Interactive Evolution of Swarms for the Visualisation of Consumptions

Catarina Maçãs[(✉)], Nuno Lourenço, and Penousal Machado

CISUC - Department of Informatics Engineering,
University of Coimbra, Coimbra, Portugal
{cmacas,naml,machado}@dei.uc.pt

Abstract. Information Visualisation studies how visual representations can help understanding hidden patterns in large amounts of data. The produced visual artefacts should have both functional and aesthetic dimensions to make the visualisation appealing to the user. However, in the Data Aesthetics field, the process of creation of visualisations is more concerned with aesthetics. Our goal for this project is to develop a framework to explore the aesthetic dimension of a functional visualisation model characterised by a series of parameters, which can make the visualisation more functional or more aesthetically appealing. In concrete, we propose a framework based on Interactive Evolutionary Computation (IEC) to evolve the parameterisation of the visualisation model, enabling the user to explore new possibilities and to create different aesthetics over the data. Our case study will be a dataset with the consumption patterns of the Portuguese people in one retail company. The developed system is able to create a wide diversity of emergent visual artefacts that can be intriguing and aesthetically appealing for the user.

Keywords: Data aesthetics · Genetic algorithms
Evolutionary computation · Swarm systems · Visualisation

1 Introduction

In a partnership with a Portuguese retail company, we had access to a high volume of data about the Portuguese's consumption. This retail company's goal is to communicate with a wider public through aesthetic experiences. For this reason, we aim to create emergent visual artefacts driven by the consumption data, positioning our work in the field of Data Aesthetics [1,2]. The dataset is rich in daily, weekly and monthly repetitions of consumption patterns, offering us the opportunity to transform the consumption of the Portuguese into visual artefacts, while exploring, highlighting and visualizing their periodic nature.

Taking into account a previous project [3], we use a swarm system to create emergent visual artefacts [4]. In the referenced work, the user had to define the parameterisations of the system to create a balance between a more functional or more aesthetically intriguing visualisation. In this project, this parameterisation

is automatically defined through the use of an evolutionary system, opening the possibilities to create a wider range of visual solutions. Whence, in this project the main concern is to create artefacts that are aesthetically appealing to the user, and not artefacts placed in the functional spectrum.

The result of this exploration is an automatic framework that is able to create a wide and diverse set of solutions. We test the validity of the system through different usage scenarios in which the system relies on the users' preferences as an input to guide the solution towards what they find attractive. We defined three objectives for the user guidance: explore specific parameterisation attributes; guide the evolution to a functional artefact; explore randomly the system.

The remainder of the article is structured as follows. In Sect. 2, we introduce the project, present the dataset used, the project objectives, and the work previously made within this project context. In Sect. 3, we present the visualisation model (Subsect. 3.1), and the Evolutionary Algorithm (Subsect. 3.2) used to generate the visual artefacts. In Sect. 4 we introduce the usage scenarios and discuss the results, effectiveness, and diversity of our approach. Finally, in Sect. 5, we define the future work.

2   Background

This article is part of a larger project developed in association with SONAE, a Portuguese retail company. Our dataset consists of 278 GB of information about customer purchases in 729 Portuguese supermarkets and hypermarkets of the company's retail chains, in a time span of 24 months (from May 2012 to April 2014).

While shopping in the retail stores, costumers tend to use their client cards to accumulate discounts and other benefits. This enables the company to create personalised discounts and to aggregate data by specific geolocations. The dataset comprises approximately 2.86 billions transactions with the following attributes: customer card id, amount spent, product designation, quantity of the purchased product, and date and time of the transaction. Each product is placed within the product hierarchy of the company, which has 6 levels—Department, BizUnit, Category, Sub-Category, Unit Base, Product.

As a guideline for the project, the company was interested in two main dimensions, one related to the analytical analysis of their data and another related to aesthetics, giving us the opportunity to freely explore the last. In the analytical dimension, a set of visualisations were already implemented concerning problems, such as the understanding of how consumption evolves through time in specific regions of Portugal or the identification of potential sites to open new supermarkets, allowing the understanding of seasonal variations [5–9].

