

Exploring Automatic Fitness Evaluation for Evolutionary Typesetting

Sérgio M. Rebelo

University of Coimbra, CISUC, DEI
Coimbra, Portugal
srebelo@dei.uc.pt

João Bicker

University of Coimbra, CISUC, DEI
Coimbra, Portugal
bicker@dei.uc.pt

Tiago Martins

University of Coimbra, CISUC, DEI
Coimbra, Portugal
tiagofm@dei.uc.pt

Penousal Machado

University of Coimbra, CISUC, DEI
Coimbra, Portugal
machado@dei.uc.pt

ABSTRACT

The recent popularity of creative coding tools and Computational Creativity approaches are promoting a paradigm shift in the creation, development and production of Graphic Design artefacts. In this work, we present an evolutionary system for the automatic typesetting of typographic posters. This system is inspired by the letterpress typesetting process of the print houses in the earlier 19th century and employs lexicon-based approaches to recognise the semantic meaning of the posters' content. During the evolutionary process, poster designs are automatically created and evaluated according to three objectives: legibility, aesthetics, and semantics. The system allows the users to express their preferences by specifying the intended visual features for the output designs, selecting the preferable fitness assignment strategy, and controlling different aspects of the evaluation strategy. We implemented three automatic strategies to evaluate the fitness of the posters: a multi-criteria hardwired fitness function, a multi-objective optimisation approach, and a hybrid strategy that combines features from the previous two strategies. The experimental results demonstrate the ability of the presented system to generate typographic posters, from scratch, and show the impact of the different evaluation strategies on the evolved poster designs. Overall, this research reveals how Evolutionary Computation approaches can be employed to develop novel co-creative typesetting tools and enable the automatic creation of customised typographic designs.

CCS CONCEPTS

• **Computing methodologies** → Information extraction; • **Theory of computation** → Evolutionary algorithms; • **Applied computing** → Publishing; Arts and humanities.

KEYWORDS

Automated Layout, Automatic Fitness Assignment, Evolutionary Computation, Generative Design, Graphic Design, Hardwired Fitness Function, Multi-Objective Optimisation

ACM Reference Format:

Sérgio M. Rebelo, Tiago Martins, João Bicker, and Penousal Machado. 2021. Exploring Automatic Fitness Evaluation for Evolutionary Typesetting. In *Creativity and Cognition (C&C '21)*, June 22–23, 2021, Virtual Event, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3450741.3465247>

1 INTRODUCTION

Posters are one of the most representative, persistent, and resilient artefacts of Graphic Design (GD) [5]. They were already present in the ancient societies, and over the ages, they adjusted to the social and technological contexts that emerged [8, 22, 24, 37]. Recently, creative coding tools and Computational Creativity (CC) approaches are becoming popular in GD and Visual Communication scenarios. The use of these approaches promotes the development of innovative tools, methods, interactions, and production processes to surpass the nowadays demand for dynamic, customised and flexible designs [2, 49, 54]. In the poster design scenario, computational approaches are recurrently employed to create and adapt posters to modern digital media, to integrate external data on the poster designs, and to develop customised and interactive experiences for the viewers.

In this paper, we present an Evolutionary Computation (EC) system to generate letterpress-inspired posters in an automatic fashion, following our research on the evolution of poster designs [50–52]. Lexicon-based approaches are employed to automatically recognise the emotions and colours related to the posters' content and a Genetic Algorithm (GA) that autonomously generates and evolves the outputs. Each generated poster is evaluated considering three objectives: legibility, *i.e.* if the content on the poster is fully readable; aesthetics, *i.e.* if the poster is visually appealing; and semantics, *i.e.* if the visual characteristics of the poster convey the semantic meaning of its content. The system assigns fitness to the evolved individuals based on three automatic strategies: a multi-criteria hardwired fitness function, a multi-objective optimisation approach, and a hybrid strategy that combines features from the multi-criteria hardwired strategy and the multi-objective optimisation strategy. Also, the system allows the users to express

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
C&C '21, June 22–23, 2021, Virtual Event, Italy
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8376-9/21/06.
<https://doi.org/10.1145/3450741.3465247>

their preferences by specifying the intended visual characteristics on the outputs, selecting the preferable fitness assignment strategy, and interactively controlling several aspects related to the evolutionary process. One can access the system, more information and materials related to this work at cdv.dei.uc.pt/evoposter/.

We conducted a series of experiments to analyse the behaviour and the results of the system when each one of the proposed fitness assignment strategies is employed. Overall, the experimental results show the potential of EC to address the contemporary needs of an effective, autonomous, scalable, and flexible approach for the generation of GD artefacts. Also, the proposed system demonstrates the potential of computational design techniques and co-creativity tools to automate several processes in GD, creating novel, customised and flexible manners to communicate with people as well as to stimulate graphic designers' creativity.

The main contributions presented herein include (i) a generative approach for the automatic typesetting of typographic posters, concerning the textual purpose and length of its content and the user preferences; (ii) an evolutionary framework for the evolution of typographic poster designs; (iii) a letterpress-inspired posters evaluation measure based on three evaluation objectives (legibility, aesthetics, and semantics); (iv) three strategies to assign the fitness of typographic poster designs (multi-criteria hardwired fitness function, multi-objective optimisation approach, and hybrid approach) and a comparative study between them; and (v) an exploration of how evolutionary and co-creative approaches may inform the contemporary GD practices.

The remainder of this paper is organised as follows. Section 2 presents the background related to the evolution and evaluation of GD artefacts. Section 3 comprehensively describes the proposed evolutionary approach. Section 4 reports the experiments conducted to analyse the results obtained by each strategy. Finally, Section 5 draws the conclusions and points the directions for future work.

2 BACKGROUND

Since the second half of the 20th century, the use of computational approaches to create GD artefacts has been increasingly explored while the interest in creative coding tools grows [6, 45, 49, 54]. In the poster and document design scenario, we highlight the experimental and pioneering works by Muriel Cooper, and her students at Visible Language Workshop in MIT (e.g. [8]), and John Maeda (e.g. [34]), who created layouts using tailor-made software and generative approaches. Over time, other designers also employed these technologies to create visuals for their designs (e.g. [43] or [21]), to automate the generation of poster designs (e.g. [30] or [43]), or to develop computer-aided design software for layout design (e.g. [11] or [12]).

