

Evolutionary Typesetting: An Automatic Approach Towards the Generation of Typographic Posters from Tweets

Sérgio M. Rebelo^[0000–0002–7276–8727], João Bicker^[0000–0002–0670–1217], and Penousal Machado^[0000–0002–6308–6484]

University of Coimbra, Centre for Informatics and Systems of the University of
Coimbra, Department of Informatics Engineering, Portugal
`{srebelo,bicker,machado}@dei.uc.pt`

Abstract. The recent developments on Artificial Intelligence are expanding the tools, methods, media, and production processes on Graphic Design. Poster designs are no exception. In this paper, we present a web system that generates letterpress-inspired typographic posters using, as content, tweets posted online. The proposed system employs Natural Language Understanding approaches to recognise the emotions, the sentiments, and the colours associated with the content. Also, the system employs an Evolutionary Computation approach to generate and evolve a population of poster designs. The outputs are evaluated according to their legibility, aesthetics, and semantics, throughout an automatic fitness assignment hybrid scheme that combines a hardwired fitness function part with a multi-objective optimisation approach part. We experimented with the system to perceive its behaviour and its ability to evolve posters from contents with distinct textual purposes and lengths.

Keywords: Evolutionary Computation · Generative Design · Poster Design · Natural Language Understanding · Twitter

1 Introduction

Posters have a ubiquitous presence in our everyday lives. They were already present in ancient societies and, throughout the centuries, have adapted to new social and technological contexts, changing their formats and purposes [4, 34]. The recent developments on Artificial Intelligence (AI) are also reflected in Graphic Design (GD) and Visual Communication, expanding their tools, methods, media, and production processes [2, 50]. In posters design, we observed the exploration of new digital media and computational techniques to create customised, distinct and interactive experiences for their viewers.

In this paper, we presented a system that generates, from scratch, typographic poster designs, similar to letterpress posters, using content posted online on Twitter. Letterpress is a printing technique that emerged on the follow-up of the Industrial Revolution (*c.* early 19th century) driven by developments of new printing technologies and the necessity of mass communication. At the time,

it became popular because it allowed a cheaper, easier and faster printing of commercial posters [17, 34].

Nevertheless, the design of letterpress posters was a process slightly different from the present-day one. At the time, designers composed the visual elements to carry out a matrix, often in collaboration with the client. The visual elements were selected from an extensive set of standard typefaces, fillets, ornaments and engravings and the philosophy of work, at the time, was to use the maximum of them [34]. The design decisions were very pragmatic: longer words and text were composed in more condensed typefaces and shorter words were composed in wider fonts. The most important parts of the content were emphasised through the use of bigger typefaces. Thus, designers needed to hold the elements firmly and strongly, imposing vertical and horizontal tension between the elements.

The present system generates outputs replicating this workflow. Briefly, the system (*i.e.* the designer) composes the content, divided into text boxes to fulfil, as much as possible, the posters' canvas (*i.e.* the matrix). The content is dynamically gathered from the Twitter API using a textual input given by the user (*i.e.* the client). In this process, an Natural Language Understanding (NLU) classifier and lexicon-based approaches recognise sentiments and emotions in the gathered content, and a Evolutionary Computation (EC) approach automatically generates and evolves the outputs. Each generated poster is evaluated according to its (I) legibility, *i.e.* how much content it is possible to read in the poster, (II) aesthetics, *i.e.* how much the design of the poster satisfies a set of aesthetics measures for typographic poster design, and (III) semantics *i.e.* how much the visual characteristics of the poster convey the semantic meaning of its content. The merit of each poster is assigned by a hybrid fitness scheme that combines a hardwired fitness function technique with a multi-objective optimisation approach. The users may communicate their preferences to the system throughout a design guidelines sheet. This system, which is available online at <http://pf.dei.uc.pt/et/>, is aligned with our previous experiments with hardwired fitness assignment methods (see [47] and [46]).

