, 2013–March 31, 2013). Most of the visits originate during this period, quickly rising from the first day, with significant lows during the holidays (24 and 31 December), then reaching a peak during the first week of January 2014 and having the number of users declining throughout January.

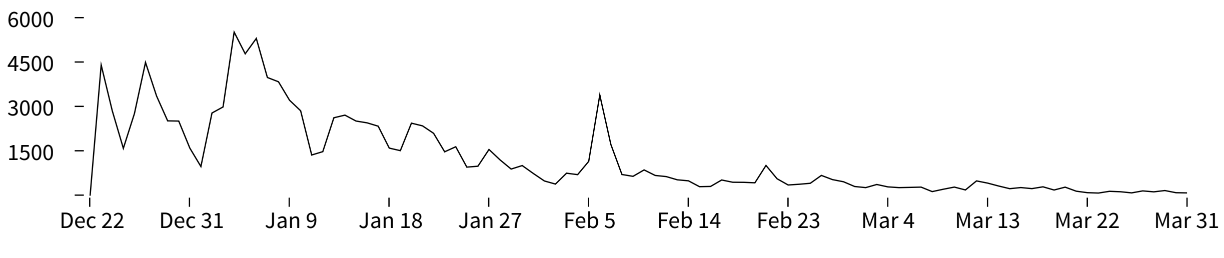

The visualization gained traction on Facebook, where it was first published. Several main traffic provenience channels are reported as direct, social, or referral, accounting for a traffic total of 98%. The distribution through traffic channels can be seen in table 6.3. Considering social traffic, 92% of it origi-

59 http://caniuse.com/#search=canvas
60 http://soulwire.github.io/sketch.js/
61 https://github.com/soulwire/Coffee-Physics

nated from Facebook. It was discovered through the Facebook Graph API that the visualization link was individually shared 47,005 times on this network. As for the referral traffic channel, it is important to state that even if 2% of the total traffic was reported as a referral from personal blogs, it was found that this represented at least 62 individual blog posts on the *Ecosystem of Corporate Politicians*.

Table 6.3 Shares of the users provenience through different channels (December 22, 2013–March 31, 2014).

Channel	*Percentage*
Direct	58%
Social	32%
Referral	8%
Others	2%

Direct traffic as reported by GA refers to when a visitor typed the website address directly into the browser, or when GA could not track the source referral, when, for example, a user clicked on a link in an email, in certain mobile applications, or in a PDF (Cutroni 2010, p.93). Graph 6.3 shows the partitioning of acquisition of new sessions through the channels as reported by GA – direct traffic (black), social (blue), and referrals (red). One important aspect of graph 6.3 is the verification of the origin of traffic from social channels, mainly from Facebook. It can be observed that direct traffic rises quickly during the first days until it surpasses the social channel as the main traffic channel. One might assume that the mediums of dissemination diversified, escaping from its inception on the social web, and spreading to other channels of communication such as email, which would justify the high amount of direct traffic. After this initial inversion, direct traffic generally remains as the main channel to acquire visitors.

Graph 6.3 shows two other interesting peaks that were caused by dissemination events, external to the social web sphere, that changed its organic evolution. The first peak refers to January 4, 2014, when *jornal i*[62] published a piece about the visualization (Reis 2014). This peak can be first observed through the referral traffic channel on January 4. This tendency seems to have had an almost immediate impact on the spike of traffic from social networks as well as a later impact on direct traffic on January 6. Nonetheless, it should be noticed that both the traffic from social and direct channels was already on the rise (if holidays are excluded), before the peak from referral traffic caused by the publication of the piece in *jornal i*. Therefore one cannot conclude how much of the traffic increase was caused by *jornal i*, and how much there would be if left to evolve organically. Although one could hypothesize that the double peak observed in aggregated traffic in graph 6.2 during 4 and 6 January was caused by a ripple effect from *jornal i* publication adding to the already-growing organic traffic.

The second peak refers to February 6, 2014 when *Sábado*[63] magazine published as well a piece on the *Ecosystem of Corporate Politicians* (Pereirinha 2014, p.82) in print and online. On the peak of 6 February all traffic channels have similar shares. Although *Sábado*'s piece seems to have increased the traffic, it did not appear to have generated traction capable of reducing the decaying tendency of the traffic. The number of printed copies of *Sábado* and the number of new users originated through *Sábado*'s online referrals were superior to the ones originated b *jornal i*'s. *Sábado* had a print run of one-hundred thousand and 1,926 new users came from its online referrals, while *jornal i* had a print run of 27,259 and provided 1,865 new users from its online referrals. In spite of this, *Sábado*'s dissemination was not able to sustain the traffic or generate a rippling effect in the social and direct traffic channels. This may be due to several causes: the novelty factor of the visualization had worn out; possibly similar audiences were targeted by both the publications; *jornal i* had a much more extensive piece, highlighting it in its front cover and dedicated two double-page spreads to describing the visualization itself; on the other hand, *Sábado* magazine dedicated a three column page, without mention on its cover and without any prominence in the index. It should be noticed as well that *jornal i* 's piece had a much more journalistic approach when reporting facts and describing the visualization in a direct way, while *Sábado* centered their piece on one of the authors of the visualization and his work in general.[64]

Graph 6.3 Number of sessions per traffic channel (December 22, 2013–March 31, 2014). — Social channel, — direct channel, — referrals channel. Two peaks in traffic can be observed that are relatable with the publications of *jornal i* and *Sábado* magazine, on January 4 and February 6 respectively.

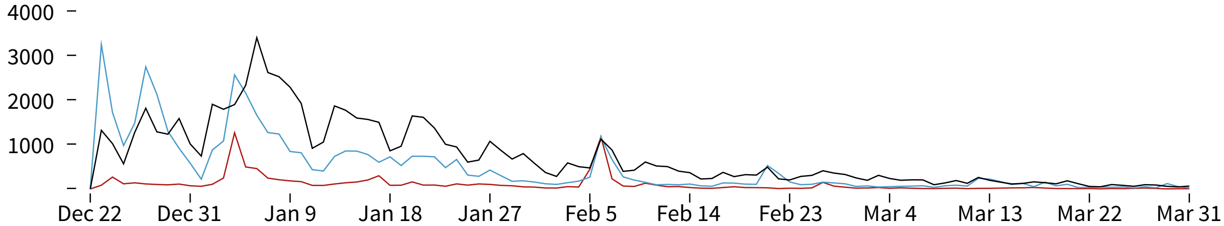

After the piece in *Sábado*'s magazine, the traffic kept declining, with little change between April 1 and October 31 of 2014. During this period, it had an average number of 64 visits per day of which 77% were new and 23% were returning users. The Latin-American Biennial of Design attributed a special award[65] on "Design of Social Movements" for the *Ecosystem of Corporate Politicians*. On November 3, 2014, the University of Coimbra disseminated the attribution of this award on its own channels[66], and was seen by a Portuguese news website[67] dedicated to technology. Such dissemination caused a peak both on direct and referral traffic channels in similar shares, and there was a smaller although distinguishable increase in traffic from social channels, which can be observed in 6.4. After this event the application received a reasonable amount of visits until the end of March, with an average of 117 users per day (79% new, 21% returning). As can be seen in graph 6.4, the period of June 1 to August 20, 2015 is the period with less traffic, having an average of 24

[62] *jornal i* is a major Portuguese newspaper with an online presence. The print run for January 4, 2014 was of 27,259.
[63] *Sábado* magazine is a major weekly Portuguese periodical in general themes. The print run on February 6, 2014 is estimated by the magazine of one hundred thousand copies. The magazine has an online presence.
[64] In contrast with *jornal i*'s tone, the title of the piece on *Sábado* magazine is loutish and vulgar, reading: "The biggest enemy of politicians is a nerd."

user visits per day (75% new, 25% returning). The traffic was steady at this low rate, with one exception that can be observed in graph 6.5. The slight increase in traffic closer to the beginning of October can be perhaps related to the Portuguese Legislative Elections in 2015 that were held on October 4. More interest on the part of the electorate in social-political themes might have caused the increase observed, which rapidly dissipated.

Graph 6.4 Number of sessions per traffic channel (April 1–December 24, 2014 and December 25, 2014–August 20, 2015. — Social channel, — direct channel, — referrals channel. A spike in traffic can be observed on November 3 as a consequence of the dissemination held by the University of the Coimbra on the visualization being a recipient of the "Design of Social Movements" award from the Latin-American Biennial of Design.