Interactive Evolutionary Computation (IEC) approaches allow the users to guide the evolution by indicating their favourite evolved individuals. This approach has been a popular strategy for the evolution of document and poster designs. Gatarski [17] developed a system to automatically evolve digital banners using the user's click-through as a fitness metric. Quiroz et al. [48] developed an approach to generate brochure documents where users guide the system by assessing only a small subset of the results. Morcillo et al. [42] created a system, the *GAUDII*, that generates single-page designs

using an interactive approach where the users define preferable design properties. Önduygu [44] developed the system *Gráphagos* that generates design compositions through the interactive evolution of certain features of visual elements. Kitamura and Kanoh [25] interactively evolved poster designs evaluating their visual characteristics. Denis Klein developed the tool *Crossing, Mixing, Mutating* Klein [26] to generate variations in a template using genetic operators. Later, Denis Klein and Lisa Reimann released an updated version of this tool as an Adobe InDesign plug-in named *Evolving Layout* [27].

In Evolutionary Art and Design, automatic fitness assignment strategies try to model and/or simulate the human evaluation [28, 33]. Over the years, different automatic approaches to assign fitness have been explored, including fitness functions based on Machine Learning techniques (e.g. [58], [9], or [63]), hardwired fitness functions (e.g. [31], [60], or [14]), multi-objective optimisation approaches (e.g. [20], [4], or [65]), etc. In the context of posters and documents evolution, as far as we know, some experiments have been developed, especially adopting hardwired fitness function approaches. Fuchs [16] evolved web page layouts based on the goal of minimising blank space. Similarly, Goldenberg [19] employed EC to automatically generate page layouts, minimising the waste of space on the page. Purvis et al. [47] automatically evolved document layouts using a multi-objective optimisation approach, taking into consideration a set of layout constraints and desired design aesthetics. Geigel and Loui [18] evolved layouts for photo books by evaluating different aesthetics criteria. Similarly, Sandhaus et al. [55] evolved photo layouts based on rules of layout design.

From our analysis of the related work, we concluded that EC approaches present a great potential for the automatization and support of the creation of poster and document designs, since they allow the replication of operations often performed by designers during their creation process. Nevertheless, their current use in this context is scarce and the existing related work presents some limitations and drawbacks. IEC strategies have been capable of evolving a wide variety of designs and allow to include the subjectivity of human evaluation on EC's systems. However, they provoke users' fatigue and, consequently, loss of interest and inconsistent evaluations [28, 62]. Alternatively, automatic evaluation strategies can achieve notable results effortlessly; nevertheless, it is more difficult for the users to express their preferences to the systems [28, 33]. This way, we believe that automatic approaches may be combined with some users' orientation, as explored in other Evolutionary Art works (e.g. [64], [32] or [35]).

3 APPROACH

The presented system generates typographic posters that organise a given textual content structured in text boxes that are placed on a canvas. The workflow of this system is inspired by the production processes of the letterpress posters in the early 19th century, developed to respond to the necessity of mass communication [22, 37]. At the time, designers composed the visual elements to carry out a matrix, often in collaboration with the clients. The visual elements were selected from an extensive set of typography resources (e.g. typefaces, fillets, ornaments, engravings, etc.) and the philosophy of design at the time was to use the maximum of them [22, 37].

This way, designers composed posters trying to fill all the space available, using condensed typefaces to longer sentences, wider typefaces to shorter sentences and emphasised the most important parts of the content by assigning them more space on the layout.

The system employs lexicon-based approaches to detect the most important parts of the text content and uses a GA to evolve a population of posters. The evaluation of posters is based on three objectives: legibility, aesthetics, and semantics. Also, we implemented three fitness assignment strategies to assess the merit of each individual: a multi-criteria hardwired fitness function strategy, a multi-objective optimisation strategy, and a hybrid strategy that combines a hardwired fitness function part with one multi-objective optimisation part.

Although the system automatically evaluates the evolved posters, the users may express their preferences by specifying several settings and intended visual features using a dedicated interface. The definition of these preferences is aligned with the specifics of typical EC systems, as well as the design and visual characteristics of letterpress posters (see [7, 37]). Thus, the user can define (i) the fitness assignment strategy employed and its preferences (*i.e.* weights or the objective(s) that should be optimised), (ii) the GA settings (*i.e.* population size, number of generations, probability of each variation operator and elite size), and (iii) visual features of the evolved posters (*i.e.* size, grid, letter case, available colours, available typeface(s) and weights). These preferences may be fine-tuned over the generations. One can incorporate new typefaces into the system by updating the typefaces settings sheet, which stores information about the typefaces used, *i.e.* available weights, pairing scores with other typefaces and paths to the typefaces files. The available colour tones can also be modified in the colours settings sheet. The users' preferences are considered by the system during its evaluation and evolution stages, overcoming some of the limitations posed by fully automatic fitness assignment strategies.

The system behaves as follows. The generative process begins with the semantic analysis of the input text content to recognise the emotions and colours associated with it. The resulting information influences the objectives evaluation. After analysing the input text, the evolutionary process begins. We start by initialising the population with random individuals, *i.e.* poster designs. Each poster of the population is then evaluated in each objective and its fitness assigned according to the fitness strategy selected. Afterwards, tournament selection is used to pick individuals, based on their fitness, from the population to form the next generation. To achieve this, crossover and mutation operators are applied on the selected individuals to generate the new offspring. The next generation will be created using an elitist process, joining these new individuals with a preset number of the best individuals of the current generation. This evolutionary process is repeated until a stop criterion is fulfilled, *e.g.* a preset number of generations is achieved or manual interruption by the user.