The outputs generated by the system fully communicate their content while achieving high levels of visual balance and expressiveness. The system addresses the contemporary phenomenon of posting, exploring the usage of posters as a canvas for personal and ephemeral expression. A relatively unexplored subject that goes against the common public and informative nature of posters [4]. Also, the system unveils how computational design techniques may expand the tools and automate some processes in GD, creating novel ways to communicate with people.

The key technical contributions presented in this paper include (I) a generative system capable of automatically designing typographic posters, regardless the length and purpose of the content, (II) a method to recognise sentiments, emotions and colours associated with tweets that combines an NLU network with lexicon-based approaches, (III) an evolutionary framework to generate typographic poster designs, and (IV) a method to evaluate typographic poster de-

signs, combining a hardwired fitness assignment method with a multi-objective optimisation approach.

The remainder of this paper is organised as follows. Section 2 summarises the related work. Section 3 comprehensively describes the system. Section 4 reports the experiments conducted to analyse the behaviour of the system and presents some results. Finally, Section 5 draws the conclusions and points the directions for future work.

2 Related Work

The use of computational processes to generate visuals already existed in the earlier times of the second half of the 20th century [23]. However, the introduction of the personal computer approximated graphic designers from these processes. From then on, a growing number of designers began to use computer programming as a tool to generate visual artefacts that solve problems in a flexible and customised way. Muriel Cooper, and her Visible Language Workshop, at MIT, and John Maeda were pioneers that explored on poster generation, using presenting tailor-made software (*e.g.* [7] and [31]).

Several graphic designers explored, afterwards, the use of these technologies. Most of these designers were focused on the generation of visuals to be used on their designs (*e.g.* [39] or [16]). Nevertheless, as far as we know, some automatic approaches were developed. LUST, in 2008, employed a generative system to generate posters using content gathered from multiple internet sources in the installation *Poster Wall for the 21st Century* featured in Graphic Design Museum in Breda (Netherlands) [29]. Between 2008 and 2019, Stephan and Haag [53] generated parametric posters optimised for cheap reproduction, using Bash scripts. In 2010, Cleveland [6] proposed a method for generating style-based layouts that explores the inter-relationships between text and graphics. Also, he presented a system to generate layouts employing these principles. Damera-Venkata et al. [9] presented, in 2011, a template-based probabilistic framework for generating document layouts for variable content. In 2014, LUST presented the interactive installation *Camera Postura* that creates posters using movies' frames where the actors' gestures are similar to the viewers' gestures [30]. In the same year, O'Donovan [40] presented an energy-based approach for designing single-page layouts. In 2016, Bleser [10] developed the *Pita Style Generator* that designs posters according to a specific style using the interactive design tool *Logic Layout*. In 2017, Zhang [55] developed a system that automatically generates banners of varied sizes using a style parameter learned from a set of training examples. More recently, in 2019, Rodenbröker [51] lectured the workshop *Programming Posters*, where the attendants designed posters using computer programming. Rebelo et al. [48] presented, in the same year, an installation that designs posters according to the state of the physical surroundings around them and, simultaneously, learns how to design successful posters for the place where it is placed. Moreover, there is an increasing interest in the employment of deep

learning approaches to solve and study layout generation problems (*e.g.* [27] or [56]).

