EvoFashion: Customising Fashion Through Evolution

Nuno Lourenço, Filipe Assunção, Catarina Maçãs, and Penousal Machado

CISUC, Department of Informatics Engineering, University of Coimbra,
Pólo II - Pinhal de Marrocos, 3030 Coimbra, Portugal
{naml,fga,cmacas,machado}@dei.uc.pt

Abstract. In today's society, where everyone desires unique and fashionable products, the ability to customise products is almost mandatory in every online store. Despite of many stores allowing the users to personalize their products, they do not always do it in the most efficient and user-friendly manner. In order to have products that reflect the user's design preferences, they have to go through a laborious process of picking the components that they want to customise. In this paper we propose a framework that aims to relieve the design burden from the user side, by automating the design process through the use of Interactive Evolutionary Computation (IEC). The framework is based on a web-interface that facilitates the interaction between the user and the evolutionary process. The user can select between two types of evolution: (i) automatic; and (ii) partially-automatic. The results show the ability of the framework to promote evolution towards solutions that reflect the user aesthetic preferences.

Keywords: Evolutionary Algorithm, Fashion Design, Interactive Evolutionary Computation, Product Customisation

1 Introduction

In today's society many online stores allow customers to personalise their products during their purchase so that each individual can carry unique, stylish and distinctive items [1]. This customisation usually includes letter engraving and selection of colours and/or materials to specific parts of the product. Nevertheless, this is often a lingering process, where the buyer is asked endless questions about his/her preferences. Additionally, the customisation of the product can have a direct impact in other aspects of the item, such as its price or availability, making the process even more difficult and not user friendly. In order to aliviate the design burden from the customers side, and ease the customisation of products, we propose a framework based on Evolutionary Algorithm (EA). We propose an automatic tool to measure the user's preferences, capable of speeding the process of personalising items and adjusting the products parameters. To allow an effortless integration with online stores, we show the viability of the approach by proposing a web framework that allows the customisation of fashion products.

Despite the wide range of practical applications of this type of frameworks, we choose the evolution of shoe design as a case-study. Our experiments consider the use of an automatic and user-guided fitness function to guide the evolutionary process. Results show the viability of the proposed framework, evincing that the combination of automatic and user-guided fitness functions is advantageous, since it alleviates the customisation burden from the user, and, at the same time, allows the emergence of unique pieces specially tailored for the user's taste.

The remainder of the paper is organised as follows. Section 2 details on methodologies applied to the customisation and optimisation of products based on EA. Our proposal is detailed in Section 3, followed by the experimental analysis (Section 4). Finally, Section 5 gathers the main conclusions and points towards future work.

2 Related Work

EAs [2] are computational models that have their origins in the early 1960's. At the time, and inspired by theories of natural-selection (Charles Darwin) and of Mendelian inheritance (Gregor Mendel), researchers started to develop different techniques to evolve artificial systems. These theories are based on the idea that the best adapted individuals have higher chances of survival and, as such, higher chances of reproducing and passing their characteristics to the next generation of individuals. As generations pass, individuals will become more adapted to the environment. EAs have been successfully applied to different domains from science to art. In the next paragraphs, focus will be given to works within the scope of this paper, i.e., the application of EAs in works concerning industry and design optimisation and customisation.

There are several examples of the application of EAs to optimise the shape and efficiency of an object. In architecture, EAs can be applied to improve the energy efficiency and design of a house. In [3], Caldas and Norford use Evolutionary Computation (EC) to evaluate several aspects concerned with the energetic performance of buildings, such as window placing and size. Still focusing on architectural efficiency aspects, in [4] Besserud and Cotten, evolved the optimal structure of a 300-meter tower, targeting the maximization of sun rays incidence for the installation of photovoltaic panels. Concerning the evolution of designs, Michalek et. al [5] automatically generate the layout of floorplans, and Juan et. al [6] propose a system that enables house customisation by combining case-based reasoning with a Genetic Algorithm (GA).