3.1 Semantic Analysis

The semantic analysis method examines the posters' text to recognise its most important parts (*i.e.* the parts that should be emphasised in the layout). This method also recognises the colours most

related to the content. The analysis is performed using lexicon-based approaches since there is needed an analysis at the word level. The process begins with the subdivision of text into lines. This division is based on the results of a Sentence Boundary Detection [53] algorithm, an optimal range of characters by line, and a random factor. Next, the resulting lines are prepared for the lexicon analyses by (i) removing the contracted word forms, (ii) replacing abbreviations and slang expressions with formal forms, (iii) replacing the words by their lemmas, (iv) replacing negations expressions by antonyms, (v) removing the stop words, and (vi) tokenising the text. Subsequently, it performs lexicon-based analyses of the text. This way, the resulting tokens are searched in a word-emotion association lexicon, developed by Mohammad and Turney [41], and in a word-colour association lexicon developed by Mohammad [39, 40]. These lexicons were created by manual annotation and include the most frequent terms in English. The word-emotion lexicon relates the words to eight basic and prototypical emotions, *i.e.* anger, anticipation, disgust, fear, joy, sadness, surprise, and trust (see [46]). The word-colour lexicon defines the intensity of the relationship between words as 11 colours, *i.e.* black, blue, brown, green, grey, orange, purple, pink, red, white, and yellow.

The results of the word-emotion analysis are grouped by line. Thus, how more occurrences are recognised in a line, more important it considered this line on the text. Also, an optimal semantic layout is defined based on the normalised emotional charge of each line. The results of the word-colour analysis are presented globally (*i.e.* the accumulative intensity of occurrences found in the text by colour) and stored in a sorted list. We developed this method using the natural language facility library *fin* [10] and the Merriam-Webster Dictionary API [38].

3.2 Representation

Each poster design comprises a set of arranged text boxes typeset in a certain colour. Thus, each poster is encoded as a sequence of arrays of numbers (*i.e.* the genotype). The first array in the sequence is a number and represents the typography colour, based on the range of colours available (*i.e.* the colour configuration gene). The following arrays are four-dimension arrays and encode the text boxes typeface, font's weight, height, and font size in percentages of the height, respectively (*i.e.* the text boxes genes). Since the content of the posters may have different lengths, the number of text boxes and, so, the size of genotype may vary. We subdivided the posters' canvas into a one-column grid with multiple rows that constraint the text boxes position and sizes. Perceptible poster designs (*i.e.* phenotypes) are generated through the rendering of the text boxes, according to the information encoded in genotype, using *P5.js* [36].

3.3 Variation Operators

The evolutionary process begins with the random initialisation of a population of posters. The colour configuration gene is randomly assigned with one colour from the range of options available. The text boxes genes are defined based on the number of lines resulting from the semantic analysis (see subsection 3.1). Thus, for each line, it defines one text box. The features of the text boxes are assigned by a bespoke method. The typeface is randomly selected from the range

of options available. The font's weight is randomly selected from the range of options available for the selected typeface. Although selected at random, the text boxes' height is defined by making sure that all the text boxes on a poster fill all the poster's height. This way, it creates a random sequence of numbers with the same length that the number of lines and its sum as the same as the number of rows of the grid in the initialisation. After, it randomly assigns one position of this list to each text box. The font size is always defined at 100% of the height.

Poster designs are evolved iteratively through the employment of crossover and mutation. Both operators preserve the validity of the generated individuals. The crossover operator employs a uniform crossover method (see [61]), which randomly selects whose parent will give its gene to the children. This operator does not crossover the values related to the height ensuring that all generated children are valid individuals, *i.e.* their text boxes fulfil all the posters' height.

The mutation operator consists of two methods: *Independent*, and *Swap*. The *Independent* mutation method randomly selects a gene and, subsequently, randomly selects a parameter in the gene for mutation. Each type of parameter has its own mutation method. If the colour configuration gene or the font's weight parameter is selected, it randomly defines a new value for the parameter, according to the options available (*i.e.* the number of colours or weights). Otherwise, if the text box's height parameter is selected, it randomly selects two genes, having one, at least, the height value bigger than one. After, it decides what will be the gene that will decrease the height and the one that will increase. This selection is performed randomly unless one of the selected genes has the value 1 (in this case, the gene with the lowest value will increase its height and the other will decrease). Finally, when the font size parameter is selected, it decreases or increases this value by 1%. The direction of this mutation is randomly calculated unless the value is already in its minimum or maximum value. The *Swap* method randomly selects two text boxes, and swaps the value of their genes.

3.4 Evaluation Objectives

The posters are evaluated based on three objectives: legibility, aesthetics, and semantics. Each objective measures a different communication feature of a poster. Their definition attempts to evaluate a typographic poster considering their efficiency of communication, along with the quality of their visual and conceptual language (see [15]).

The legibility objective measures how readable is the content on the posters. Thus, it measures if the width of the content, typeset on the text box, is inside of the width of the poster and if the negative space (*i.e.* space in background colour) is minimised as possible. The legibility value of each text box is the difference between the width of the posters and the width of the text box content when typeset. After, this difference is normalised to assign a poor assessment when text width exceeds the poster width and, at the same time, progressively prejudice the text boxes when the negative space surpasses a certain target amount. The overall legibility value is the weighted arithmetic mean of the value of all text boxes. The weight of each text box in the mean is automatically calculated based on its height.

The aesthetics objective measures how much of the poster design is visually appealing according to a set of aesthetic measures for typographic poster designs. These measures are based on the works of Harrington et al. [23], Bringham [7] and Lupton [29]. Thus, the aesthetics of a poster is evaluated according to its (i) *alignment*, (ii) *regularity*, (iii) *balance*, (iv) *negative space fraction*, (v) *composition security*, and (vi) *typography pairing*. The overall value of the aesthetics objective is the arithmetic mean of these attributes. The *alignment* attribute measures how justified is the text, by calculating the distance between the horizontal position of the left edges in neighbouring text boxes. The closer the vertical distance between text boxes, the higher is the alignment score. The overall alignment measure is the arithmetic mean of all distances. The *regularity* attribute measures how regular is the vertical placement of the text boxes, by calculating the distances between the vertical positions of the top edges in neighbouring text boxes. The *balance* attribute measures how much of the poster is centrally balanced. The centre balance of a poster is the difference between the centre of its visual weight and its visual centre and calculated as proposed by Harrington et al. [23]. The *negative space fraction* is the absolute difference between the current percentage of coloured in background colour and a certain target percentage threshold. The *composition security* attribute measures if the text boxes positioned near the edges of the poster are secure and do not appear to fall off. This value is the minimum value of the top and bottom edges of all text boxes. The *typography pairing attribute* is the arithmetic mean of the pairing scores of the typeface used on the poster. If only one typeface is used on the poster, this measure is not considered.