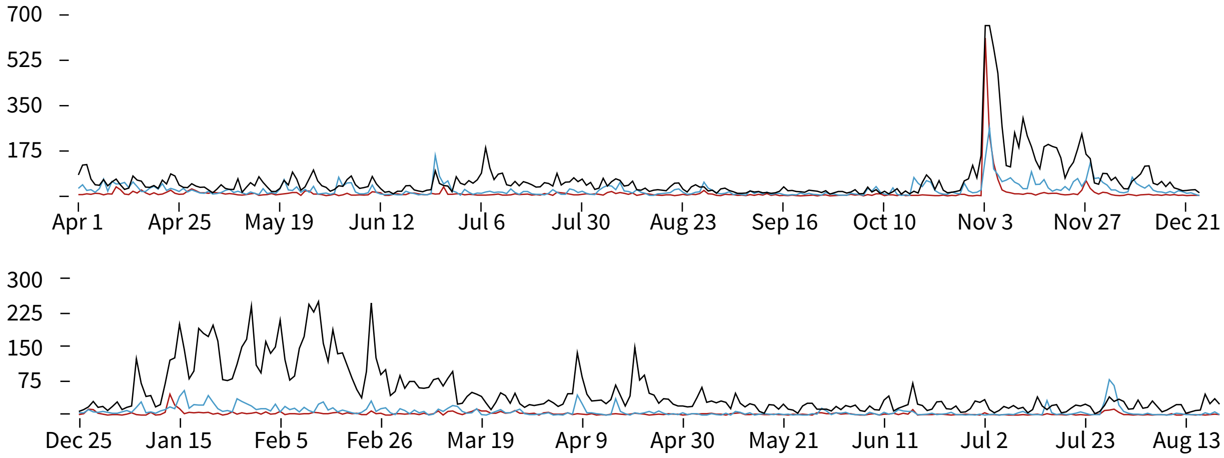

Graph 6.5 Number of sessions (August 21 – October 15, 2015). A slight increase in traffic with a peak on October 4 might be due to the Portuguese Legislative Elections of 2015 that were held precisely on that day.

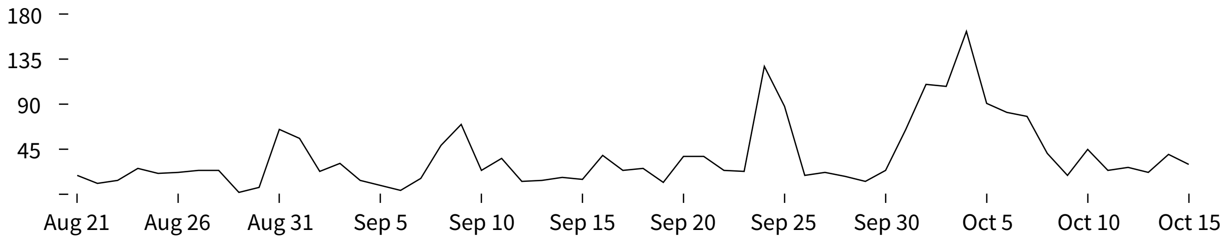

6.4.2 Engagement and time spent on site

A metric usually used to measure engagement with a website is the bounce rate. According to GA, the bounce rate is the percentage of single-page sessions, where the user left the site from the entrance page without interacting with the page[68]. Clicking on politicians and companies in the application is not an organic interaction for GA since it does not generate additional page

65 http://galerias.bid-dimad.org/bid_14/?p=3166

66 http://noticias.uc.pt/destaque/1/bienal-ibero-americana-de-design-distingue-projeto-da-uc-que-estabelece-a-teia-de-ligacoes-entre-politicos-e-empresas/

67 http://exameinformatica.sapo.pt/noticias/internet/2014-11-03-As-relacoes-entre-politicos-e-empresas-num-unico-site

views. This is a natural limitation in how GA tries to measure the presence of users on a website and how GA calculates time spent on a website. The time spend on a website can only be computed between two page views of the same user's session (Cutroni 2010, p.74). So, if a user entered and left through the same page, independent of the time the user spent there, GA can only register the entrance page view and computes time spent on a site as zero and classifies the visit as a bounce. This is the most common reported behavior for the *Ecosystem of Corporate Politicians* since it was implemented as one single web page. Exceptions exist when users follow a link in the informative panel of the visualization to another page in the same domain (pmcruz.com). In these rare cases the visits are not classified as a bounce and the time spent on the visualization can be measured. Having these limitations, the reported bounce rate of 83% since the publication of the visualization until August 20, 2015 is an inaccurate overestimate and would signify that most of the users that reach the application do not engage with it and leave right away.

Until August 20, 2015, 28,891 sessions were classified as non-bouncing (17% of the total of sessions) from 26,692 different users. These sessions had an actual average session duration of 00:05:08. When observing these sessions, 88% of them also left the website through the landing page without visiting more pages in the domain before leaving. One might ask how GA could have measured a positive time spent on the site for these sessions if there were no apparent additional page views. The answer lies in the fact that each of these sessions amounting to 88% had exactly two page views in the same landing-exit page. The two page views might be caused by a user clicking twice on the web visualization link, by re-entering the address in the browser or simply by refreshing the browser at any point in time. Graph 6.6 reports the time spent on the website from the users that entered and left through the same main page, resulting in an average time spent of 00:04:55. It is therefore a measure of the amount of time between hits, for example how long a user spent on the visualization before refreshing it. There is no way at this point to ensure that the users indeed interacted with the visualization or if they just left the page open. The system takes more than ten seconds to be available for exploration, assuming that the users quickly push the exploration button in the initial screen. Therefore the amount of users that spent this time between hits perhaps wanted to read the initial information again or observe the first moment of the visualization. About 71% of the sessions in graph 6.6 had more than two minutes between hits.

[68] https://support.google.com/analytics/answer/1009409?hl=en

Graph 6.6 Time spent on site in seconds by the 88% of the non-bouncing sessions that entered and left through the same main entrance of the visualization (December 23, 2013–August 20, 2015). The average session duration is 00:04:55.

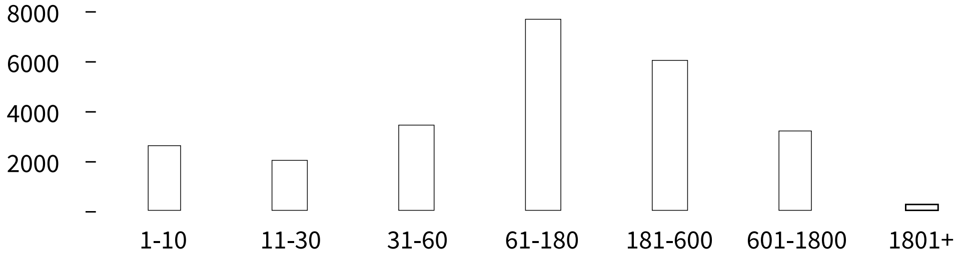

The remaining 12% of sessions that were classified as non-bounces, followed links in the informative panel to other pages in the same domain. Graph 6.7 shows the time spent on the visualization for these sessions, showing the amount of time that a session was opened on the main visualization before the user jumped to another page in the same domain. These sessions had a generous average session duration of 00:07:44, where 89% of the durations were superior to two minutes. These times are arguably sufficient for a high amount of interactions with the visualization, but one should notice that analyzing this sample of 12% of non-bounces might be biased, since it is analyzing a sample of users that already showed some degree of engagement with the visualization, or at least interest or curiosity, since they clicked on the informative panel, wanting to know more about the visualization and ultimately know more about the authors of the visualization.

Graph 6.7 Time spent on site in seconds by the 12% of the non-bouncing sessions that clicked on an additional page in the same domain (December 23, 2013–August 20, 2015). The average session duration is 00:07:44.

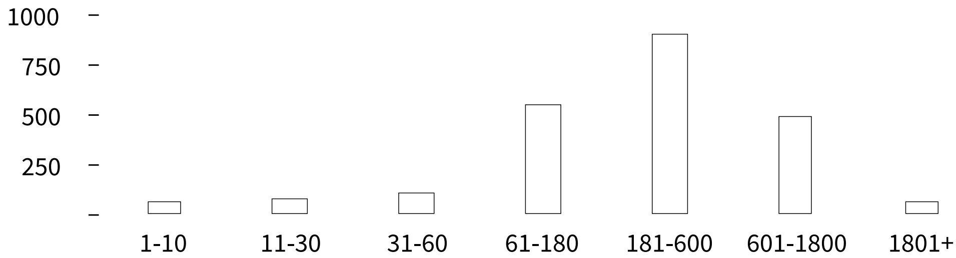

On August 21, 2015, a script[69] was installed in the visualization that fed GA with tracking events: an event was generated for each five seconds that a session was open on the visualization; an event was generated each time a user clicked on a company; and an event was generated each time a user clicked on a politician. These events enable GA to compute the time spent on the site and the bounce rates in a much more accurate way, as it does not rely solely on page views for tracking events. August 21 to October 15, 2015 shows a sample of 2,163 sessions from 1,861 unique users, where 83% were new and 17% returning. The bounce rate in these more accurate measuring conditions was of 43%, against the previous reported rate of 83%, which is a drop of 52%. For this shorter period, one can assert that 57% of sessions represented an concerted amount of time spent on the website and did not just left. The non-bouncing visits were 89% new visitors and 11% returning. Although the previously reported bounce rate of 83% until the monitoring of tracking events is an overestimate, it is hard to estimate the real percentage of users that did not bounce. The preceding month (July 20-August 20, 2015) the measuring of tracking events was reporting a bouncing rate of 91% by GA. This reinforces the inaccurate nature of GA's bouncing rate reports and thus suggests that the real bouncing rate should be viewed much closer to 43% than to 83%.