EC has been used with success in several creative domains related to GD, such as the generation of pictorial symbols (*e.g.* [11], [8]), type designs (*e.g.* [25], [32]) or web designs (*e.g.* [41]), *etc.* Lewis [26] presents a good overview of the field. Although few, some related work exists in the context of document and poster design. In 2002, Gatarski [14] developed a system to evolve automatically digital banners using the user’s click-through as fitness metric. In the same year, Goldenberg [15] employed an EC to automatically generate page layouts minimising the waste of space on the page. Purvis et al. [44] presented, in 2003, a multi-objective optimisation approach to automatically generate document layouts that satisfies certain content and layout constraints, as well as certain desired design aesthetics. Quiroz et al. [45] developed, in 2009, an evolutionary approach to generate brochure documents where the user guides the system assessing only a small subset of the results. In 2010, Morcillo et al. [38] created *GAUDII*, an evolutionary system that generates single-page designs using Interactive Evolutionary Computation (IEC) to define the design properties. Önduygu [42] developed, soon thereafter, *Gráphagos*, an IEC system that generates design compositions through the evolution of position, scale and rotation of a set of visual elements. In 2011, Kitamura and Kanoh [20] evolved poster designs using IEC to assess visual features on the generated posters (such as the size or colour). In 2012, Denis Klein developed *Crossing, Mixing, Mutating* [21], a tool to generate variations of a template using genetic operators. Subsequently, in 2016, he and Lisa Reimann updated the tool and released it as an Adobe InDesign plug-in named *Evolving Layout* [24].

3 The System

The presented system is a web application that generates letterpress-inspired typographic posters, displaying content posted online on Twitter. The system generates posters by organising a certain content, divided into text boxes, to fulfil the available space on a canvas. It performs this task through the employment of 3 modules: (I) Input Processing; (II) Evolution; and (III) Evaluation.

The system dynamically gathers the content using Twitter API, based on a textual input given by the user. The Input Processing module is responsible to gather the content, processing it, and recognising its most important parts. Thus, it recognises the sentiments, emotions and colours associated with the content, employing an NLU classifier and lexicon-based approaches. The Evolution module randomly initialises a new population of candidate solutions and uses a Genetic Algorithm (GA) to evolve this population. In this work, a candidate solution, or an individual, consists of a poster design. The evolution is performed by the iterative employment of variation operators (*i.e.* mutation and recombination) on the most promising poster designs. These posters are selected by tournament based on their fitness. This method practices elitism, persevering the best individual of each generation to the next generation. The fitness of

each poster is assigned by the Evaluation module through a hybrid method that blends characteristics from hardwired and multi-objective optimisation assignments. The fitness calculation is based on 3 objectives: (I) legibility; (II) aesthetics; and (III) semantics. The legibility objective assesses how much content it is possible to read on the poster. The aesthetics objective assesses how much the individuals fulfil a set of aesthetics measures for typographic poster designs. The semantics objective assesses how much the individuals' visual characteristics convey the semantic meaning of its content.

The users of the system can communicate their preferences to the system by fulfilling the design guidelines sheet. In this sheet, users can define a set of core variables for the functioning of the system. These variables are the weights of the evaluation of each objective, the GA setup parameters, the available typefaces and their scores of pairing, as well as their emotional scores. The system will take into consideration these variables in its generation and evaluation processes. A schematic of the system is overviewed in Figure 1.

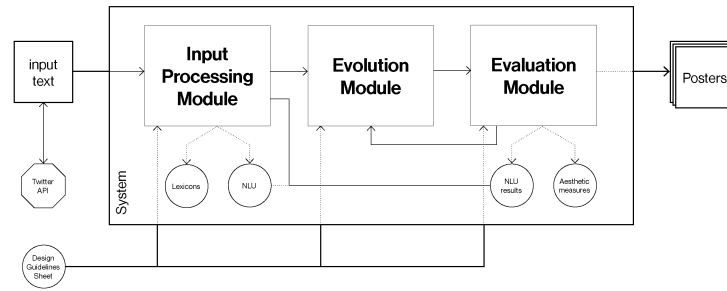


Fig. 1. Schematic of the system's architecture.

3.1 Input Processing Module

The Input Processing module dynamically gathers and processes the content to be placed on poster. The content is gathered using Twitter API based on a textual input given by users throughout a specific form (see Figure 2). This way, when a user inputs a text string, a standard search in Twitter API is performed and the tweet that better meets the query is returned. In this form, the users can select what type of search they prefer and define the search's specifications. The system performs 3 types of searches: (I) general search, *i.e.* to search by tweets that contain the inputted text; (II) username search, *i.e.* to search by tweets of a specific Twitter user; and (III) hashtag search, *i.e.* to search by tweets with a specific hashtag. Also, users may select if the tweet returned will be selected based on its popularity, newness, or both.