The semantics objective measures how much of the posters' visual characteristics convey the semantic meaning of their content. In the context of this work, the posters should emphasise the most important parts of their content on the layout. We considered that the most important parts are those with more word-emotions occurrences. Also, the typography colour should be related to the content, once the use of colours that match the message can strengthen the communication of the same message [40]. Thus, the semantics value of a poster is the arithmetic mean between the evaluation of the appropriateness of (i) *layout* and (ii) *typography colour*. The *appropriateness of the layout* is computed considering the optimal semantic layout. This way, for each text box, it calculates the distance between its current height and its optimal height. The overall value is the arithmetic mean of all distances. The *appropriateness of a typography colour* is calculated, taking into consideration the word-colour analysis. One poster's typography colour conveys the semantic meaning if it is the colour most related to it. Thus, its value is the normalised distance, in the number of positions on the sorted list, between the employed colour and the optimal colour. Optimal semantic layout and the word-colour analysis results sorted list are defined earlier (see subsection 3.1).

3.5 Fitness Assignment Strategies

We implemented three automatic strategies to fitness the individuals generated by the system: multi-criteria hardwired fitness function approach, multi-objective optimisation approach, and hybrid approach. The definition of the quality of a poster is a contextual and subjective task influence by multiple factors. In this sense,

Table 1: Experimental parameters.

Parameter	Value	Parameter	Value
Generations	250	Poster margin	15px
Population size	30	Poster grid	26 × 1
Crossover probability	0.8	Visual centre vertical offset	1/12
Mutation probability	0.7	Target percentage of negative space	50%
Phenotype size	298px × 420px		

graphic designers are the most responsible for evaluating their creative process, editing (or curating) iteratively the outcomes based on the briefing, the tradition and the clients' expectations [15, 57]. The developed fitness assignment strategies are aligned with these characteristics of the GD process, supplying different manners to edit the system's creative process and distinct ways of combining the three evaluation objectives. Also, the users can define what is the fitness assignment strategy used and swaps it over the generation. This way, they may decide whose strategy is more proper, in each moment, according to their empirical observation of the evolution.

The multi-criteria hardwired fitness assignment strategy considers that quality posters are those that fully satisfy all the objectives. However, each objective can have a different weight on the evaluation. Therefore, this strategy considers that the quality of a poster is measured by the weighted arithmetic mean of the three objectives. Users may define the weight of each objective on the mean and fine-tune it over the generation.

The multi-objective optimisation approach strategy implements an Elitist Non-dominated Sorting Genetic Algorithm (NGGA-II) [13] to fitness the posters. This strategy considers that quality posters are those that promote more innovative and distinctive relationships between the three objectives (either optimisation, minimisation or balance). The fitness of each individual is defined by the rank of the front, where it is placed in the non-domination sort of the population. It organises the population in fronts as proposed by Deb et al. [13]. The crowding distance is also considered in the selection of the individuals during the tournament. Since this strategy is very destructive for the population, we implemented a bespoke elite mechanism that stores in the elite pool all the individuals on the non-dominated front (guaranteeing that, at least, 50% of the individuals on population are from the new offspring). The users may also define the objectives that they want to optimise. In these scenarios, it stores only the individuals that fulfil these preferences.

The hybrid approach combines a hardwired fitness function part with a multi-objective optimisation approach part. This strategy considers that quality posters should always be readable and present innovative and distinctive relationships between the semantics and aesthetics objectives (either balance or optimisation/minimisation). Thus, it fitness each poster by calculating the weighted arithmetic mean of legibility measure (the hardwired part) with the relation between the semantics and aesthetics objectives (the multi-objective part). The relation between the aesthetics and the semantics objectives is computed similarly to the multi-objective optimisation approach strategy; however, it only uses two objectives. Crowding distance is also used in case of a tie in the tournament. Although a

minimum size has defined, the elite pool size may vary in the cases that the best-ranked individuals share the same legibility evaluation and be placed in the non-dominated front concerning aesthetics and semantics objectives. As in the multi-criteria hardwired strategy, users may define, and refine, the weights of each part of the mean.

4 EXPERIMENTATION

We conduct three experiments on the present system to study how evolution is affected by the proposed fitness assignment strategies. Although the system evolution can be affected by contextual, textual, and user-related factors, here, we are only focused on understanding how each evaluation strategy influences the evolution. This way, we conducted these experiments by evolving posters from the same contents and under the same settings, albeit changing the fitness assignment strategy. In these experiments, any preference is not changed during evolution and all the colours available are used. The posters were generated using the type family *Bureau Grot*, by David Berlow (1989), loaded using Adobe Typekit service [1]. Concerning the posters' content, we used three text quotes with different lengths and purposes: (i) *"How amazing! How many wonderful creatures there are here! Mankind is so beautiful! Oh, what a wonderful new world, that has such people in it!"* by Shakespeare [56, Scheme 1, Act 5]; (ii) *"There is no document of civilization which is not at the same time a document of barbarism"* by Benjamin [3, p. 256]; and (iii) *"the work of memory [...] collapses time"* by Sontag [59, p. 115]. The experimental setup was defined by empirical exploration and summarised in table 1. The project's repository, demonstration videos, more examples and the settings sheets employed in these experiments are accessible in cdv.dei.uc.pt/evoposter/.

4.1 Multi-criteria Hardwired Fitness Function Strategy

The multi-criteria hardwired fitness function strategy evaluates each poster by the weighted arithmetic mean of the three objectives. We empirically defined the weight of each objective as 90% to legibility, 5% to aesthetics, and 5% to semantics. The elite size was defined at 1. We established a high legibility weight to force the system to generate readable poster designs, whatever the length and semantic meaning of the text. Figure 1 displays some outputs generated by the system, employing this evaluation strategy.