During this period from August 21 to October 15, 2015, the average session duration of non-bouncing sessions was 00:04:46. This is similar to the one reported by August 20, 2015 (graph 6.6 – or the non-bouncing sessions with the same landing and exit page). This puts forward a question related to how well the sample on graph 6.6 represents the real time spent on the visualization by users. The samples on graph 6.6 and 6.8 have different characterizations. Graph 6.6 shows sessions that had the same landing and exit page and had exactly two page views, implying a double hit behavior from the user. That is arguably the reason why the number of sessions that had a duration of 1-10 seconds and 11-30 seconds in graph 6.6 is comparatively more pronounced than in graph 6.8. The sample in graph 6.8 is much smaller, but is characterized by 96% of sessions that entered and left through the same page, against the 88% that sample 6.6 represents. The sample on graph 6.8 has a double hit behavior as well, but it accounts for only 22% of the sessions. This leads to the suggestion that the double hit behavior – a second refresh, a second enter in the browser address bar, or a second click on the same link to the application – is not a factor that can significantly bias the sample, and thus grap 6.6 can be seen as the result of a robust measure and distribution of the actual time on the visualization by the users.

Graph 6.8 Time spent on site in seconds by the non-bouncing sessions (August 21–October 15, 2015). During this period, the presence of sessions on the visualization was monitored as tracking events. The average session duration is 00:04:46.

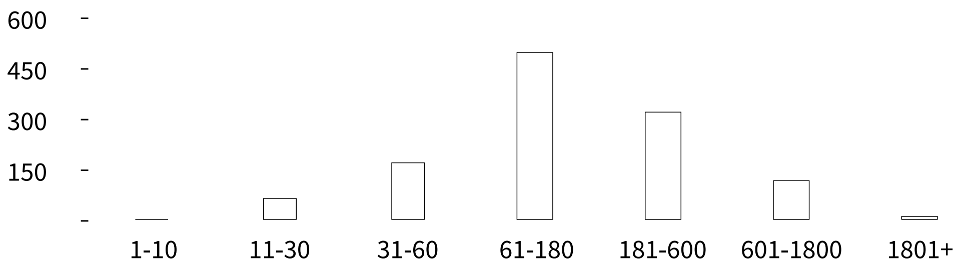

Another factor that is important to assess is if these non-bouncing sessions from August 21 to October 15, 2015, with an average duration of 00:04:46, were indeed accompanied by user interactions. It was discovered that for this sample, 92% of the sessions involved some type of interaction with a politician or a company, generating a total of 9,409 interaction events. These 1,123 sessions with at least one interaction event with either a company or a politician, had an average of 7.6 interaction events per session. From the total in-

[69] http://riveted.parsnip.io

teraction events, 71% were with companies and 29% with politicians. As can be seen in graph 6.9, 51% of the interaction events with companies refer to the top ten companies that had the most politicians (table 6.1), out of a total of 354 companies. This behavior is expected since the top companies also have larger areas and hence are emphasized by the visualization, being prone to receive more clicks.

Graph 6.9 The 20 companies with the most interaction events. The top ten companies with the most interaction events also coincide with the companies that have the most politicians and hence have larger sizes.

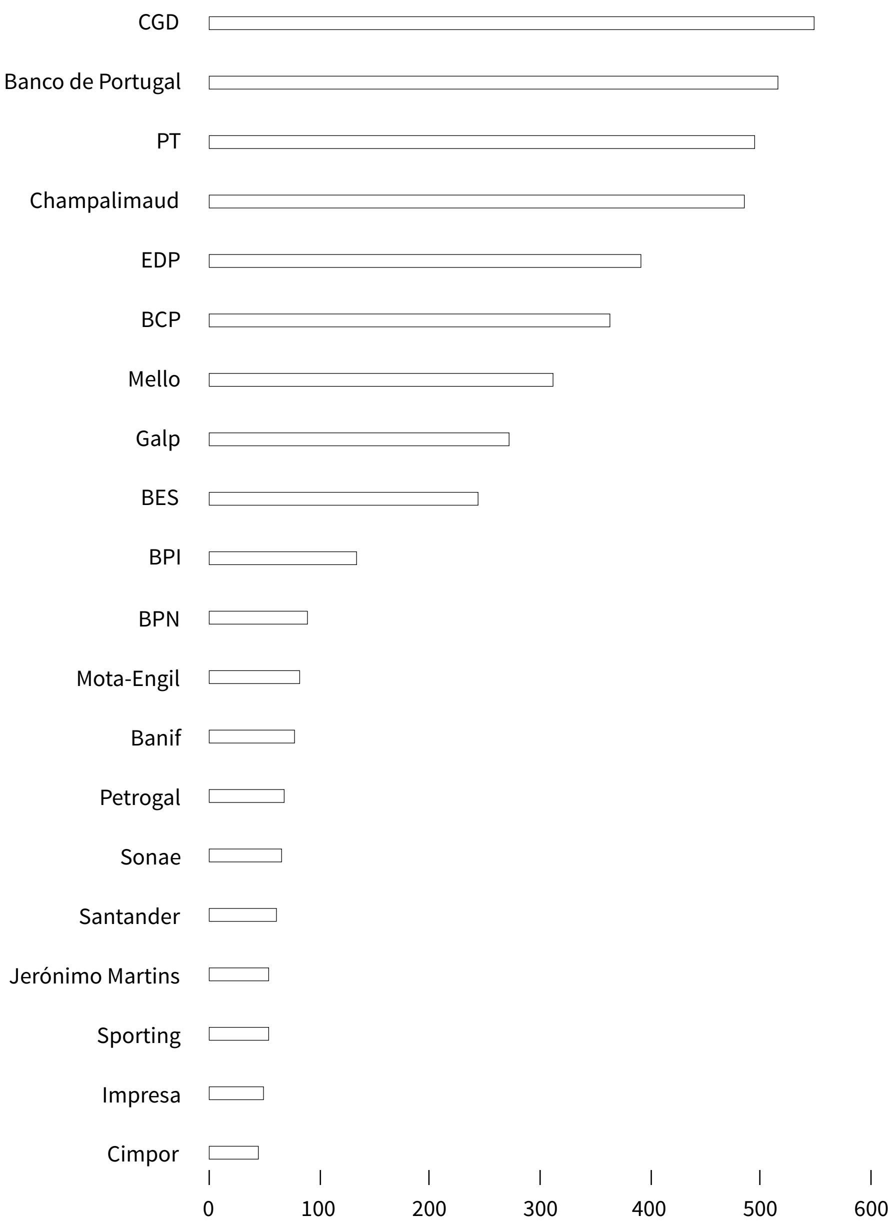

There is a more even distribution of interaction events with politicians than for the companies (see graph 6.10). The 20 politicians (15% of the total 130 politicians) that were more interacted with represented 34% of the total interaction events. In fact, only three of the 20 politicians with the most clicks, can also be found in the top ten list of politicians with the most positions in companies (table 6.2). This lead to the conclusion that users were not actively and visually searching for the politicians that held more positions. The users can select specific politicians by name, and therefore politicians can be selected by general popularity, given the users' personal degree of information regarding this theme, which can be related with the popularity of a certain politician. One of these cases seems to be the then Prime Minister of Portugal, Pedro Passos Coelho, who received the most interaction events despite not being one of the politicians with the most positions in companies (Pedro Passos Coelho had ten positions in eight different companies).

Graph 6.10 The 20 politicians with the most interaction events. The events are more evenly distributed compared to companies. Only three of these politicians also features in the top ten list of politicians with the most positions in companies.

6.5 Conclusions

This chapter described a visualization of the relations of Portuguese politicians and companies. The *Ecosystem of Corporate Politicians* is a concrete application of an approach on semantic figurative metaphors in visualization. The system is interactive and invites the users to a ludic exploration. It relies on the modification of a radial network, which, instead of having a strict radial organization of nodes, creates a circular belt of packing bubbles. Furthermore, it modifies a typical radial network model by not drawing edges that connect companies, but instead depict flying politicians that connect them through movement. The main metaphorical visual cue is the portrayal of politicians as generic insects or mollusks that frenetically jump among companies. Each insect is characterized by a color referring to the politician's political affiliation, and by a rounder or slimmer shape that hints the number of positions that a politician held. The politicians' representation and movement refer both to insects due to their frenetic jumping behavior, and to mollusks due to their flexible shapes. This metaphorical and authorial intent was chosen to have a satirical representation of this theme as a way to enhance memorability and engagement. One of the main objectives of the visualization was to reach a wide and broad audience, while observing that the audience indeed engaged with it.

From December 23, 2013 to September 25, 2015, the visualization received 175,514 sessions from 145,822 users. The majority of the traffic was generated in the first three months of the visualization's publication. Most traffic was generated organically through shares from Facebook (the total number of shares accounts for 47,005) and simultaneously propelled by a piece in a major Portuguese newspaper. The visualization was the central theme of 62 different posts on several personal blogs. Portugal was the geographical source of 88% of traffic. With these numbers, and in the context of the main target audience[70] for this visualization, it is defended that the visualization had a significant dissemination.