In the aesthetic dimension of the project, we have already developed a set of works. Although they are intended to explore the aesthetics of visualisation, they are also concerned with functionality, trying not to overpass the barrier of legibility. In a previous work [10], we explored how the visualisation can morph depending on the data, creating movement and highlighting the rhythm and disruptions of the normal consumption patterns. We also explored the qualitative

representations of data, providing an overview of how data behaves through time [3]. In the latter, we apply a swarm based system as a method to create emergent visualisations of the consumption values with the intent to convey meaningful information and, at the same time, explore the boundaries between Information visualisation and Data Aesthetics. The application of swarm systems to visualize data can also be seen in [11–13]. Additionally, in the field of Generative Art, swarm forces were also used in a variety of projects [14–16]. We focused on the ability of this emergent system to communicate information while engaging the viewer with organic visuals [4]. Additionally, with the different parameterisations, we were able to create a set of renderings with different levels of legibility and attractiveness. The approach presented in this article builds on this work, by developing a framework that is able to evolve the configuration of the visualisation model through the use of an Evolutionary Algorithm (EA) [17]. Our goal is to create new, diverse, and surprising visual artefacts, which, in spite of not being completely functional, are engaging and entertaining for the user.

3   Evolutionary Approach

To improve the swarm system detailed in [3], we apply an EA to increase the degrees of freedom in the creation of visual models, enabling a diverse range of solutions. Our intention with this approach is to develop visual artefacts that are continuously readjusting to the intentions of the user and to enable the user to explore artefacts not imagined by him/her. In this way, the company can deploy a system that will evolve and adapt to different audiences' aesthetic preferences.

3.1   Visualisation Model

The visualisation model that serves as the basis for this project, consists of a swarm system that, through different forces of separation, attraction, and cohesion, creates emergent visualisations [4] about the Portuguese's consumption routines. This system is constituted by several boids—artificial objects that simulate the flocking behaviour of birds—in an environment (i.e., the canvas), that react to the changes in consumption over time.

For this project, we aggregate all the transactions of the dataset by the highest level in the hierarchy, i.e., the Department. There are a total of seven different Departments in the dataset: Grocery, Fresh Food, Food & Bakery, Home, Leisure, Textile, and Health. Each transaction has the hour, minute, and second of purchase. However, the representation with this degree of detail would be too subtle to the human eye. Therefore, we aggregate the data in intervals of two hours, resulting in 12 intervals per day. Although the transactions occur only between 8 am and 10 pm, we keep the representation of the 24 h aggregated every two hours so it is easier to distinguish different days of consumption.

Swarming Forces. The swarm system simulates the behaviour of multiple boids [18]. Each boid is represented through a circle and is described by properties such as velocity, position, size, and colour. Only the position and velocity of the boids are affected by the swarming forces.

Based on the work of Reynolds [18], each boid follows three basic rules: (i) cohesion; (ii) separation; and (iii) alignment. To explore the system, we applied different values to each force, so it was possible to create different outputs. If two neighbouring boids are from the same Department, we apply higher forces of attraction and lower forces of separation, but, if they are from different Departments, the attraction force is lower then the previously defined, and the separation force higher.

To prevent the boids from randomly moving on the canvas and to enhance their periodic behaviour, we defined a target that all boids should look for. Hence, in addition to the previously described forces, all boids are under the influence of an attraction force towards this moving target. The target boid, although not represented visually, starts from the centre of the canvas and swirls around, creating a spiral with equal distances between each lap.

Since we want to represent a time-series data, the representation of time must be added to the visual artefacts. We consider that each lap of the swirling boids represents one month of data. Then, depending on the angle of the boid with the center of the canvas, we obtain the day and corresponding hour of the month. To do so, the 360° are divided by the 31 days multiplied by 12 h ($2\Pi \div (31 \times 12)$), since, as previously stated, the data is aggregated in intervals of two hours. Note that all laps have 31 days. By doing so, all months start with the same angle (at the top of the circle), and, if they have less than 31 days, the consumption values are null during those nonexistent days.

Grocery   **Leisure**   **Food & Backary**   **House**   **Fresh Food**   **Health**   **Textile**

Fig. 1. Colours used to distinguish the boids representing one of the seven Departments in the company's product hierarchy. (Color figure online)

Rendering. As stated before, each boid is represented through a circle that has a specific size and colour. While colour identifies to which Department the boid belongs to (Fig. 1), the size represents the consumption in a certain time: the bigger the circle, the bigger the consumption value. As the boids wander through the canvas, they leave an imprint of their shape, enabling the user to see its path, and consequently, the consumption values. The boids' size is mapped to a predefined minimum and maximum radii, that can represent the minimum and maximum sale of each individual Department (local normalisation), or the minimum and maximum sale of all Departments (global normalisation). Additionally, we defined three different styles to represent the boids: through a filled circle, through the outline of the circle, and through a line that connects all boids of the same department. Each style is painted with the colour of the corresponding Department (Fig. 1). Additionally, we implemented a mechanism to

sort the circles in depth according to their radii. With this, the smaller circles are drawn over the larger ones and are never hidden by the larger ones.