This strategy allows the generation of readable and finished posters from scratch. Visually observing the outputs, one can note that the visual diversity between posters with the same content is often reduced, both at chromatic and layout levels. We observed that the results' diversity is directly related to data from semantic analysis. Thus, texts with a weak emotional charge or a uniform

distribution of emotional charge generate populations of posters with more diversity and *vice versa*. Regarding the typography colour, there is similar behaviour. Moreover, we noticed that the system generates faster legible outputs when the text is longer since the number of possible layouts decreases. Due to the same reason, longer posters' contents generate more similar populations.

Figure 2 demonstrates the progression of individuals' fitness in the population and the corresponding objectives' evaluation. These charts display that high fitness values are reached in a few generations and, in the initial stages, the evolution is faster than in the later stages. Although the fitness does not reach the maximum theoretical value, the best individuals in the population achieve a high fitness value and high evaluation in legibility and semantics (both above 0.9 out of 1). We observed that high aesthetics evaluations are not too frequently achieved on the results, because the evaluation of this objective is more dependent on the features of the content (e.g. size, lines, number of words per line, etc.). Also, this objective tends to be constrained by legibility, since the readable layouts often cannot satisfy all the aesthetics measures. Nevertheless, it is visible an increase in the aesthetic evaluations over the generations.

4.2 Multi-objective Optimisation Strategy

The multi-objective optimisation strategy assesses each poster based on their non-domination on population. Therefore, this strategy promotes the novelty search, looking for unique designs. Figure 3 present some typical posters generated by the system, employing this evaluation strategy.

This strategy cannot guarantee the generation of fully readable and finished posters (see the bottom line of Figure 3). This occurs because, in some individuals, it minimises the legibility over the other objectives (i.e. it optimises the aesthetics and/or the semantics). Due to the same reason, legible posters sometimes are discarded from the population in favour of unreadable ones. We noted that the probability of generating legible posters is lower when text is not too small or not too long (c. 70–130 characters). This happens because this kind of text does not impose significant layout restrictions (such happens when the text is longer) and not enable flexibility in layout variations (such happens when the text is shorter). Visually observing the results, it is possible to see that longer texts generate more centred compositions than in the other strategies (this type of compositions has a higher aesthetic evaluation). When the user requests the optimisation of the legibility, the system appears to often produce readable outputs (see Figure 4).

Figure 5 illustrates the progression of objectives evaluation in the best-ranked individual and the population over the generations. These graphs illustrate that the use of this strategy produces an inconstant and destructive evolution. However, we observed that



Figure 1: Typical best-ranked individuals evolved by a multi-criteria hardwired approach.




Figure 2: Progression of the posters' fitness (left) and objectives' evaluation (right) over the generations. In the left chart, the solid black line presents the fitness of the best-ranked individual, the solid grey line presents the fitness of the worst individual, and the dotted line presents the average fitness of the population. In the right chart, the blue, green and red lines display the evaluation of legibility, aesthetics, and semantics objectives, respectively. The solid lines represent the best individual and the dotted lines the population's average. The visualised data is the average of 60 runs.



Figure 3: Typical best-ranked individuals evolved by a multi-objective optimisation approach.



Figure 4: Typical best-ranked individuals evolved by a multi-objective optimisation approach, optimising the legibility objective. These posters were generated aside from the presented experiments.




Figure 5: Progression of the objectives evaluation of the best-ranked individual (left) and population (right). Legibility, aesthetics, and semantics objectives are represented by blue, red, and green lines, respectively. The visualised data is the average of 60 runs.



Figure 6: Typical best-ranked individuals evolved by a hybrid approach.

the objectives evaluation increase during the evolution (both in the best-ranked individuals and the population). We also notice that the aesthetic evaluation, on average, maintains a low evaluation over the generations. As previously mentioned, this is related to the specific dependencies of this evaluation. Nevertheless, high aesthetic value individuals are frequently visualised in the population when this strategy is employed. Although the system achieves high semantics and legibility evaluations, evolution is unsteady.

4.3 Hybrid Strategy

The hybrid strategy evaluates each poster by the weighted arithmetic mean of two parts: a hardwired part (the legibility objective), and a multi-objective optimisation part (the relation between aesthetics and semantics objectives). The value of the multi-objective part is defined by the ranking of the front where the individual belongs, after the organisation of population based on the dominance (*i.e.* non-dominant posters have maximum evaluation). We defined the minimum elite size as 1. We empirically defined the weight of 90% to hardwired part and 10% to multi-objective part. We determined a high legibility weight to force the system to generate readable designs, whatever the length and semantic meaning of the text. Figure 6 unveils some outputs generated by the system, employing this evaluation strategy.

This strategy to fitness also allows the generation of readable and finished poster designs from scratch. We noted that the visual diversity of individuals is slightly higher, compared with a multi-criteria hardwired fitness function strategy. The posters also are more related to letterpress' aesthetics. The increase of the visual diversity is related to the multi-objective approach part since the best-ranked individuals do not need to fully satisfy the aesthetics and semantics objectives, but only present a non-dominant relationship between these two objectives. We observed that this strategy evolves well-evaluated poster designs faster when the content is longer. However, how much longer is the content, fewer is the visual diversity of outputs.

Figure 7 illustrates the progression of individuals' fitness in the population and the corresponding objectives' evaluation. To fully understand the progression fitness, the two charts should be viewed together, since the relation between aesthetics and semantics is not a quantitative value, *i.e.* it corresponds to the rank of the front that the individual belongs to. Although the fitness value does not reach the maximum theoretical value, we notice that high fitness values are attained in a few generations. Also, we observe that similar, or higher, fitness and legibility evaluations are achieved, when compared to the multi-criteria hardwired strategy; however, the evolution is slower. On average, the aesthetics objective maintains lower evaluation values. As previously mentioned, this is related




Figure 7: Progression of the posters' fitness (left) and objectives' evaluation (right) over the generations. In the left chart, the solid black line presents the fitness of the best-ranked individual(s), the solid grey line presents the fitness of the worse individual, and the dotted line displays the average fitness of the individuals in the population. In the right chart, the blue, green, and red lines display the evaluation of legibility, aesthetics, and semantics objectives, respectively. The solid lines present the evaluation of the best individuals and the dotted lines the average population. The visualised data is the average of 60 runs.

to the specific dependencies of the evaluation of this objective. On the other hand, the semantics evaluation value is lower and more unstable, when compared with the other strategies. We also note that this strategy is not so destructive as the multi-objective optimisation approach. One can observe that in the earlier stages of evolution, the evaluation of aesthetics and semantics achieves similar, or sometimes higher values, than in some later generations. It is also perceptible that in some earlier stages the average evaluation of these objectives in the population is above the evaluation of the fittest individuals in the population. The reason is that legibility acts like a constraint, disregarding individuals that promote relations between aesthetics and semantics objectives that do not favour the posters' legibility. Nevertheless, the values of the evaluation of aesthetics and semantics objectives are always directly related to the characteristics of the content.