Several other works related with the optimisation and customisation of products are present in the literature. For example, in [7] the design of antennas is addressed, with the objective of increasing the efficiency of the design process. In [8], [9] and [10] several methodologies for design of layouts are introduced. Morcillo et. al [8] used the principles and techniques of EC together with Fuzzy Logic to automatically obtain the layout of a single-page design. Morcillo's approach was not able to learn different styles, and in 2014 O'Donovan et. al [9] developed an approach to automatically create graphic design layouts through a new energy-based model derived from design principles. With this method, it was possible to arrange a set of elements in a particular style, previously learned from the analysis of designs presented to the system. In [10] the layout of a set of album images is evolved.

In the majority of the aforementioned works, it is possible to define a fitness function that measures how far we are from our objective. However, when dealing with the design of wearables or with customisable products that are to be purchased by masses, the definition of quality is subjective to the buyer and to its notion of aesthetics, making it much harder to define. With the increasing importance of online shopping, the application of EC to allow and facilitate the customisation and personification of products, have emerged. Whilst the majority of the previous works use automatic fitness assigning schemes, when dealing with fashion it is common to involve the customer in the evolutionary process. Such methodologies are normally referred to as Interactive Evolutionary Computation (IEC) [11] approaches, where the user is involved in the evolutionary process to define the fitness of a certain candidate solutions.

An example of the application of IEC to fashion is in the combination of clothing parts to generate new pieces. In [12] Kim et al. used IEC to generate unique dress designs from separated parts, previously classified in three categories. Through user interaction, it is possible to set the fitness and create outputs that reflect personal preferences. Khajeh et al. [13] also applied this method to enhance the creativity of the designer in order to produce novel sets of clothes design. In their work, clothes components and fabric patterns were designed and encoded separately. Then, through user interaction and the application of fashion design rules, the system combines the components and the fabric patterns. The system proved to be efficient in the generation of fashion designs, minimising the cost and time according to the user preferences. Another important aspect in fashion is concerned with reduction of fabric waste; in [14] and [15] the nesting problem is addressed. Still focusing on fashion, but more specifically on shoe design, in [16], Shimoyama et. al use EC to optimise the sole structure of a shoe. In [17] Dasan, aided by IEC evolves shoe models according to the user preference.

A problem inherent to the majority of IEC approaches concerns user weariness and consistency, i.e., it is difficult for the user to efficiently rank and compare all the individuals in large populations, maintaining the same classification criteria throughout the entire process. Therefore, in the current work we aim at developing a framework that reduces the need for interaction, but at the same time allows customers to taylor their products. Similarly to other approaches, the user interacts with the system to feed his(the preferences. However, in addition to common IEC approaches, we follow the same guidelines behind the work by Machado et. al [18]. The user will not need to classify every single candidate solution and will be able to set filter values, reducing the necessary number of clicks and interaction. An archive where the user can store the best solutions according to his/her preferences is also provided.

Front		Side		Sole		Lining		Zipper Tape		Heel		Straps		Straps Tips		Hardware	
3	11	4	22	0	2	0	1	0	1	4	2	0	2	0	0	0	1

Fig. 1: On top, the genotype of a candidate solution; On the bottom, the resulting phenotype.

3 Framework

EvoFashion is a customisation framework that uses principles of IEC to promote the personalisation of fashion products. Despite the wide range of applicability in fashion domains, we will focus our attention on the evolution of the design of shoe models due to the availability of an engine capable of rendering them [1]. In the following sub-sections we detail the main components of the developed EA and the web interface.

3.1 Representation

Solutions are encoded as a set of integers with the same length as the number of parameters allowed in the customisation of a certain product. Each gene (i.e., integer) has a value in the $[0, num\_possibilities]$ interval, where $num\_possibilities$ is the maximum number of different possibilities for that parameter.

Our experiments will focus on the customisation of a specific shoe model, with 9 customisable parts. Each part allows personalisation of two properties: material and colour. Consequently, the candidate solutions will be composed by $9 \times 2 = 18$ integers. An example of a possible solution along with its phenotypic representation is depicted in Figure 1.