3.2   Evolutionary Algorithm

The visualisation model described above is easy to understand and use, yet it requires the definition of several parameters to create visual artefacts that can be appealing for the user. To aid in this task, we propose a framework based on EAs [19]. To evolve the swarm system, we will be searching for the best combination of the following system's parameters: (i) the separation, alignment, and cohesion forces; (ii) the minimum and maximum radius; (iii) the use of a global or local normalisation; and (iv) the representation modes (lines, circles, transparency, sorted circles). Note that the boids' size is always mapped according to the consumption value on the data. In the following subsections, we present the parameters used to evolve the visual artefacts and the used genetic operators.

Representation. Each EA solution is encoded as a set of values that correspond to the number of parameters needed by the swarm system. In concrete, we have 10 different parameters that are required by the visual model:

– **Separation Force**: a real value between 0 and 3;
– **Alignment Force**: a real value between 0 and 3;
– **Cohesion Force**: a real value between 0 and 3;
– **Boid Render**: one of the following options: lines, circles and filled circles;
– **Transparency**: a boolean value that enables transparency of the boids;
– **Size Ordered**: a boolean value. If this value is true, the visualisation model sorts the boids by radius, i.e., boids with smaller radius will be on top of boids with larger radius;
– **Mapped**: a boolean value. If it is true, it indicates that the separation force is mapped depending on the radius of the boids;
– **Normalisation**: a boolean value that enables normalisation based on the maximum sales values;
– **Maximum Radius**: a real value between 30 and 80 that corresponds to the maximum radius of the boids;
– **Minimum Radius**: a real value between 0.1 and 15 that corresponds to the minimum radius of the boids.

An example of a possible solution alongside with its phenotype representation is depicted in Fig. 2.

Genetic Operators. To promote the evolution and the proper exploration of the search space we rely on two operators: recombination and mutation. The recombination operator is the uniform crossover and combines two solutions by creating a random mask of the same size of the genotype, and then swap the genetic material according to the previously generated mask. Regarding the mutation operator, we apply a per gene mutation to the candidate solutions. This allows the algorithm to change, from generation to generation, a significant percentage of the genes to other valid ones.
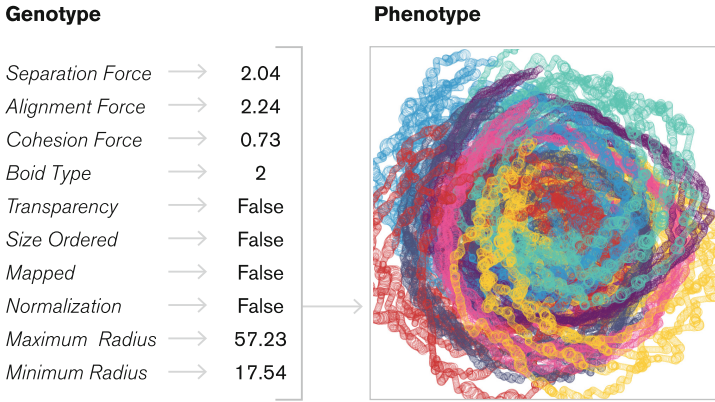
Genotype   **Phenotype**

Separation Force	$\longrightarrow$	2.04
Alignment Force	$\longrightarrow$	2.24
Cohesion Force	$\longrightarrow$	0.73
Boid Type	$\longrightarrow$	2
Transparency	$\longrightarrow$	False
Size Ordered	$\longrightarrow$	False
Mapped	$\longrightarrow$	False
Normalization	$\longrightarrow$	False
Maximum Radius	$\longrightarrow$	57.23
Minimum Radius	$\longrightarrow$	17.54

Fig. 2. On the left, the genotype of a possible solution; On the right, the resulting phenotype, i.e., the visualisation model.