5 CONCLUSION

We presented an evolutionary system for the automatic typeset of letterpress-inspired poster designs. We proposed a poster evaluation measure based on three objectives: legibility, semantics and aesthetics. Each objective assesses one different communication characteristic of the generated posters. Also, three automatic strategies were presented to calculate the fitness of the individuals: multi-criteria hardwired fitness function, multi-objective optimisation approach, and hybrid fitness function. We conducted a series of experiments to explore and evaluate the behaviour of the system using each one of the proposed fitness assignment strategies.

The experimental results demonstrated that the multi-criteria hardwired and hybrid strategies allow the autonomous generation of finished and communicative posters. Overall, the two strategies achieved comparable results. The hybrid strategy achieved slightly better results in terms of visual diversity and legibility evaluation. On the other hand, the multi-criteria hardwired strategy achieved slightly better results in terms of semantic evaluation and evolutionary convergence speed. The multi-objective optimisation strategy, in contrast, generated unstable populations and did not ensure the generation of finished and communicative posters, especially when

the content was not too long or too shorter. Nevertheless, this strategy developed individuals with higher aesthetics evaluations and created more visual diverse populations.

The results also revealed a high potential for the combined use of the three fitness assignment strategies during evolution and this way achieve reliable results that meet different goals. For instance, the user may use the hybrid strategy for multi-purpose needs, the multi-objective strategy to increase visual diversity within the population, and the multi-criteria hardwired strategy to fine-tune the best individuals on the population. However, further studies are needed to understand how the swap between strategies should be performed automatically. Currently, users can already perform this change interactively based on their observation of the evolutionary process.

This system suggests how creative coding tools and CC approaches, especially EC, can be employed with success to automatically generate communication artefacts. Also, it reveals the potential of these approaches in GD future practices. Besides the ability of the system to generate finished and communicative posters from scratch, we also observed that the system is a functional co-creativity tool. The system generates posters in an automatic and customised way, considering the users' preferences to circumvent the common limitation of the automatic evolutionary systems and the user fatigue typically caused by IEC. We believe that systems like this are useful tools for stimulating and foster the users' creativity and inspiration (mostly graphic designers), especially in the earlier and most exploratory stages of their design processes. They enable the exploration of multiple conceptual and visual perspectives when users design posters, in an easier, faster and effortless manner.

Future work on this system will focus on (i) further evaluation of the quality of generated outputs to understand if they are aligned with the expectation of users and stakeholders; (ii) further study and testing of the use of this system as a co-creativity tool by users; (iii) implementation of an adaptive evaluation method that automatically suggests and/or swaps between fitness assignment strategies; (iv) creation of an interface that enables the user to interactively define the aesthetics principles related to the evaluation of aesthetics objective; (v) adaptation of the system evaluation to enable the evolution of posters inspired in other aesthetics styles; (vi) implementation of a method to the automatic gathering of the typefaces and the definition of their pairing (using e.g. Adobe Typekit Web-fonts or Google Fonts APIs); and (vii) adaptation of the system to include images and illustrations.

ACKNOWLEDGMENTS

This work is partially supported by national funds through the Foundation for Science and Technology (FCT), Portugal, within the scope of the project UID/CEC/00326/2019. The first author is funded by FCT under the grant SFRH/BD/132728/2017.