3.2 Genetic Operators

To promote the evolution and the proper exploration of the problem domain we rely on recombination and mutation. We use uniform crossover to recombine two parents. Firstly we create 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

[1] In the current work we have used the my-swear platform to render the generated shoes, which can be found in `https://www.my-swear.com/`.

solutions, which allows the algorithm to change, from generation to generation, a percentage of the genes to other valid ones.

3.3 Fitness Assessment

The main goal of the framework is to promote the evolution of candidate solutions towards regions of the search space that the customer sees as aesthetic. To accomplish that, practitioners often use IEC systems, asking the user to rank the evolved solutions. However, this might lead to some disparities, as the user needs to look at every single evolved solution and compare it with the remaining ones. For that reason, we have used principles of partially interactive computation, which combines automatic fitness components, such as filters, with the interaction between the framework and the user. Three automatic fitness components are used, and they work as filters for the colour, material and price of the evolved models. Further details are presented in the next paragraphs.

Colour – the user defines a colour that he/she likes and the goal of the evolution is to promote the convergence of all parts of the generated shoes to that colour. For that, the Root Mean Square Error (RMSE) is used to compute the distance between a snapshot of the generated shoe models and the target colour, defined by the user;

Material – similar to the colour, but for materials, i.e., the objective is to converge to an individual where all shoe parts are made of the material selected by the user. As with colour, the evolution is guided using an error that is to be minimised and represents the percentage of shoe parts that are not made using the selected material;

Price – by selecting a price range the framework promotes the emergence of candidate solutions that are within that interval, using the RMSE to evaluate the distance between the evolved shoe models' price and the closer bound of the target price range.

From the previous, we propose the following model to compute the quality of each individual:

$$fitness = \frac{1}{1 + automatic\_fitness} + user\_input,$$

where,

$$automatic\_fitness = colour + material + price.$$

That is, the fitness is made of two distinct parts: (i) the automatic evolution towards user selected criteria; and (ii) a user input that works as a weight for ranking the individuals. Both parts are optional, but at least one of the fitness components must be provided, otherwise all individuals would be assigned the same fitness value. In the first part, if one of the automatic components is not defined its RMSE value is 0.

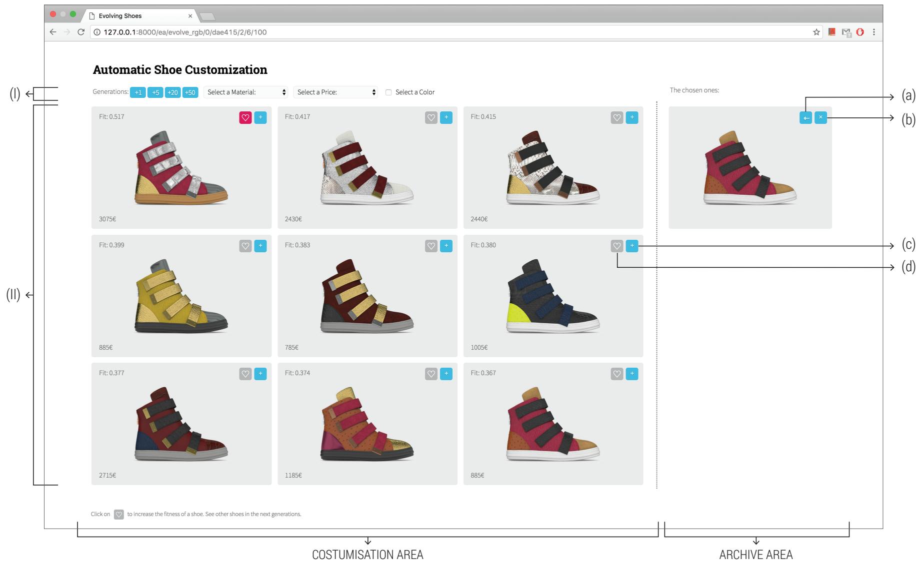

Fig. 2: The Customisation Area, on the left side, is divided in two sub-areas: the evolutionary process management (I) and the visualisation of the population (II). In the last, the user can increase the quality of individual shoes (d) and/or add them to the archive (c). In the Archive Area, the user can: (a) add the saved shoes back to the population; and (b) remove archive members.