Other significant contribution to the traffic increase was the publication of another piece in a major Portuguese magazine, a weekly periodical on general themes. The traffic decreased continuously after these first months, with exceptions for periods marked by the dissemination of the award for "Design of Social Movements" from the Latin-American Biennial of Design and the 2015 Portuguese Legislative Elections of 2015. Nonetheless, these exceptions were not able to push the traffic to levels of several thousand daily sessions found in the first months. This is probably due to the novelty momentum that the visualization incurred immediately after its launch.

By monitoring a sample of the visits on the visualization, a bounce rate of 43% was observed, resulting in 57% of the sessions staying on the visualization for an average time of 00:04:46. This suggests that the users engaged with the visualization since it is enough time to interact with several companies and politicians. For this sample, an average of 7.6 interaction events per session was observed, of which 71% were with companies and 29% with politicians.

This apparent lack of interaction with the politicians results from the system behavior and the application of figurative metaphors in this visualization – the politicians' frenetic jumping behavior and dimensions are less perceptible, and interaction itself is harder, in spite of the mouse attracting the closest politician to spiral around the cursor. Ultimately it is a design decision that forces the user to explore the system. Furthermore, the companies are more responsive to user interaction, because their size dramatically increases when the mouse passes over them and are therefore easier to click on . In this sense, it is argued that the interaction with the packed bubbles model is a seemingly more familiar way to explore the system. Nonetheless, this is not in detriment to the chosen design of the politician's portrayal and behavior, which are central to the visualization. The most distinguishable semantic and metaphorical cues are implemented on top of the politicians. These are subtleties that complement the information and add an authorial tone to the visualization. The mere suggestion that fat politicians frequent more companies and that the user can intervene in their eternal routine by disrupting their movement exists to provoke the audience and possibly raise awareness of this theme.

Of course, there are more clear ways to depict a networked structure of these relations among companies and politicians, as suggested in section 6.3.2. Nevertheless, the design decisions that were taken for this visualization, more than being tied to metaphorical and satirical intents, are centered on a notion of obscurantism with this type of information. This type of information is available, but rarely aggregated, it is free and public, but politically inconvenient. It is risky to suggest that such information is deliberately hidden, but it is safe to state that it is often neglected and passively masked, since its sources are not unified, do not store information persistently, and are often cumbersome to access. The mainstream media themselves also lack a widely available systematic treatment and dissemination of such information. This is to say that having a slight degree of difficulty in exploring this particular visualization can be utilized with a ludic result while enhancing engagement for the users. Inviting the user to explore the visualization is forcing the user to discover the relations for themselves, while giving the satisfaction of untangling the network of political-corporate relations.

The website received 107 voluntary and spontaneous comments in its dedicated comments section[71]. The comments are positive, with 95% of them addressing the visualization directly in an enthusiastic and supportive way. The rest of the comments are from internal discussions or simply judgments on politicians with corporate careers. Although this is a small set of comments, received without any formal requirements or supervision, its spontaneous nature should be seen as a valuable way to assess if the designed metaphorical cues were assimilated by the audience. Among the 107 comments, 11 address the main metaphorical cue of the visualization, i.e., politicians as generic insects or mollusks. As mentioned previously, this figurative approach still employs an abstract language and therefore is open to personal interpretations.

[70] Portugal had a total of 10.4 million inhabitants in 2014, of whom 66% have access to the internet. Source: INE, PORDATA,

From this subset of comments, politicians were generalized as parasites, ants, cockroaches, worms, and crabs. These comments are a strong suggestion that the audience indeed assimilated the main metaphors implemented in the visualization. The following list shows 11 comments from different users that classify the politicians as an insect or mollusk, some sort of infestation, or simply as an organic version of the visualization.

— "(…) here is the degree of equality among the Portuguese politicians and parasites (…)"
— "A simple homemade formicide, and we're done!!!"
— "(…) Get rid of the cockroaches for good!!"
— "(…) The symbiosis between politicians and corporations!"
— "(…) Perfect analogy between cockroaches and other parasitic insects! (…)"
— "(…) It is confirmation that Portugal needs a disinfestation…"
— "(…) I would like to suggest less repulsive worms (…), it repulses me to interact with these insects (…)"
— "(…) Between politicians and cockroaches, I prefer the latter (…)"
— "(…) The diagnosis is done. Treatment: disinfestation!!!"
— "So many parasites!"
— "I would suggest crab lice."

It is argued that having such a metaphorical conceptualization of the graphical elements in the visualization contributed to its impact and memorability. Another four users were complimentary about the visualization, praising it as public service.

After the visualization's publication, an extended study and database of some of the authors of "Donos de Portugal" was made available (Louçã et al. 2014). This new dataset now gathers information from 769 politicians and their ties with corporate entities. It is a much vaster dataset, and the model designed for the *Ecosystem of Corporate Politicians* would not be suitable to represent such an amount of information, both from the perspective of organization of visual space, as well as due to computational limitations in the real-time simulation of several hundred politicians. This is to say that the visualization, naturally, has limitations. However, it can be adapted by, for example, depicting corporate relations for each constitutional government individually. But what would that do to the visualization? Transform it into a model where politicians are constantly represented by generic insects? This is to say that the *Ecosystem of Corporate Politicians* was not designed for robust utilitarian applications, but rather as a satirical and exploratory piece, truthful to information, in the context of casual visualization. Repetition weakens effective satire, and with the metaphors assimilated, novelty is lost together with the objective of reaching wide audiences with refreshing approaches. Audiences can be connected with data by employing more figurative representations, specifically designed for a dataset and a target audience.

[71] The comments can be viewed in the visualization webpage http://pmcruz.com/eco. Every comment except one, all are from Portuguese users and written in Portuguese.

7

Discussion

Semantic figurative metaphors can be used in the context of casual visualization, where communication with large audiences is favored in detriment of more utilitarian usages, aiming for contemplative and momentary uses, trying to provoke awareness and reflective insights. They humanize visualizations because they are less abstract and more figurative, connecting strongly with the audience in an emotional way. They are not models, but an approach to visualization. Each non-data metaphorical aspect should be tailored specifically for a dataset, for a structural model and for a message, and thus it is rarely generalizable. They should have well-defined communicative authorial purposes and not a mere add-on: implementing semantic figurative metaphors should be part of the design process since its inception.

Semantic figurative metaphors carry the weight of authorial intent in visualization, which is subjective in itself. Additionally, they explicitly add non-data elements to visualizations in order to augment their metaphorical value. Adding non-data aspects to visualization comes at the expense of data inaccuracies and may cause confusion for the viewers when distinguishing what is data related from what is not.

Within this context, the main objective of this dissertation was to to show that these visual metaphors can be explored together with the distortion of visualization models, to the point of deliberately inserting inaccuracies in data representation, in order to approximate the audience to the visualization while being faithful to the data.

This dissertation presented three works that explore semantic figurative metaphors. These works use computational systems that adapt to data and deliberately insert inaccuracies and distortions in its representation. These evasions of familiar visualization models are used to portray authorial messages while still being clear and faithful to data. The three works placed an emphasis on the modifications of the structural models as well as on the addition of non-data related visual cues, in order to implement specific semantic figurative metaphors.

The visualization of the Empires' decline uses a conceptual framework for generative storytelling, translating countries into actors, and growths and disintegrations into time events that model the system's interactions. The simulation presents different results for every run, telling the same story with slightly different expressions. The first implementation of this visualization modifies the structural metaphor of packed bubbles and added visual cues that allude to the metaphor of disintegration and competition. The second instantiation, being a refinement of the first, uses Dorling cartograms as a starting point in order to include geographic information. This second implementation shows which parts of the world were once part of one of the four empires, bringing the narrative to contemporaneity. As in packed bubbles, Dorling cartograms were implemented in an animated way that includes behaviors that metaphorically allude to thrusts of independence. The visual cues added on these models result from the implementation of empires as soft bodies. This portrayal is a figurative way of depicting competition by deforming the empires as they collide. The distortions yield on the empire's shapes also convey a sense of fragility and ephemerality. Distorting the shapes of the empires incurs errors in information representation. These errors were measured, and by comparing them with the errors that humans make in area estimation, it is argued that they are negligible.

The Empire's visualization was widely adopted, and viewers engaged with it through comments and likes. Most of the comments are eulogistic with adjectives such as "beautiful", "original", "hypnotic", "attractive", "dramatic", "organic, "visually stunning", "historically compelling", "efficient", "instructional", or "mesmerizing". This suggests that both the narrative options and the visual metaphors built into the visualization played an important role in its adoption. Additionally, viewers debated how the independence dates were chosen and suggested that more empires could be portrayed. These suggestions question how data was collected and how the story was extracted from data. This shows that indeed the viewers observed the visualization in detail, which is a strong suggestion of true engagement. On a conceptual level, the metaphors chosen were not put to questioning. Of the 477 total comments received, only six explicitly expressed that they did not understand what the areas represented or if the interactions or "wobbles" had any data significance. Using complex soft bodies to show transcontinental land areas is unorthodox, but its value to confer a dramatic effect lies in this unfamiliarity. The non-data movements that can cause such oscillations are used to strengthen the metaphoric value of the visualization and confer a dramatic tone, which is suggested here as being the mainly responsible for the visualization's wide adoption.