The present module analysis the gathered content to recognise sentiments and emotions. It performs this analysis on global and local levels. The global

analysis recognises sentiments and emotions transmitted by the text as a whole. We achieve this by implementing a Bayes NLU classifier that recognises sentiments and the intensity of certain emotions in the tweets. The sentiment analysis is performed in multiple languages and on the positive-negative axis. Furthermore, this network was trained to identify the intensity of certain emotions on text using a manually annotated data set of tweets, in the English language, collected by Mohammad and Bravo-Marquez [36]. Currently, this model can recognise 4 emotions (anger, fear, joy, and sadness) and their intensity on the text. The NLU classifier was implemented and trained using the tools available at NLP.JS library [3]. The local analysis recognises emotions and sentiments in every word on the text. Thus, after tokenising and lemmatising the text, each word is searched in a word-emotion association lexicon. The used lexicon, developed by Mohammad and Turney [37], enables to recognise 8 basic and prototypical emotions (*i.e.* anger, anticipation, disgust, fear, joy, sadness and surprise) [43] and 2 sentiments (*i.e.* positive and negative). The local analysis, therefore, perceives what are the more emotional parts of the text and, so, the parts that should be emphasised in the outputs. Currently, this analysis also uses more established emotional data, enabling to bridge the current limitations of the global analysis.

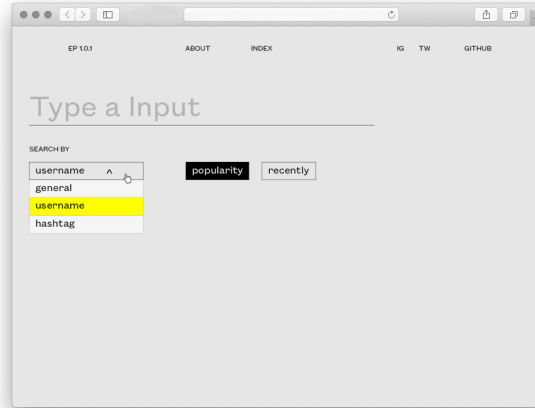


Fig. 2. Screenshot of the interaction form. The users can input a text string in the text area on the top. In the drop-down selector, below at left, the user may select the type of search. In the buttons, at the right, the user may define the way how the tweet is selected.

The module also analyses the content to recognise colours associated with it. Similarly to the local analysis, this analysis is performed using a word-colour association lexicon by Mohammad [35]. The used lexicon relates several words with 11 colours: black; blue; brown; green; grey; orange; purple; pink; red; white;

and yellow. In the end, this analysis creates an annotated map that describes the intensity of the relation between the colours and the content. The intensity of each colour is the sum of its scores of association with the words of the content.

At the end of the analysis, the present module splits the text into lines as follows. First, it performs a Sentence Boundary Detection [49] method to divide the content into sentences. After, it subdivides into lines the sentences that are lengthier than a predefined maximum characters threshold. In this subdivision, an optimal characters length threshold defines the probability of breaking the sentences. This probability increases when the sentence's length approaches a maximum characters threshold. The optimal threshold is randomly defined in each subdivision based on a predefined range, creating, therefore, more organic subdivisions of the content. The maximum threshold and the optimal range are defined in the design guidelines sheet.

3.2 Evolution Module

The Evolution module implements a GA to create a population of poster designs at random and, subsequently, evolve them using variation operators, *i.e.* crossover and mutation (*e.g.* see [12] for further information about GAs). Also, this module provides the necessary methods to render and export the outputs.