3.4 Archive

The archive works as a repository for the individuals that the user finds aesthetically appealing. If, during the evolutionary process, the user finds a good solution, he can store it in the archive and re-introduce it later in the population. The archive also helps avoiding convergence, by increasing the diversity at the population level; for that, if solutions start being too similar, without having achieved the desired result, the user just needs to repopulate the generation with archived members.

3.5 Interface

To enable the visualisation and exploration of the different individuals and to allow the interaction with the user, we developed a web-based application. This application allows the user to define different filters, such as colour or price, to improve the fitness of the individuals by clicking on them, and to store them in the archive. Figure 2 shows the different areas of the interface and buttons functionality.

To simplify navigation, the interface is a single page divided into two main areas: the Customisation Area and the Archive Area. This way, the user is always

Table 1: Experimental parameters.

Parameter	Value
Number of runs	30
Population size	12
Number of generations	50
Crossover rate	70%
Mutation rate	10% per gene
Elite size	1 individual

aware of the status of the evolutionary process and can easily manage the archive and change the filters.

The Customisation Area is divided into two sub-areas. In the first one, the user can manage the evolution, by increasing the number of generations or modifying the filters parameters (such as colour, material or target price). In the second area, the user can visualise the outcomes of the evolution and, if the candidate solutions correspond to the user's preferences, increase their fitness and/or save them to the archive.

In the archive area, the user can visualise the saved individuals, remove them, or reintroduce them in the current population.

4 Experimental Results

We will use EvoFashion to conduct two different experiments. First, we will address the ability of the automatic fitness components to promote evolution towards a target colour, material, price or a combination of multiple parameters. Next, the automatic fitness components will be combined with user interaction, through IEC.

4.1 Experimental Setup

Table 1 details the framework parameters used in the experiments conducted in the following sections. All the values, except the number of generations, are kept fixed throughout the experiments. When performing tests using interactive fitness the number of maximum generations is not fixed, as the stopping criteria depends on the user preferences. In order to prevent user exhaustion, and to allow the majority of the population to be displayed in the screen at the same time, we keep the population size small, i.e., 12 individuals. For that reason, we use high mutation and crossover rates so that the individuals in each generation have noticeable differences. This enables a faster convergence towards feasible solutions, i.e., the emergence of individuals that are considered of high quality by the user. Parent selection is performed using roulette wheel.

The shoe models that are going to be optimised allow the customisation of the following parts: front, side, sole, lining, zipper tape, heel, straps and strap tips, and hardware. For each shoe part it is possible to define two parameters: colour and material.

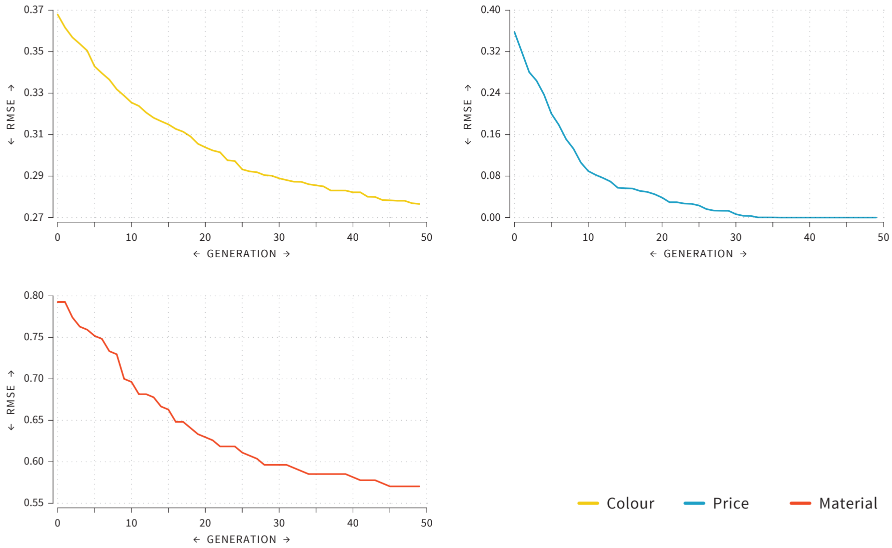

Fig. 3: Evolution of the RMSE of the best individuals across 50 generations for each of the automatic fitness components: colour, price and material. The results are averages of 30 independent runs.