REFERENCES

- [1] Adobe. 2020. Adobe Typekit Web Fonts. Retrieved 20 April 2020 from <https://fonts.adobe.com/typekit/>
- [2] Helen Armstrong and Zvezdana Stojmirovic. 2011. *Participate: Designing with User-Generated Content*. Princeton Architectural Press, New York, NY.
- [3] Walter Benjamin. 1968. *Illuminations*. Schocken Books, New York, NY.
- [4] Steve Bergen and Brian J. Ross. 2012. Aesthetic 3D Model Evolution. In *Evolutionary and Biologically Inspired Music, Sound, Art and Design*, Penousal Machado, Juan Romero, and Adrian Carballal (Eds.). Springer, Berlin, Germany, 11–22.
- [5] Andrew Blauvelt. 2011. The Persistence of Posters. In *Graphic Design: Now in Production*, Andrew Blauvelt and Ellen Lupton (Eds.). Walker Art Center, Minneapolis, MN, Chapter 11, 92–111.
- [6] Margaret A Boden and Ernest A Edmonds. 2009. What is generative art? *Digital Creativity* 20, 1-2 (2009), 21–46.
- [7] Robert Bringhurst. 2005. *The Elements of Typographic Style* (3rd ed.). Hartley & Marks, Vancouver, Canada.
- [8] Muriel Cooper. 1989. Computers and Design. *Design Quarterly* 1, 142 (1989), 1–31.
- [9] João Correia, Penousal Machado, Juan Romero, and Adrián Carballal. 2013. Evolving Figurative Images Using Expression-Based Evolutionary Art. In *Proceedings of the fourth International Conference on Computational Creativity (ICCC)*. University of Sydney, Sydney, Australia, 24–31.
- [10] Alex Corvi. 2020. Fin Natural Language Processor. Retrieved 8 February 2020 from <https://finnlp.gitbook.io/fin/>
- [11] Niranjan Damera-Venkata, José Bento, and Eamonn O'Brien-Strain. 2011. Probabilistic Document Model for Automated Document Composition. In *Proceedings of the 11th ACM Symposium on Document Engineering*. ACM, New York, NY, 3–12.
- [12] Frederik De Bleser. 2016. *Generative Design: The NodeBox Perspective*. Ph.D. Dissertation. University of Antwerp, Antwerp, Belgium.
- [13] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197. <https://doi.org/10.1109/4235.996017>
- [14] Eelco den Heijer and A. Eiben. 2013. Using Scalable Vector Graphics to evolve art. *International Journal of Arts and Technology* 9, 1 (01 2013), 59–85. <https://doi.org/10.1504/IJART.2016.075408>
- [15] Jorge Frascara. 1988. Graphic design: Fine Art or Social Science? *Design Issues* 5, 1 (Autumn 1988), 18–29.
- [16] M. Fuchs. 2000. An evolutionary approach to support Web page design. In *Proceedings of the 2000 Congress on Evolutionary Computation. CEC00 (Cat. No.00TH8512)*, Vol. 2. IEEE, New York, NY, 1312–1319 vol.2. <https://doi.org/10.1109/CEC.2000.870803>
- [17] Richard Gatarski. 2002. Breed better banners: Design Automation Through On-line Interaction. *Journal of Interactive Marketing* 16, 1 (November 2002), 2–13.
- [18] Joe Geigel and Alexander C. P. Loui. 2000. Automatic Page Layout using Genetic Algorithms for Electronic Albuming. In *Internet Imaging II*. International Society for Optics and Photonics, SPIE, San Jose, CA, 79–90.
- [19] Eldan Goldenberg. 2002. *Automatic Layout of Variable-content Print Data*. Master's thesis. University of Sussex, Brighton, United Kingdom.
- [20] Gary R Greenfield. 2003. Evolving aesthetic images using multiobjective optimization. In *The 2003 Congress on Evolutionary Computation, 2003. CEC'03*, Vol. 3. IEEE, IEEE, New York, NY, 1903–1909.
- [21] Benedikt Groß and Julia Laub. 2007. Diploma – Generative Systeme Posters. Retrieved 20 April 2020 from <https://benedikt-gross.de/projects/diploma-generative-systeme-posters/>
- [22] Elizabeth E. Guffey. 2014. *Posters: A Global History*. Reaktion Books, London, United Kingdom.
- [23] Steven J. Harrington, J. Fernando Naveda, Rhys Price Jones, Paul Roetling, and Nishant Thakkar. 2004. Aesthetic Measures for Automated Document Layout. In *Proceedings of the 2004 ACM Symposium on Document Engineering*. ACM, New York, NY, 109–111.
- [24] Richard Hollis. 2002. *Graphic Design: A Concise History* (2nd ed.). Thames Hudson, New York, NY.
- [25] Sonomi Kitamura and Hitoshi Kanoh. 2011. Developing Support System for Making Posters with Interactive Evolutionary Computation. In *2011 Fourth International Symposium on Computational Intelligence and Design*. IEEE, New York, NY, 48–51.
- [26] Denis Klein. 2012. Crossing, Mixing, Mutation. Retrieved 11 May 2020 from <http://www.gutenberg-intermedia.de/wissenschaft-gestaltung/denis-klein-crossing-mixing-mutation/>
- [27] LESS. 2016. Evolving Layout. Retrieved 11 May 2020 from <http://www.evolvinglayout.com/>
- [28] Matthew Lewis. 2008. Evolutionary Visual Art and Design. In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, Juan Romero and Penousal Machado (Eds.). Springer-Verlag, Heidelberg, Germany, Chapter 1, 3–37.
- [29] Ellen Lupton. 2010. *Thinking With Type: A Critical Guide for Designers, Writers, Editors, & Students* (2nd ed.). Princeton Architectural Press, New York, NY.
- [30] LUST. 2008. Graphic Design Museum: Poster Wall for the 21st Century. Retrieved 11 May 2020 from <https://lust.nl/#projects-3041/>
- [31] Penousal Machado and Amílcar Cardoso. 2002. All the truth about NeVAr. *Applied Intelligence* 16, 2 (2002), 101–118.
- [32] Penousal Machado, Tiago Martins, Hugo Amaro, and Pedro H. Abreu. 2016. Beyond Interactive Evolution: Expressing Intentions Through Fitness Functions. *Leonardo* 49, 3 (2016), 251–256. https://doi.org/10.1162/LEON_a_01103