Each poster is a set of arranged text boxes. The text boxes are encoded as a sequence of numeric arrays (*i.e.* the genotype). The first array in the sequence encodes the poster's background colour (*i.e.* the colour configuration gene). The following arrays encode the text boxes (*i.e.* the text boxes genes). Figure 3 schemes the genotype of the individuals. The colour configuration gene is a two-position length array, where the first position encodes the base colour and the second encodes percentage of the tint. The text boxes genes are four-position length arrays that encode the typeface, the font's weight, the height, and the font size in percentages of the height, respectively. Since posters' contents may have different lengths, the number of text boxes and, consequently, the size of genotype may vary.




Fig. 3. The genotype is encoded as a sequence of numeric arrays. The first array (the colour configuration gene) encodes the background colour of the output, *i.e.* the base colour (C1) and the percentage of tint (C2). The following arrays (text boxes configuration genes) encode the properties of the text boxes, *i.e.* the typeface (TA1), the font's weight (TA2), the height (TA3) and the font size (TA4).

The posters' canvas is subdivided in a one-column grid with multiple rows. This grid constraints the text boxes position and sizes. The canvas has, by de-

Variation Poster designs are evolved iteratively through the employment of crossover and mutation operators. Both operators are designed to preserve the validity of the generated individuals.

The crossover operator exchanges genes between two parents to generate new posters designs. Thus, it randomly selects two parents regarding their fitness and, after, employs a uniform crossover method, *i.e.* it randomly selects, for each gene, what will be the parent that will give the gene to the children [54]. This operator does not crossover the genes related to the height of the text boxes, ensuring that the generated individuals are always valid (*i.e.* posters where the text boxes fill all the available canvas space).

The mutator operators perform random modifications in several parts of the individuals' genotype. We designed these operators ensuring that they cover all the search space. This resulted in 3 operators: (I) independent; (II) specialised and (III) swap. The mutations are performed based on a certain probability. Thus, the system for each candidate solution, in the new offspring, randomly defines if it will be mutated and, subsequently, selects the mutation operator. Each operator has the same probability of being employed.

The independent mutation operator random selects a gene and, subsequently, a parameter inside of the gene and mutates it. The only exception is the font-size parameter, which can not be selected by this operator. Each parameter has its own range of values. Thus, a customised method is employed depending on the parameter that is selected. If the base colour or the typeface parameter are selected, *i.e.* an unsubordinated parameter, the new value is changed at random according to the options available. The modification of these parameters may also require the change of their subordinated parameters (*i.e.* the percentage of tint or the font's weight). Thus, if the resulting mutated unsubordinated parameter value does not support the value of the subordinated parameter, a new value is defined at random. On the other hand, if a subordinated parameter is selected, the value is randomly modified according to the possibilities available for the unsubordinated parameter. If the text box height parameter is selected, two genes are randomly selected, having one, at least, its value bigger than 1. After, the module decides what will be the gene that will decrease the height and the one that will increase. This selection is made at random, unless one of the selected genes has a value of 1. In this case, the gene with a value of 1 will increase its height and the other will decrease.

The specialised mutation operator is a custom implementation of the independent mutation method. This operator mutates the highest text box on the individual and the font size parameter. We empirically observed that the evolution of the posters often was slow and easily stabilised in specific designs. This was mostly because a text box was much bigger than the average and/or some content was too long for text box and, so, it was not fully displayed, even when it was composed in the most condensed typeface. This way, we designed bespoke methods to perform regular modifications on these parameters, allowing a faster and more diverse evolution. When this mutator is selected one method is randomly chosen and performed in the individual. The mutator of the highest

text box decreases the height of the highest text box and increase the height of another randomly selected text box in the individual. The mutation of the font size randomly selects a text box gene and decreases, or increases, the value of its font size parameter in 1%. The direction of the mutation is defined at random unless the parameter value is the maximum, *i.e.* 100%, (the value will only be decreased) or the minimum, *i.e.* 30% (the value will only be increased).