4.2 Automatic Evolution

The first experiments focus only on the use of the automatic fitness components. Therefore, we have analysed the ability of each fitness function component to promote the emergence of candidate solutions that depict the desired characteristics for a specific target value. The targets for colour, price and material were set to $\#dae415$ ( ), > 7000 and crocodile, respectively.

Figure 3 shows the evolution of the RMSE of the best individuals across 50 generations for each of the automatic fitness components. The results are averages of 30 independent runs. A perusal analysis of the results confirms that in all scenarios the system is capable of promoting the convergence towards regions of the search space that have the desired characteristics, i.e., yellow shoe models, with prices above 7000 or with the majority of the shoe parts made of crocodile, respectively for the evolution guided by colour, price and material.

A more detailed analysis of the charts reveal that when evolving towards a target colour (in this case, $\#dae415$) the convergence speed is lower than in the remaining experiments. Moreover, while in the experiments concerning evolution based on price and material evolution attains results close to the best possible solution, with colour this seems not to be the case. In the evolution towards crocodile materials, the best possible solution has a RMSE of $5/9$ because only 4 shoe parts out of the 9 can be made of crocodile. However, it is expected that RMSE values attained with colour are higher than those reached by the remaining metrics. The reason for that is related to the values that each of

Fig. 4: Example of the initial population of one evolutionary run focusing the evolution towards a target colour: $\#dae415$ ( ) . Individuals are ordered from left to right and from top to bottom according to their fitness values.

Fig. 5: Example of the last population for the same run of Figure 4. Individuals are ordered from left to right and from top to bottom according to their fitness values.

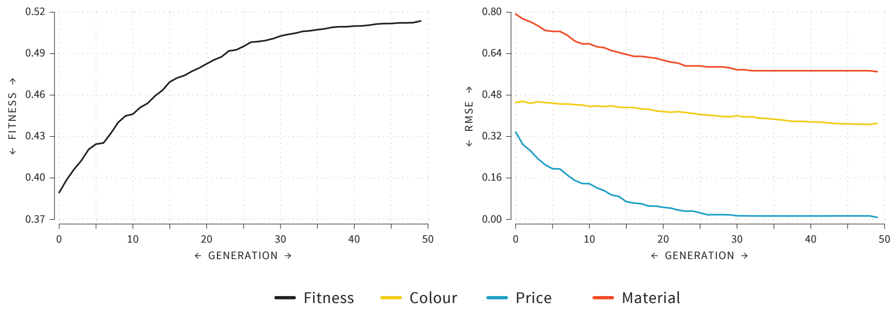

Fig. 6: On the left, the evolution of the fitness of the best individuals across 50 generations with a combination of the the automatic fitness components: colour ($\#dae415$), price (> 7000) and material equal crocodile. On the right, the evolution of each of the automatic fitness components. The results are averages of 30 independent runs.

the variables can assume and how the error is computed: in the colour guided evolution, the error is computed as a distance to a target colour, which is specified by the user. But, the specified colour may not be allowed to all shoe parts, or may have slightly different RGB colour values depending on the used material. As such, the goal of this fitness component is not to attain a value of 0, but rather an approximation of the input colour.

Figures 4 and 5 show, respectively, the initial and last populations of a randomly chosen run aiming at converging the colours of the different shoe parts towards $\#dae415$. Candidate solutions are ordered from the best to the worst individuals in terms of RMSE, i.e., the first individuals are those closer to having all shoe parts with the same colour, and equal to the target one. Looking at the images of the shoe models it is possible to observe that in the initial generation the shoes depict a great variance, having different colours and materials. Additionally, only 2 out of 12 shoes depict parts with yellow colours. On the contrary, in the last generation, 9 out of the 12 shoes have at least one part of the shoe with a colour similar to the target one, being that the first 4 presented models are almost entirely yellow.