A second exploration regards the metaphorical mappings of traffic visualization. In these mappings, data was represented with its geographic properties. Representing vehicles and velocities on a map, through a common cartographic perspective, makes the structural metaphor behind these visualizations. This model was made less familiar because it distorts spatial dimensions in order to convey traffic information. For this, cartograms were used because they contrast familiar maps in unfamiliar representations of space, hence being more dramatic and having the ability to "shock" the audience. In order to implement this for this specific case, a new edge-based cartogram technique was devised, instead of the more common area-value approaches. This edge-based technique is formalized in a model that can be applied to other cities in order to show traffic information. The edges that represent a road network compress when velocities are high as if distances were perceptually farther, and distends when velocities are low and traffic is congested. The devised model for edge-based cartograms is implemented through a mesh of simulated springs. The cartogram is contiguous as it preserves topology, and it is dynamic because it can interpolate and continuously adapt to new data states, resulting in cartograms that are animated in simulation-time. This contrasts with most cartogram applications that only aim to provide a static version of a map, and need more computing power than the edge-based approach (Keim et al. 2004). The cartogram portrays the rhythm of a city with more impact on the audience, making the city grow, shrink, and deform as time passes by. Inherently, this adaption of the undistorted map employs a time-distance metaphor to show data.

As with any contiguous cartogram, the edge-based cartogram shows several errors when representing information of the data values and the original map. The distortions of the original map should not be exaggerated to a point that impairs the map's recognizability. The cartogram model yields different parameterizations that result in higher or lower errors. The model was tested for two cities and the corresponding traffic datasets. A total of 1728 parameterizations were simulated for each city in order to extract the best configurations according to a quality measure that accounts for the errors.

Two additional visualizations were applied on top of the cartogram model. The first is a visualization of trajectories that shows vehicle trails from the last 30 minutes and uses a metaphorical color encoding related to traffic signals: low velocities for reds and high speeds for greens and cyans. The second visualization model is more figurative and embodies the addition of visual cues contemplated in semantic figurative metaphors. This visualization shows the city as a living organism, more specifically, as a system of blood vessels: vessels pulse according to traffic speed, they get thicker according to traffic volume, have a stream of cells with speed and density directly related with the speed and volume of traffic, and are darker or lighter depending on traffic velocity, as if darker reds meant stagnant blood.

The blood vessels visualization implemented on top of the cartogram culminated in a less elegant artifact from a minimalist point of view. In this application, the authorial intent superimposes the functionality of data portrayal because the motion of the vessels and cells, together with the expansions and contractions, add non-data visual aspects that are perhaps too distracting. With this, the limits of distortion and addition of visual cues in semantic figurative metaphors have perhaps been reached. Although permissible in an authorial context, there is a point when the complexity of the artifact turns into visual confusion.

The "Ecosystem of Corporate Politicians" is a specific application of a semantic figurative metaphor. The system is interactive and invites users to a ludic exploration. It modifies the structural model of a radial network, which instead of having a strict radial organization of nodes, creates a circular belt of packing bubbles. Moreover, it also modifies the radial network model by removing edges that connect companies. In place of these edges, it adds flying politicians that connect companies through movement. Politicians are represented as generic insects or mollusks that frenetically jump among companies. This is a specific visual cue added to the model, which has a satirical intent. The visualization suggests in a subtle way that fat politicians frequent more companies and that the user has the power to intervene by disrupting their routine behaviors.

The visualization reached a wide audience of more than one hundred thousand users. The bulk of the traffic was generated in the first three months of the visualization's publication. Most traffic was generated through spontaneous shares on Facebook and simultaneously propelled by a piece on a major Portuguese newspaper. By observing a sample of the visits to the visualization, 57% of the hits were non-bounces, spending an average time of 00:04:46 interacting with the visualization. Each user in this sample interacted 7.6 times on average with the visualization, suggesting that users were engaged. Furthermore, the application received 107 voluntary and spontaneous comments. The comments are positive, with 95% of them addressing the visualization directly in an enthusiastic and supportive way. Among the 107 comments, 11 address the main metaphorical cue of the visualization, mentioning them as insects, mollusks, parasites, ants, cockroaches, worms, and crabs. This proves that the added visual cues, although more figurative, are abstract enough to enable users to devise specific interpretations. These comments strongly suggest that the users assimilated the main metaphors of the visualization.

Despite the Politicians visualization's satirical tone, its wide adoption is perhaps due to a notion of obscurantism in this type of theme. The information portrayed in this visualization is politically inconvenient, and is often neglected and passively masked. The everyday user does not have easy access to this type of information. By inserting degrees of difficulty, such as having to catch or chase a politician, the users are in fact unveiling information, giving them the satisfaction that they are exposing politically sensitive information that is not readily available.

All the three works culminated in visualizations that implement semantic figurative metaphors. Doing so encompassed two steps: the modification of the structural metaphor and the addition of visual cues. Table 7.1 synthesizes, across the three works, the structural metaphors used, how they were modified, and which specific visual cues were added in order to implement semantic figurative metaphors.

Table 7.1 Synthesis of the structural modifications and visual cues added to the visualizations.

	Structural metaphors	*Structural modifications*	*Figurative visual cues added*
Empires	Packed bubbles	New nations do not pack themselves and are propelled through thrusts	Circles as soft bodies
	Dorling cartogram	Nations have to travel to their locations and can be disrupted by others	
Traffic	Trajectories visualization	Space distortion using edge-based cartograms	Blood vessels with pulsing motion and streams of cells
Politicians	Radial network	Packed nodes and removal of edges	Politicians depicted as insects

All three works distorted visualization models that introduced inaccuracies in data portrayal (e.g. the area oscillations of the empires; the intrinsic value error of a cartogram; or the erasure of edges that connect companies). These inaccuracies were controlled in order to enhance the communication of certain authorial messages from the data. The empire's and the politician's visualizations both show engagement from the audience, having brought the audience close to the data. Moreover, the reception among users and viewers indicates that such engagement was due to the implementation of figurative metaphors. Each of the three described works could have used even more intense and elaborated expressions with the excuse of an authorial intent, but they were in the most part restrained to moderated visual cues.

These applications intentionally distorted shapes, subverted the use of color, added elements not present in data and manipulated time and space to present stronger messages. It is agreed that areas in the Empires visualization are not always accurate; that using cartograms necessarily brings an error to information portrayal; that using texture together with thickness to show traffic volume in the vessels visualization is redundant; that displaying velocity through motion is not particularly orthodox; and that by not drawing edges on a network the visualization of patterns might be lost. All this is a consequence of implementing certain semantic figurative metaphors, which at the expense of these inaccuracies try to intensify the communication with the audience. It is an exploratory side of visualization, and should be handled with care, but never at the expense of telling an untruth. Employing semantic figurative metaphors should be handled soberly and with restraint.

Dissemination

This dissertation describes semantic figurative metaphors as an approach to visualization, framing three works in this context, and contributes a rationale to the design options (Cruz 2015).

The Empires' work resulted in a system that is able to portray the decline of four overseas empires with varying degrees of expressiveness. The first version of the system (Cruz and Machado 2010b) is referenced (Jorge et al. 2012;

Jorge and Chambel 2014) and refined into a more detailed second version that displays more information (Cruz and Machado 2010a). The second version won the ACM SIGGRAPH Student Research Competition (2010) and was selected by a jury on the ACM SIGGRAPH Computer Animation Festival[72] and Electronics Alive VI[73]. This work implements a conceptual framework for generative storytelling (Cruz and Machado 2011a) and was mentioned as an example of storytelling in visualization in the *The Data Journalism Handbook* (Gray et al. 2012) and *Informotion* (Finke and Manger 2012, pp.160-161). The second version of the visualization is being used as teaching material in a Brazilian school manual (Santos 2013).

The city traffic work contributed with a model for edge-based cartograms (Cruz et al. 2015). Two visualizations were developed to be applied over the cartogram: the visualization of trajectories and the blood vessels visualization. The visualization of trajectories of Lisbon was part of several exhibitions such as MoMA's *Talk to Me*[74] (Antonelli and Carmody 2011, p.122), *Mediateca Expandida: Habitar*[75] (Vicente and Girardin 2010, pp.86-87), *The Art of Networks*[76], *MOVE: The Transportation Expo*[77], and *Nightime-Dreamreal*[78]. It was also mentioned in *Visual Complexity* (Lima 2011, p.147) and *Design for Information* (Meirelles 2013, pp.158-159). The visualization of trajectories applied to the city of London was featured in Wired UK (Hussey 2010).