- [33] Penousal Machado, Juan Romero, and Bill Manaris. 2008. Experiments in Computational Aesthetics. In *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, Juan Romero and Penousal Machado (Eds.). Springer, Berlin, Germany, 381–415.
- [34] John Maeda. 2000. *Maeda@Media*. Thames & Hudson, London, United Kingdom.
- [35] Tiago Martins, João Correia, Ernesto Costa, and Penousal Machado. 2019. Evolving Stencils for Typefaces: Combining Machine Learning, User's Preferences and Novelty. *Complexity* 2019 (March 2019), 16 pages.
- [36] Lauren McCarthy, Processing Foundation, and NYU ITP. 2020. P5.js. Retrieved 11 May 2020 from <https://p5js.org/>
- [37] Philip B. Meggs and Alston W. Purvis. 2016. *Meggs' History of Graphic Design* (6th ed.). John Wiley & Sons, New York, NY.
- [38] Merriam-Webster Dictionary. 2021. Merriam-Webster Dictionary API. Retrieved 8 February 2020 from <https://dictionaryapi.com/>
- [39] Saif M. Mohammad. 2011. Colourful Language: Measuring Word-Colour Associations. In *Proceedings of the Second Workshop on Cognitive Modeling and Computational Linguistics*. ACL, Portland, OR, 97–106.
- [40] Saif M. Mohammad. 2011. Even the Abstract Have Colour: Consensus in Word-Colour Associations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*. ACL, Portland, OR, 368–373.
- [41] Saif M. Mohammad and Peter D. Turney. 2012. Crowdsourcing a Word-Emotion Association Lexicon. *Computational Intelligence* 29, 3 (September 2012), 436–465.
- [42] Carlos Gonzalez Morcillo, Victor Jose Martin, David Vallejo Fernandez, Jose Jesus Castro Sanchez, and Javier Alonso Albusac. 2010. Gaudii: An Automated Graphic Design Expert System. In *Proceedings of the Twenty-Second Conference on Innovative Applications of Artificial Intelligence*. AAAI, Atlanta, GA, 1775–1780.
- [43] Boris Müller. 2002. Poetry on the Road 2002–2013. Retrieved 11 May 2020 from <https://www.esono.com/boris/projects/poetry02/>
- [44] Deniz Cem Öndünyü. 2010. *Graphagos: Evolutionary Algorithm as a Model for the Creative Process and as a Tool to Create Graphic Design Products*. Master's thesis. Sabanci University.
- [45] Matt Pearson. 2011. *Generative art: a practical guide using processing*. Manning, Shelter Island, NY.
- [46] Robert Plutchik. 1980. A General Psychoevolutionary Theory of Emotion. In *Theories of Emotion*, Robert Plutchik and Henry Kellerman (Eds.). Academic Press, Cambridge, MA, Chapter 1, 3–33.
- [47] Lisa Purvis, Steven Harrington, Barry O'Sullivan, and Eugene C. Freuder. 2003. Creating Personalized Documents: an Optimization Approach. In *Proceedings of the 2003 ACM Symposium on Document Engineering*. ACM, New York, NY, 68–77.
- [48] Juan C. Quiroz, Amit Banerjee, Sushil J. Louis, and Sergiu M. Dascau. 2009. Document Design with Interactive Evolution. In *New Directions in Intelligent Interactive Multimedia Systems and Services - 2*, Ernesto Damiani, Jechang Jeong, Robert J. Howlett, and Lakhmi C. Jain (Eds.). Springer-Verlag, Heidelberg, Germany, Chapter 29, 309–319.
- [49] C. Reas, C. McWilliams, and LUST. 2010. *Form + Code: In Design, Art, and Architecture*. Princeton Architectural Press, New York, NY.
- [50] S. M. Rebelo, J. Bicker, and P. Machado. 2020. Evolutionary Experiments in Typesetting of Letterpress-Inspired Posters. In *Proceedings of the Eleventh International Conference on Computational Creativity, September 7–11, 2020*, F. A. Cardoso, P. Machado, T. Veale, and J. M. Cunha (Eds.). Association for Computational Creativity, Coimbra, Portugal, 110–114.
- [51] S. M. Rebelo, J. Bicker, and P. Machado. 2021. Evolutionary Typesetting: An Automatic Approach Towards the Generation of Typographic Posters from Tweets. In *Interactivity and Game Creation*, A. Brooks, E. I. Brooks, and D. Jonathan (Eds.). Springer, Cham, Germany, 343–362.
- [52] S. M. Rebelo, C. M. Fonseca, J. Bicker, and P. Machado. 2018. Evolutionary Experiments in the Development of Typographical Posters. In *6th Conference on Computation, Communication, Aesthetics & X (xCoAx 2018)*, A. Rangel, L. Ribas, M. Verdicchio, and M. Carvalhais (Eds.). Universidade do Porto, Madrid, Spain, 65–75.
- [53] Jeffrey C. Reynar and Adwait Ratnaparkhi. 1997. A Maximum Entropy Approach to Identifying Sentence Boundaries. In *Proceedings of the Fifth Conference on Applied Natural Language Processing*. ACL, Washington, DC, 16–19.
- [54] Andrew Richardson. 2016. *Data-Driven Graphic Design: Creative Coding for Visual Communication*. Bloomsbury Publishing, London, United Kingdom.
- [55] Philipp Sandhaus, Mohammad Rabbath, and Susanne Boll. 2011. Employing Aesthetic Principles for Automatic Photo Book Layout. In *MMM 2011: Advances in Multimedia Modeling*, Kuo-Tien Lee, Wen-Hsiang Tsai, Hong-Yuan Mark Liao, Tsuhan Chen, Jun-Wei Hsieh, and Chien-Cheng Tseng (Eds.). Springer-Verlag, Berlin, Germany, 84–95.
- [56] William Shakespeare. 2006. *The tempest*. Yale University Press, London, United Kingdom.
- [57] Adrian Shaughnessy. 2010. *How to be a graphic designer without losing your soul* (2nd ed.). Laurence King Publishing, London, United Kingdom.
- [58] Baluja Shumeet, DDean Pomerleu, and Jochem Todd. 1994. Towards Automated Artificial Evolution for Computer-generated Images. *Connection Science* 6, 2-3 (1994), 325–354. <https://doi.org/10.1080/09540099408915729>
- [59] Susan Sontag. 2002. *Under the Sign of Saturn: Essays*. Picador, New York, NY.
- [60] Nils Svängård and Peter Nordin. 2004. Automated Aesthetic Selection of Evolutionary Art by Distance Based Classification of Genomes and Phenomes Using the Universal Similarity Metric. In *Applications of Evolutionary Computing*, Günther R. Raidl, Stefano Cagnoni, Jürgen Branke, David Wolfe Corne, Rolf Drechsler, Yaochu Jin, Colin G. Johnson, Penousal Machado, Elena Marchiori, Franz Rothlauf, George D. Smith, and Giovanni Squillero (Eds.). Springer, Berlin, Germany, 447–456.
- [61] Gilbert Syswerda. 1989. Uniform Crossover in Genetic Algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers Inc., Fairfax, VA, 2–9.
- [62] H. Takagi. 2001. Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proc. IEEE* 89, 9 (Sep. 2001), 1275–1296.
- [63] Fazle Tanjil and Brian J. Ross. 2019. Deep Learning Concepts for Evolutionary Art. In *Computational Intelligence in Music, Sound, Art and Design*, Anikó Ekárt, Antonios Liapis, and Maria Luz Castro Pena (Eds.). Springer International Publishing, Cham, 1–17.
- [64] J. Ventrella. 1995. Disney meets Darwin—the evolution of funny animated figures. In *Proceedings Computer Animation '95*. IEEE, New York, NY, 35–43. <https://doi.org/10.1109/CA.1995.393550>
- [65] Adriano Vinhas, Filipe Assunção, João Correia, Anikó Ekárt, and Penousal Machado. 2016. Fitness and Novelty in Evolutionary Art. In *Evolutionary and Biologically Inspired Music, Sound, Art and Design – 5th International Conference, EvoMUSART 2016, Porto, Portugal, March 30 – April 1, 2016, Proceedings*. Springer, Cham, Germany, 225–240. https://doi.org/10.1007/978-3-319-31008-4_16