From the above conducted experiments we conclude that when we only aim at tackling one of the automatic fitness components the framework is able to find adequate solutions. In the upcoming paragraphs, we will analyse the ability of the framework to obtain shoe models subjected to several restrictions. More precisely, we will promote the evolution towards models that are yellow, have a price above 7000 and are made of crocodile.

Figure 6 depicts the results obtained when using all automatic fitness components. From the left figure it is possible to acknowledge that across generations evolution is promoted. To explain how better solutions are being attained, we need to look at the image on the right, which shows the decomposition of the automatic fitness function into the different components. The fitness value is to

Fig. 7: On the left, the best candidate solutions of the 6 evolutionary runs that obtained the highest fitness results. On the right, the opposite, i.e., the best candidate solutions of the 6 evolutionary runs that obtained the lowest fitness results. Results are from the 30 evolutionary runs guided using the combined fitness function.

be maximised and the errors of the different fitness components are to be minimised. Thus, in the left chart the higher the values the better, and on the right hand side, the lower the better.

An analysis of the right panel shows that the automatic fitness components that have a greater contribution to the fitness evolution are price and material, which attain results closer to the best possible: 0 in the price error, and $^5/_9$ in the material error (only 4 of the 9 shoe model parts can be made of crocodile). The colour evolution is much slower, ending with an error of approximately 0.4, which is not far from the initial 0.45. However, looking at the candidate solutions of the evolutionary runs that reach the best and worst results (see Figure 7) it is possible to se that even the solutions from the runs that obtain the worst results have at least one of the shoe parts coloured in yellow. As such, we conclude that the colour objective is also being accomplished. Because of the used material, the RGB components in the crocodile material are different than the ones of the selected colour by the user. Therefore, it is impossible to get shoe parts painted with the selected colour, due to conflicting evolutionary objectives. As a curiosity, the shoe with the highest fitness in Figure 7 (top left corner) was capable of achieving all evolutionary objectives: it has a price of 7450, all shoe parts use crocodile, and the shoe model is entirely yellow.

4.3 Partially Interactive Evolution

Another aspect that needs to be analysed is the ability of the user to interact with the framework. To accomplish this task, we developed an easy to use web-interface (Section 3.5). The user sees the results of the evolutionary in a grid layout, and may provide feedback regarding the quality of each of the generated solutions. For that, the user presses the improve fitness button as many times as he/she likes, i.e., each time the button is pressed, the fitness value is increase by an a priori defined value. Additionally, if the user finds a shoe design aesthetically pleasing, he/she can store it on the archive. Notwithstanding, the user can combine his/her own preferences with the aforementioned automatic fitness components.

To demonstrate that the framework is able to comply with the user preferences, following the steps described above, please refer to the video posted in `http://goo.gl/hoOIPj`.

5 Conclusions and Future Work

Motivated by the increase in online shopping and by the desire of customers to personalise the products they buy we propose a partially interactive evolutionary computation framework for the customisation of fashion products. The framework is based on a web-interface which allows the user to visualise the evolutionary process and select the products he/she finds aesthetically pleasing. In addition to user attribute fitness, the framework also lets the user to specify filter values, which are used to automatically guide evolution. There is an archive where best solutions can be saved, and later re-introduced in the evolutionary process. The archive has two main goals; first, it allows the customer to keep track of the most aesthetic solutions; secondly, it avoids the loss of population diversity.

Results show the ability of the framework to successfully customise fashion products. In concrete, we have targeted the personalisation of shoe models. In the first experiments, focus was given to the evolution only considering the automatic fitness components, i.e., the evolution of candidate solutions towards a specific colour, price range or material. Then, tests regarding the evolution of shoe models that satisfy more than one of the previous conditions were conducted. In both scenarios it was possible to obtain the expected results, and the algorithm successfully converged to regions of the search space where the shoe model had the desired characteristics.

After assessing the ability of the automatic fitness function components, we performed experiments where the user was asked to evaluate the generated shoe models. Results show that, by using an interactive fitness function, the user is able to promote solutions with the defined parameters and that, when his/her needs are satisfied by a given solution, he/she can increase its fitness, promoting the shoe model characteristics through the next generations.