A first approach to the edge-based cartogram was featured at the festival *Enter 5*[79] (Sedlák 2011, p.68). A first version of the blood vessels visualization over the cartogram (Cruz and Machado 2011b) was a semi-finalist in the ACM SIGGRAPH Student Research Competition (2011) and received an honorable mention in the *MiniMax Mapping Context* (2011) of École Polythechnique Fédérale de Lausanne. This initial approach was also described and framed in the context of a figurative approach to visualization (Cruz and Machado 2014) and was used as an example to introduce the concept of analogy in a Singaporean secondary education school manual (Doyle et al. 2011), as well as being referenced in Gondin et al. (2014) and Jorge et al. (2014). The latest version of the blood vessels visualization applied over the cartogram was also framed in the context of semantic figurative metaphors (Cruz and Machado 2016).

The *Ecosystem of corporate politicians* (Cruz et al. 2014) was made available together with its database under a Creative Commons License[80]. The database was used for social network analysis (Geadas et al. 2015), and the visualization is going to be applied to Danish political-corporate relations[81]. The visualization received the award for *Design and Social Movements* in the Ibero-American Biennial of Design[82] ("4ª bienal iberoamericana de diseño - Imaginación Colectiva" 2014, p.50) and was featured in the exhibition *The Art of Networks II*[83].

---

72 *Electronic Theatre - ACM SIGGRAPH Computer Animation Festival*, 25-29 July 2010, Los Angeles.

73 *Electronics Alive VI*, 21 January - 24 February 2011, University of Tampa, Tampa FL.

74 *Talk to Me: Design and the Communication between People and Objects*, 24 July - 7 November 2011, MoMA Museum of Modern Art, New York.

75 *Mediateca Expandida: Habitar*, 27 May - 8 November 2011, LABoral Art and Industrial Creation Centre, Gíjon.

76 *The Art of Networks*, 8 March 2012, Foosaner Art Museum, Melbourne FL.

77 *MOVE: The Transportation Expo*, Institute Without Boundaries and Evergreen Canada, Evergreen Brick Works, 24 May - 8 Oct 2012, Toronto.

[78] *Nightime-Dreamreal*, Power Station of Art, 3 December - 30 March 2013, Shanghai.
[79] *Enter 5: Datapolis, international art sci tech biennale*, 14-17 April 2011, National Technical Library, Prague.
[80] https://github.com/pmcruz/eco
[81] http://www.magtelite.dk
[82] *4th Ibero-American Biennial of Design*, 25-29 November 2014, Madrid.
[83] *The Art of Networks II, 26 March - 31 May 2015*, New York Hall of Science, New York.

References

4ª bienal iberoamericana de diseño – Imaginación Colectiva, 2014. Madrid: DIMAD.

Andrienko, N. and Andrienko, G., 2013. Visual analytics of movement: an overview of methods, tools and procedures. *Information Visualization*, 12(1), pp. 3-24.

Antonelli, P. and Carmody, K., 2011. *Talk to Me – Design and the Communication between People and Objects.* P. Antonelli and K. Carmody, eds., New York: Museum of Modern Art.

Bailey, W.B., 1911. Apportionment Map of the United States. *The Independent*, 70(3253).

Bal, M., 1997. *Narratology Introduction to the Theory of Narrative.* 2nd ed., Toronto: University of Toronto Press.

Barber, C.B., Dobkin, D.P. and Huhdanpaa, H., 1996. The quickhull algorithm for convex hulls. *Transactions on Mathematical Software*, 22(4), pp. 469-483.

Bateman, S. et al., 2010. Useful junk? The effects of visual embellishment on comprehension and memorability of charts. In ACM Conference on Human Factors in Computing Systems. New York: ACM, pp. 2573-2582.

Bertin, J., 1967. *Sémiologie graphique: Les diagrammes, Les réseaux, Les cartes.* Paris: La Haye Mouton, Gauthier-Villars.

Bertin, J., 2010. *Semiology of Graphics: Diagrams, Networks, Maps*. Madison, Wisconsin: ESRI.

Bianchi, M.T. and Viana, R.C., 2012. Estudo exploratório sobre as ligações políticas das empresas cotadas em Portugal. In XV Encontro AECA. Ofir-Esposende.

Black, M., 1962. *Models and Metaphors – Studies in Language and Philosophy*. Ithaca, New York: Cornell University Press.

Borgo, R. et al., 2012. An Empirical Study on Using Visual Embellishments in Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), pp. 2759-2768.

Borkin, M. et al., 2013. What Makes a Visualization Memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12), pp. 2306-2315.

Brasil: Graphicos Econômicos - Estatísticas, 1929. *Brasil: Graphicos Econômicos - Estatísticas*. Rio de Janeiro: Instituto de Expansão Comercial.

Brinton, W.C., 1939. Pictorial unit bar charts. In *Graphic Presentation*. New York: Brinton Associates, pp. 121-131.

Burke, C., Kindel, E. and Walker, S., 2014. *Isotype: Design and Contexts 1925-1971*. London: Hyphen Press.

Byrde, K.C., 1920. *Our Flags and Their Significance*. Bristol: The Bristol British-Israel Association. Available at: https://archive.org/details/ourflagstheirsigoobyrd.

Byron, L. and Wattenberg, M., 2008. Stacked Graphs – Geometry & Aesthetics. *IEEE Transactions on Visualization and Computer Graphics*, 14(6), pp. 1245-1252.

Card, S.K. and Mackinlay, J.D., 1997. The structure of the information visualization design space. In IEEE Symposium on Information Visualization. Phoenix: IEEE, pp. 92-99.

Card, S.K., Mackinlay, J. and Shneiderman, B., 1999. *Readings in information visualization: using vision to think*. San Francisco: Morgan Kaufmann Publishers Inc.

Catmull, E. and Rom, R., 1974. A class of local interpolating splines. In R. Barnhill and R. Riesenfeld, eds. *Computer Aided Geometric Design*. Salt Lake City: Academic Press, pp. 317-326.

Cawthon, N. and Vande Moere, A.V., 2007. The effect of aesthetic on the usability of data visualization. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. Zurich: IEEE, pp. 637-648.

Chamfort, S.R.N., 1796. *Maximes, pensées, caractères et anecdotes*. London.

Chen, C., 2005. Top 10 unsolved information visualization problems. *IEEE Computer Graphics and Applications*, 25(4), pp. 12-16.

Chen, X., 2011. *Seeing differently: cartography for subjective maps based on dynamic urban data*. Cambridge MA: Massachusetts Institute of Technology.

Chernoff, H., 1973. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68(342), pp. 361-368.

Clark, J.W., 1977. Time-Distance Transformations of Transportation Networks. *Geographical Analysis*, 9(2), pp. 195-205.

Cleveland, W.S. and McGill, R., 1984. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, 79(387), pp. 531-554.

Cleveland, W.S., Harris, C.S. and McGill, R., 1982. Judgments of Circle Sizes on Statistical Maps. *Journal of the American Statistical Association*, 77(379), pp. 541-547.

Costa, J. et al., 2010. Política e negócios. In *Os Donos de Portugal*. Lisbon: Edições Afrontamento, pp. 345-381.

Cox, D., 2006. Metaphoric mappings: The art of visualization. In P.A. Fishwick, ed. *Aesthetic Computing*. Cambridge: The MIT Press, pp. 89-114.

Cruz, P., 2015. Wrongfully right: applications of semantic figurative metaphors in information visualization. In Proceedings of the IEEE VIS Arts Program. Chicago, pp. 14-21.

Cruz, P. and Machado, P., 2010a. Visualizing empires decline. In ACM SIGGRAPH Posters. Los Angeles: ACM, p. 154.

Cruz, P. and Machado, P., 2010b. Visualizing the Portuguese Empire Expansion and Decline. In Proceedings of the 5th International Conference on Digital Arts. Guimarães, pp. 65-71.

Cruz, P. and Machado, P., 2011a. Generative Storytelling for Information Visualization. *IEEE Computer Graphics and Applications*, 31(2), pp. 80-85.

Cruz, P. and Machado, P., 2011b. Visualizing the circulatory problems of Lisbon. In ACM SIGGRAPH Posters. Vancouver: ACM, p. 92.

Cruz, P. and Machado, P., 2014. City Portraits and Caricatures. In D. Offenhuber and C. Ratti, eds. *Decoding the City: How Big Data Can Change Urbanism*. Berlin: Birkhäuser, pp. 96-109.

Cruz, P. and Machado, P., 2016. Pulsing Blood Vessels: A Figurative Approach to Traffic Visualization. *IEEE Computer Graphics and Applications*, 36(2), pp. 16-21.

Cruz, P., Costa, C. and Machado, P., 2014. An Ecosystem of Corporate Politicians. In Proceedings of xCoAx Computation, Communication, Aesthetics and X. Porto, pp. 146-156.

Cruz, P., Cruz, A. and Machado, P., 2015. Contiguous animated edge-based cartograms for traffic visualization. *IEEE Computer Graphics and Applications*, 35(5), pp. 6-13.

Cutroni, J., 2010. *Google Analytics*. Sebastopol, CA: O'Reilly.

Delaunay, B., 1934. Sur la sphere vide. I. Vinogradov, ed. *Bulletin de l'Academie des Sciences de l'URSS. Classe des sciences mathematiques et naturelles*, 6, pp. 793-800.