4   Usage Scenarios

Through Interactive Evolutionary Computation (IEC) [20], we explored the evolutionary system based on the user's preferences. We implemented a simple interface to enable the interaction with the system[1]. These explorations were based on three different objectives. In the first exploration, the goal is to evolve solutions with specific parameterisation attributes: the boids must be represented with filled circles, use the local normalisation, and have a zigzagging pattern. In the second, the goal is to attain solutions that must intersect the functional dimension, enabling the readability of the artefacts. Finally, for the third exploration, there is no predefined objective, so the solutions must be diversified according to the user's taste. For the first two explorations, the user must have some experience on how the system works, its parameters and the data. In the last, the user can have no experience with the data nor with the parameterisations, the user only explores the system and creates artefacts suitable for his/her own taste. As the generations evolve, the user chooses the visual artefacts that visually intrigue, amaze, and/or correspond with his/her preferences.

To avoid user fatigue, we used a reduced number of individuals per generation. For each exploration, the number of individuals is set to 20 and the maximum generations to 10 (this value can be increased). Additionally, the user can stop the evolutionary process at any given time if the visual artefacts are in accordance to his/her expectations.

In the first exploration, and as the user has already a target solution in mind, it is possible to perceive that the evolutionary system evolved correctly towards the user predefined goal (Fig. 3). As the individuals are being created, the user selects only the ones with filled circles, and the ones where all colours appeared balanced, corresponding to the local normalisation where the values are

---

[1] A video of this interface can be seen in: https://vimeo.com/289093672.

Fig. 3. Chosen individuals in: (left) first generation, (center) third generation, and (right) sixth generation.


Fig. 4. Chosen individuals in: (left) first generation, (center) third generation, and (right) sixth generation.

normalised by Department. When using the global normalisation, the artefacts colours consist mainly of blues and greens because these colours correspond to the Departments with higher consumption values, and thus, the ones with more visual presence.

For the second exploration, and with the intention to guide the evolutionary system to generate artefacts which can be aesthetic and at the same time functional, the system also proved to be capable of evolving good solutions (Fig. 4). In the last generations, the majority of the individuals were readable. This means that they had less clutter and lower separations forces, which enabled them to have a more strict behaviour as they swirl, not deviating from the spiral path. Additionally, the majority of the generated artefacts are similar to the explorations made in the previous project [3].

For the last exploration, the user had no predefined objective. As the artefacts were being presented, the user chose freely the preferred ones. In the end, the user managed to guide the system through a specific style, which was appealing for him/her—the use of lines to represent the Departments (Fig. 5).

As a summary, we found that the system was capable of discovering artefacts similar to the ones created by human designers, but also to find new ones. Furthermore, the system can be guided by the user, showing artefacts similar to the user choices, but, at the same time, giving to the user others to enhance his/her choice.

Fig. 5. Chosen individuals in: (left) first generation, (center) third generation, and (right) sixth generation.


Fig. 6. Different individuals with similar parameterisations.

4.1   Results' Diversity

One of the important aspects that we looked for in the previous experiments was the capacity of the system to generate and evolve a diversified range of visual artefacts. The random initialisation of the visualisation models creates a wide variety of behaviours. For example, if the forces are too strong, the boids will deviate from the spiral path, creating random zigzagging patterns. In a similar way, these forces can also cause the boids to stagnate at their position in the centre. In Fig. 6, it is possible to see how the system generates a set of individuals, which are different but based on the choices made previously by the user.

As the system evolves the parameters, without the direct intervention of the user, it can generate solutions which the user was not expecting and has not seen before. Even an experienced user, that knows the parameters and the system itself, can be surprised by the solutions found[2].

5   Conclusion

Over the last few years, we have seen an exponential increase in data in all sectors of business. Companies have seen this as an opportunity to improve their operations and increase the satisfaction of their costumers. In previous works, we have proposed, developed, and implemented visualisation models for one of the biggest retail companies in Portugal. We have developed a set of artefacts that refined the way the company looked at their data in order to improve their

---

[2] In https://cdv.dei.uc.pt/ie-of-swarms-in-visualisation it is possible to see other visual solutions.

business strategy. In this article, we extend a previous work [3] and explore an automatic manipulation of a swarm system through the application of an Evolutionary Algorithm. Our main goal was to explore the aesthetic dimension of the consumption data provided by the Portuguese retail company.

As future work, we intend to improve the interactive evolutionary system by enabling it to learn the previous choices of the user and guide the evolution based on those parameters. With this, we can also augment the number of individuals per generations and, in order to diminish user fatigue, show only the fittest (based on previous choices) for the user to choose.