Next steps to expand on this work will focus on experimenting with the customisation of different fashion products, using different rendering engines. Additionally, we will look into methodologies for automatically updating the archive. An example of such approaches is the work of Vinhas et. al [19], where the an archive of candidate solutions is updated according to the new solutions' novelty degree, i.e., for individuals to be added to the archive they have to be different from those that are already in the archive, and their quality must be superior to a given minimum threshold. Approaches for automatically initialising the first population from the archive, as well as other automatic fitness function metrics will also be investigated.

Acknowledgments

All shoe images are copyright of MYSWEAR[2], and are used as fair use for academic purposes only. We gratefully acknowledge the support of NVIDIA Corporation for the donation of a Titan X GPU. We would also like to thank Tiago Martins for all the patience making the charts herein presented.

References

1. Pine, B.J.: Markets of one: Creating customer-unique value through mass customization. Harvard Business Press (2000)
2. Eiben, A.E., Smith, J.E.: Introduction to evolutionary computing. Volume 53. Springer (2003)
3. Caldas, L.G., Norford, L.K.: A genetic algorithm tool for design optimization. In: ACADIA'99. Proc. of 1999 Conf. of the Association for Computer-Aided Design in Architecture, Salt Lake City, UT. (1999)
4. Besserud, K., Cotten, J.: Architectural genomics. (2008)
5. Michalek, J., Choudhary, R., Papalambros, P.: Architectural layout design optimization. Engineering optimization **34**(5) (2002) 461–484
6. Juan, Y.K., Shih, S.G., Perng, Y.H.: Decision support for housing customization: A hybrid approach using case-based reasoning and genetic algorithm. Expert Systems with Applications **31**(1) (2006) 83–93
7. Hornby, G.S., Globus, A., Linden, D.S., Lohn, J.D.: Automated antenna design with evolutionary algorithms. In: AIAA Space. (2006) 19–21
8. C. Gonzalez-Morcillo, V. J. Martin, D.V.F.J.J.C.S., Albusac, J.: Gaudii: An automated graphic design expert system. In: IAAI. (2010)
9. O'Donovan, P., Agarwala, A., Hertzmann, A.: Learning layouts for single-pagegraphic designs. IEEE transactions on visualization and computer graphics **20**(8) (2014) 1200–1213
10. Geigel, J., Loui, A.: Using genetic algorithms for album page layouts. IEEE multimedia **10**(4) (2003) 16–27

[2] `https://www.my-swear.com/`

11. Takagi, H.: Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation. Proceedings of the IEEE **89**(9) (2001) 1275–1296
12. Kim, H.S., Cho, S.B.: Application of interactive genetic algorithm to fashion design. Engineering applications of artificial intelligence **13**(6) (2000) 635–644
13. Khajeh, M., Payvandy, P., Derakhshan, S.J.: Fashion set design with an emphasis on fabric composition using the interactive genetic algorithm. Fashion and Textiles **3**(1) (2016) 1–16
14. Crispin, A., Clay, P., Taylor, G., Bayes, T., Reedman, D.: Genetic algorithm coding methods for leather nesting. Applied Intelligence **23**(1) (2005) 9–20
15. Shiyou, Y., Guangzheng, N., Yan, L., Renyuan, T.: Shape optimization of pole shoes in harmonic exciting synchronous generators using a stochastic algorithm. IEEE Transactions on Magnetics **33**(2) (1997) 1920–1923
16. Shimoyama, K., Seo, K., Nishiwaki, T., Jeong, S., Obayashi, S.: Design optimization of a sport shoe sole structure by evolutionary computation and finite element method analysis. Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology **225**(4) (2011) 179–188
17. Dasan, A.: The evolutionary design of generative shoes. `http://arandasan.co.uk/Evolve` (2012) Accessed: 2016-11-06.
18. Machado, P., Romero, J., Cardoso, A., Santos, A.: Partially interactive evolutionary artists. New Generation Computing **23**(2) (2005) 143–155
19. Vinhas, A., Assunção, F., Correia, J., Ekárt, A., Machado, P.: Fitness and novelty in evolutionary art. In: International Conference on Evolutionary and Biologically Inspired Music and Art, Springer (2016) 225–240