Dorling, D., 1996. Area cartograms: their use and creation. *Concepts and Techniques in Modern Geography*, 59.

Doyle, P., Cheng, E. and Koh, J., 2011. *All About English, Unit 2 – Using Language to Inform*. Singapore: Pearson Education South Asia.

Fallman, D., 2003. Design-oriented human-computer interaction. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. Ft. Lauderdale, FL: ACM, pp. 225-232.

Fallman, D., 2007. Why Research-Oriented Design Isn't Design-Oriented Research: On the Tensions Between Design and Research in an Implicit Design Discipline. *Knowledge, Technology & Policy*, 20(3), pp. 193-200.

Few, S., 2009. *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. Oakland: Analytics Press.

Few, S., 2011a. Benefitting InfoVis with Visual Difficulties? – Provocation Without a Cause. *Visual Business Intelligence*. Available at: https://www.perceptualedge.com/articles/visual_business_intelligence/visual_difficulties.pdf [Accessed August 18, 2015].

Few, S., 2011b. The Chartjunk Debate – A Close Examination of Recent Findings. *Visual Business Intelligence*. Available at: http://www.perceptualedge.com/articles/visual_business_intelligence/the_chartjunk_debate.pdf [Accessed August 4, 2015].

Few, S., 2013. Chart Junk: A Magnet for Misguided Research. *Visual Business Intelligence*. Available at: https://www.perceptualedge.com/blog/?p=1770 [Accessed August 4, 2015].

Finke, T. and Manger, S., 2012. *Informotion: Animated Infographics*. Berlin: Gestalten.

Flannery, J.J., 1971. The relative effectiveness of some common graduated point symbols in the presentation of quantitative data. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 8(2), pp. 96-109.

Fox, J. and Daily, M.J., 2005. Method and apparatus for geographic shape preservation for identification. US Pattent Office: US 6,847,888 B2.

Friendly, M., 2008. A Brief History of Data Visualization. In *Handbook of Data Visualization*. Berlin: Springer, pp. 15-56.

Fruchterman, T.M.J. and Reingold, E.M., 1991. Graph drawing by force-directed placement. *Software: Practice and Experience*, 21(11), pp. 1129-1164.

Fry, B.J., 2004. *Computational information design*. PhD dissertation, Massachusetts Institute of Technology.

Funkhouser, H.G., 1937. Historical development of the graphical representation of statistical data. *Osiris*, 3, pp. 269-404.

Furnas, G.W., 1986. Generalized fisheye views. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. Boston: ACM, pp. 16-23.

Gardner, M., 1970. Mathematical Games. *Scientific American*, 223, pp. 120-123.

Gastner, M.T. and Newman, M.E.J., 2004. Diffusion-based method for producing density-equalizing maps. *Proceedings of the National Academy of Sciences of the United States of America*, 101(20), pp. 7499-7504.

Gastner, M.T. and Newman, M.E.J., 2005. Density-equalizing map projections: Diffusion-based algorithm and applications. In Proceedings of the 8th International Conference on GeoComputation. Michigan.

Geadas, B. et al., 2015. The hidden connections – Social Network Analysis Portuguese companies and political figures (1950-2014). Report to NOVA Information Management School. Lisbon.

Gershon, N.D. and Page, W., 2001. What storytelling can do for information visualization. *Communications of the ACM*, 44(8), pp. 31-37.

Gondim, H.W.A.S., do Nascimento, H.A.D. and Reilly, D., 2014. Visualizing Large Scale Vehicle Traffic Network Data – A Survey of the State-of-the-art. In Proceedings of the 5th International Conference on Information Visualization Theory and Applications. Lisbon, pp. 337-346.

Gray, J., Chambers, L. and Bounegru, L., 2012. *The data journalism handbook*. Sebastopol, CA: O'Reilly.

Haroz, S., Kosara, R. and Franconeri, S.L., 2015. ISOTYPE Visualization – Working Memory, Performance, and Engagement with Pictographs. In Proceedings of the 33rd ACM Conference on Human Factors in Computing Systems. Seoul: ACM, pp. 1191-1200.

Havre, S., Hetzler, B. and Nowell, L., 2000. ThemeRiver: Visualizing theme changes over time. In Proceedings of the International Symposium on Spatial Data Handling. Salt Lake City: IEEE, pp. 115-123.

Härö, A.S., 1968. Area cartogram of the SMSA population of the United States. *Annals of the Association of American Geographers*, 58(3), pp. 452-460.

Heath, C. and Heath, D., 2008. *Made to Stick*. New York: Random House.

Heer, J. et al., 2008. Creation and Collaboration: Engaging New Audiences for Information Visualization. *Information Visualization*, 4950, pp. 92-133.

Heller, S., 2006. *Nigel Holmes: On Information Design*. New York: Jorge Pinto Books Inc.

Henriques, R., Bação, F. and Lobo, V., 2009. Carto-SOM: cartogram creation using self-organizing maps. *International Journal of Geographical Information Science*, 23(4), pp. 483-511.

Holmes, N., 1984. *Designer's Guide to Creating Charts and Diagrams*. New York: Watson-Guptill Publications.

Holmes, N., 2014. Nigel Holmes on Using Humor to Illustrate Data That Tells a Story. *AIGA*. Available at: http://www.aiga.org/nigel-holmes-on-using-humor-to-illustrate-data-that-tells-a-story/ [Accessed August 2, 2015].

Holmes, N., 2015. Using Humor to Inform. Presented in OPENVIS Conference. Boston, MA.

Holten, D., 2006. Hierarchical Edge Bundles: Visualization of Adjacency Relations in Hierarchical Data. *IEEE Transactions on Visualization and Computer Graphics*, 12(5), pp. 741-748.

Holten, D. and Wijk, J., 2009. Force-Directed Edge Bundling for Graph Visualization. In IEEE-VGTC Symposium on Visualization. Berlin: IEEE, pp. 983-998.

House, D.H. and Kocmoud, C.J., 1998. Continuous cartogram construction. In Proceedings of the Conference on Visualization. Durham, NC: IEEE, pp. 197-204.

Hullman, J., Adar, E. and Shah, P., 2011. Benefitting InfoVis with Visual Difficulties. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), pp. 2213-2222.

Hussey, M., 2010. The city's heartbeat. *Wired UK*, 12, pp. 40-41.

Indurkhya, B., 1992. *Metaphor and Cognition – An Interactionist Approach*. Dordrecht: Kluwer Academic Publishers.

Jakobsen, T., 2001. Advanced character physics. In Game Developers Conference. San Francisco, pp. 383-401.

Jorge, A. and Chambel, T., 2014. Visualizations in Time for a New Look at the Movies. *International Journal of Creative Interfaces and Computer Graphics*, 5(2), pp. 40-61.

Jorge, A., Gil, N. and Chambel, T., 2012. Time for a New Look at the Movies through Visualization. In Proceedings of the 6th International Conference on Digital Arts. Faro, pp. 269-278.

Jorge, A., Serra, S. and Chambel, T., 2014. Interactive Visualizations of Video Tours in Space and Time. In Proceedings of the 28th International BCS Human Computer Interaction Conference on HCI. Southport, UK, pp. 329-334.

Judelman, G., 2004. Aesthetics and Inspiration for Visualization Design: Bridging the Gap between Art and Science. In 8th International Conference on Information Visualization. London, pp. 245-250.

Kamvar, S.D. and Harris, J., 2011. We feel fine and searching the emotional web. In Proceedings of the 4th ACM International Conference on Web Search and Data Mining. Hong Kong, pp. 117-126.

Karsten, K.G., 1923. *Charts and Graphs: An introduction to graphic methods in the control and analysis of statistics*. New York: Prentice-Hall Inc.

Karsten, K.G., 1925. Map. US Pattent Office: US US1556609 A.

Kaspar, S., Fabrikant, S.I. and Freckmann, P., 2011. Empirical Study of Cartograms. In Proceedings of the 25th International Cartographic Conference, CO-112. Paris.

Keahey, T.A. and Robertson, E.L., 1996. Techniques for non-linear magnification transformations. In Proceedings of the IEEE Symposium on Information Visualization. San Francisco: IEEE, pp. 38-45.

Keim, D., North, S.C. and Panse, C., 2005a. Method for generating contiguous cartograms. US Pattent Office: US 6,853,386 B1.

Keim, D., North, S.C. and Panse, C., 2005b. System and method for contiguous images from polygons. US Pattent Office: US 2005/0062763 A1.

Keim, D.A. et al., 2006. Challenges in Visual Data Analysis. *Proceedings of the IEEE Conference on Information Visualization*. Baltimore: IEEE, pp. 9-16.

Keim, D.A., North, S.C. and Panse, C., 2004. CartoDraw: a fast algorithm for generating contiguous cartograms. *IEEE Transactions on Visualization and Computer Graphics*, 10(1), pp. 95-110.

Knott, G.D., 2000. *Interpolating cubic splines*. Boston: Birkhäuser.

Kocmoud, C.J., 1997. Existing continuous cartogram methods. In *Constructing Continuous Cartograms: a constraint-based approach*. MSc thesis, Texas A&M University, pp. 15-25.