References

1. Manovich, L.: Information and form (2000). http://manovich.net/index.php/projects/information-and-form
2. Moere, A.V.: Aesthetic data visualization as a resource for educating creative design. In: Dong, A., Moere, A.V., Gero, J.S. (eds.) Computer-Aided Architectural Design Futures, CAADFutures 2007, pp. 71–84. Springer, Dodretch (2007). https://doi.org/10.1007/978-1-4020-6528-6_6
3. Maçãs, C., Cruz, P., Martins, P., Machado, P.: Swarm systems in the visualization of consumption patterns. In: Yang, Q., Wooldridge, M. (eds.) Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, 25–31 July 2015, Buenos Aires, Argentina, pp. 2466–2472. AAAI Press (2015)
4. Jones, D.: Swarm aesthetics: a critical appraisal of swarming structures in art practice. Master's thesis, MA Sonic Arts (2007)
5. Maçãs, C., Cruz, P., Polisciuc, E., Amaro, H., Machado, P.: ISO-edges for the geovisualization of consumptions. In: Braz, J., Kerren, A., Linsen, L. (eds.) IVAPP 2016 - Proceedings of the 11th Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, pp. 220–227. SciTePress (2016)
6. Polisciuc, E., et al.: Arc and swarm-based representations of customer's flows among supermarkets. In: Braz, J., Kerren, A., Linsen, L. (eds.) IVAPP 2015 - Proceedings of the 6th International Conference on Information Visualization Theory and Applications, 11–14 March 2015, Berlin, Germany, pp. 300–306. SciTePress (2015)
7. Polisciuc, E., Maçãs, C., Assunção, F., Machado, P.: Hexagonal gridded maps and information layers: a novel approach for the exploration and analysis of retail data. In: SIGGRAPH ASIA 2016 Symposium on Visualization, pp. 6:1–6:8. ACM (2016)
8. Polisciuc, E., Cruz, P., Maçãs, C., Amaro, H., Machado, P.: Flow map of products transported among warehouses and supermarkets. In: Braz, J., Kerren, A., Linsen, L. (eds.) Proceedings of the 11th Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, vol. 2, pp. 177–186. SciTePress (2016)
9. Maçãs, C., et al.: Time-series application on big data - visualization of consumption in supermarkets. In: Braz, J., Kerren, A., Linsen, L. (eds.) IVAPP 2015 - Proceedings of the 6th International Conference on Information Visualization Theory and Applications, 11–14 March 2015, Berlin, Germany, pp. 239–246. SciTePress (2015)

10. Maçãs, C., Machado, P.: The rhythm of consumptions. In: Expressive 2016 - Proceedings of the Joint Symposium Expressive 2016, co-located with the Eurographics 2016. ACM (2016)

11. Moere, A.V., Lau, A.: In-formation flocking: an approach to data visualization using multi-agent formation behavior. In: Randall, M., Abbass, H.A., Wiles, J. (eds.) Australian Conference on Artificial Life, ACAL 2007. LNCS, vol. 4828, pp. 292–304. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76931-6_26

12. Ogawa, M., Ma, K.L.: code_swarm: a design study in organic software visualization. IEEE Trans. Vis. Comput. Graph. **15**(6), 1097–1104 (2009)

13. Kamvar, S.D., Harris, J.: We feel fine and searching the emotional web. In: Proceedings of the fourth ACM International Conference on Web Search and Data Mining, pp. 117–126. ACM (2011)

14. Bornhofen, S., Gardeux, V., Machizaud, A.: From swarm art toward ecosystem art. Int. J. Swarm Intell. Res. (IJSIR) **3**(3), 1–18 (2012)

15. Greenfield, G., Machado, P.: Swarm art: Curator's introduction. Leonardo **47**(1), 5–7 (2014)

16. Shiffman, D.: Swarm (2004). http://shiffman.net/projects/swarm/

17. Eiben, A.E., Smith, J.E.: Introduction to Evolutionary Computing. NCS, vol. 53. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-662-05094-1

18. Reynolds, C.W.: Flocks, herds and schools: a distributed behavioral model. In: ACM SIGGRAPH Computer Graphics, vol. 21, pp. 25–34. ACM (1987)

19. Eiben, A.E., Smith, J.E.: Introduction to Evolutionary Computing. NCS. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-44874-8

20. Takagi, H.: Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. Proc. IEEE **89**(9), 1275–1296 (2001)