Kocmoud, C.J., 1998. A constraint-based approach to constructing continuous cartograms. In Proceedings of the International Symposium on Spatial Data Handling. Vancouver, pp. 236-246.

Krygier, J. and Wood, D., 2011. *Making Maps: A Visual Guide to Map Design for GIS.* 2nd ed., New York: Guilford Press.

Lakoff, G., 1993. The contemporary theory of metaphor. In *Metaphor and Thought*. Cambridge, UK: Cambridge University Press, pp. 202-251.

Lakoff, G. and Johnsen, M., 1980. *Metaphors we live by*. Chicago: The University of Chicago Press.

Lima, M., 2011. *Visual Complexity*. New York: Princeton Architectural Press.

Louçã, F., Lopes, J.T. and Costa, J., 2014. *Os Burgueses.* Lisbon: Bertrand.

Lupton, E. and Phillips, J.C., 2008. *Graphic Design The New Basics*. New York: Princeton Architectural Press.

Lynn, S.K., 2010. Decision-Making and Learning: The Peak Shift Behavioral Response. M. Breed and J. Moore, eds. *Encyclopedia of Animal Behavior*, 1, pp. 450-475.

Mackinlay, J., 1986. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2), pp. 110-141.

Mahyar, N., Kim, S.H. and Kwon, B.C., 2015. Towards a Taxonomy for Evaluating User Engagement in Information Visualization. In IEEE Vis 2015 Workshop Personal Visualization: Exploring Data in Everyday Life. Chicago.

Manovich, L., 2011. What is visualisation? *Visual Studies*, 26(1), pp. 36-49.

McNally, R., 1897. *Universal Atlas of the World*. Chicago and New York: Rand, McNally & Co. Available at: http://digicoll.library.wisc.edu/cgi-bin/History/History-idx?id=History.RandWorldAtlas.

Meirelles, I., 2013. *Design for Information*. Beverly, MA: Rockport Publishers.

Muller, C., 2008. *Metaphors Dead and Alive, Sleeping and Waking: A Dynamic View*, Chicago and London: University of Chicago Press.

Munzner, T., 2008. Process and Pitfalls in Writing Information Visualization Research Papers. In A. Kerren et al., eds. *Information Visualization: Human Centered Issues and Perspectives*. Berlin: Springer, pp. 134-153.

Norman, D.A., 1999. Affordance, conventions, and design. *Interactions*, 6(3), pp. 38-43.

Offenhuber, D., 2002. Wegzeit: the geometry of relative distance. In ACM SIGGRAPH Conference Abstracts and Applications. New York: ACM, p. 280.

Offenhuber, D., 2010. Visual anecdote. *Leonardo*, 43(4), pp. 367-374.

Onken, W.H., Jr and Braymer, D.H., 1921. Electrical importance of the various states. *Electrical World*, 77(12), pp. 650-651.

Ortony, A., 1993. Why Metaphors Are Necessary and Not Just Nice. *Educational Theory*, 25(1), pp. 45-53.

Parkin, G.R., 1892. *Imperial Federation – The Problem of National Unity*. London: Macmillan and Co.

Paullin, C.O., 1932. *Atlas of the Historical Geography of the United States*. J.K. Wright, ed., Baltimore: Carnegie Institution of Washington and American Geographical Society of New York.

Pereirinha, T., 2014. O maior inimigo dos políticos é um nerd. *Sábado*, 510, pp. 82-83.

Pinheiro, R.B., 1885. Apresentação. *Pontos nos ii*, 1, p. 1.

Pinheiro, R.B., 1900a. A Burocracia: a Grande Rata. *A Parodia*, 32, p. 1.

Pinheiro, R.B., 1900b. A Economia: a Galinha Choca. *A Parodia*, 4, p. 1.

Pinheiro, R.B., 1900c. A Politica: a Grande Porca. *A Parodia*, 1, p. 1.

Pinheiro, R.B., 1900d. A Rethorica Parlamentar: o Grande Papagaio. *A Parodia*, 18, p. 1.

Playfair, W., 1801. *The statistical breviary: shewing the resources of every state and kingdom in Europe*. London.

Playfair, W., 1805. *An Inquiry Into the Permanent Causes of the Decline and Fall of Powerful and Wealthy Nations*. London: Greenland and Norris.

Polisciuc, E. et al., 2015. Arc and Swarm-based Representations of Customer's Flows among Supermarkets. In Proceedings of the 6th International Conference on Information Visualization Theory and Applications. Berlin, pp. 300-306.

Pousman, Z., Stasko, J.T. and Mateas, M., 2007. Casual Information Visualization: Depictions of Data in Everyday Life. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), pp. 1145-1152.

Press, W.H. et al., 2007. *Numerical Recipes.* 3rd ed., Cambridge, UK: Cambridge University Press.

Purchase, H.C. et al., 2008. Theoretical Foundations of Information Visualization. In *Information Visualization: Human centered issues and perspectives*. LNCS 4950. Berlin: Springer, pp. 46-64.

Raisz, E., 1934. The rectangular statistical cartogram. *Geographical Review*, 24(2), pp. 292-296.

Raisz, E., 1948. *General Cartography.* 2nd ed., New York, Toronto, London: McGraw-Hill Book Co. Inc.

Ramachandran, V.S. and Hirstein, W., 1999. The science of art: A neurological theory of aesthetic experience. *Journal of Consciousness Studies,* 6(7), pp. 15-51.

Reis, M.F., 2014. Os donos de Portugal, agora em versão ecossistema interactivo. *Jornal i*, 5(1462), pp. 16-19.

Richards, I.A., 1936. *The Philosophy of Rhethoric*. New York: Oxford University Press.

Sagar, B.S.D., 2014. Cartograms via mathematical morphology. *Information Visualization*, 13(1), pp. 42-58.

Santos, D., 2013. *Geografia das redes – o mundo e seus lugares*. São Paulo: Editora do Brasil.

Sedlák, P., 2011. *Enter 5 – Datopolis*. Prague: CIANT.

Skau, D., Harrison, L. and Kosara, R., 2015. An Evaluation of the Impact of Visual Embellishments in Bar Charts. *Computer Graphics Forum*, 34(3), pp. 221-230.

Snyder, J.P., 1987. Mercator projection. In *Map projections – A working manual*. Washington D.C.: U.S. Geological Survey, pp. 39-47.

Sugiura, K., 1973. Japan time-travel map. *Shukan Asahi*.

Sun, H. and Li, Z., 2010. Effectiveness of Cartogram for the Representation of Spatial Data. *The Cartographic Journal*, 47(1), pp. 12-21.

Süli, E. and Mayers, D., 2003. *An Introduction to Numerical Analysis*. Cambridge: Cambridge University Press.

Tao, M., 2010. *Using Cartograms in Disease Mapping*. PhD dissertation, The University of Sheffield.

Tobler, W.R., 1973. A Continuous Transformation Useful for Districting. *Annals of the New York Academy of Sciences*, 219(1), pp. 215-220.

Tobler, W.R., 1986. Pseudo-Cartograms. *Cartography and Geographic Information Science*, 13(1), pp. 43-50.

Tobler, W.R., 2004. Thirty five years of computer cartograms. *Annals of the Association of American Geographers*, 94(1), pp. 58-73.

Tufte, E.R., 1983. *The visual display of quantitative information*. Cheshire, CT: Graphics Press.

Tufte, E.R., 1990. *Envisioning Information*. Chesire, CT: Graphics Press.

Vande Moere, A.V. and Purchase, H., 2011. On the role of design in information visualization. *Information Visualization*, 10(4), pp. 356-371.

Vicente, J.L. and Girardin, F. eds., 2010. *Mediateca Expandida – Habitar*. Gijón: LABoral Centro de Arte y Creación Cultural.

Viegas, F. and Wattenberg, M., 2007. Artistic data visualization: Beyond visual analytics. *Online Communities and Social Computing*, LCNS 4564, pp. 182-191.

Viegas, F. et al., 2007. ManyEyes: a Site for Visualization at Internet Scale. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), pp. 1121-1128.

Ware, C., 2004. *Information Visualization – Perception for Design.* 2nd ed., San Francisco: Elsevier.

Wattenberg, M., 2002. Arc diagrams: visualizing structure in strings. In Proceedings of the IEEE Symposium on Information Visualization. Boston: IEEE, pp. 110-116.

Wood, J. et al., 2012. Sketchy Rendering for Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), pp. 2749-2758.

Wu, Y.H. and Hung, M.C., 2010. Non-Connective Linear Cartograms for Mapping Traffic Conditions. *Cartographic Perspectives*, 65, pp. 33-50.

Zacks, J. and Tversky, B., 1999. Bars and lines: A study of graphic communication. *Memory & Cognition*, 27(6), pp. 1073-1079.

Ziemkiewicz, C., 2010. *Understanding the structure of information visualization through visual metaphors*. PhD dissertation, University of North Carolina.

Ziemkiewicz, C. and Kosara, R., 2008. The Shaping of Information by Visual Metaphors. *IEEE Transactions on Visualization and Computer Graphics*, 14(6), pp. 1269-1276.