The swap mutation operator, as the name indicates, randomly selects two text boxes, in the same individual, and swaps the value of their genes.

Export The presented module also provides the necessary means to export the outputs, when necessary. The users may export one or multiple outputs at any time during the evolutionary process. When the user exports an output, 4 files are downloaded: the vector (SVG file) and raster (PNG file) versions of the phenotype, the genotype of the output (JSON file) and the content (text file).

3.3 Evaluation Module

The Evaluation module assesses the generated outputs according to 3 objectives: (I) legibility, *i.e.* how much content it is possible to read on the poster; (II) aesthetics, *i.e.* how much of the poster design satisfies a set of aesthetics measures for typographic poster design; and (III) semantics, *i.e.* how much of the poster's visual characteristics convey the semantic meaning of its content. Each objective has its own evaluation method.

The present module implements an automatic fitness assignment hybrid scheme that combines a hardwired fitness function part with a multi-objective optimisation approach part. The fitness of each poster is calculated by a weighted arithmetic mean of legibility (the hardwired part) with the relation between the semantics and aesthetics objectives (the multi-objective part). This way, it values fully legible posters while, simultaneously, searches by different relations between the aesthetics and semantics objectives. This relation can be a balance or the optimisation of one objective over another. The overall fitness value ranges from 0 (bad) to 1 (good).

The relation between the aesthetics and the semantics objectives is calculated in the following way. First, the module sorts the entire population based on the non-domination of each individual (see [52]). A solution dominates others if it is better in at least one objective and same as, or better, in another objective. A non-dominated solution is, therefore, a solution that is not dominated by any other solution in the population. Next, the population is organised in fronts. The first front is populated by the non-dominant set, the second front by the solutions dominated by the ones only in the first front and the fronts goes so on. The value of this relation for an individual is the rank of the front where it is placed, normalised according to the number of fronts.

This evaluation process is in accordance with the discussions about how we should measure the quality in GD. Posters' evaluation is a subjective task influenced by multiple factors such as the purpose, the *zeitgeist*, the context,

the target people, *etc.* (see [34] and [17]). In this sense, the present module evaluates the outputs, considering that a poster should communicate its content above all. Besides that, there is no direct way to evaluate the quality of a poster and it (like other GD’s artefacts) should not only be evaluated by its aesthetics [13]. This way, like graphic designers, the system designs posters by trying to balance the pure aesthetics with the semantics meaning of the content, valuing one characteristic over the other to understand where the borderline is, when the overall quality of a poster begins to decrease.

Legibility The legibility objective measures how much of the content is legible on the poster. The overall legibility value of a poster is related to the legibility of its text boxes. The legibility of a text box is the difference between its target width (*i.e.* the available width of the poster) and the width of its content when rendered. This difference is calculated considering that the content should be always rendered inside of the poster and the negative space inside of the text box (*i.e.* the space coloured with the background colour) should, as much as possible, be minimised. Thus, the legibility of a text box is this difference mapped to assign a poor assessment when the rendered text exceeds the width of the poster and, simultaneously, to prejudice progressively the assessment as soon as the amount of negative space inside text box surpasses a certain threshold. The overall legibility value of a poster is the weighted arithmetic mean of the value of text boxes. The weight of each text box in the mean is given according to its height. In the end, the value is normalised between 0 (bad) and 1 (good).

Aesthetics The aesthetics objective measures how much the poster design satisfies a set of aesthetic measures for typographic posters’ design. We defined these measures based on the work of Harrington et al. [18]. However, we decided not to measure the white-space free-flow, the proportion, and the uniform separation because these measures are not applicable to the generated outputs since they are composed in a one-column grid. The outputs are typographic posters and a good pairing of the typefaces is a key factor to create harmonious layouts [5, 28]. In this sense, we also added a font pairing measure. The aesthetics of a poster is, then, evaluated according to (I) the alignment, (II) the regularity, (III) the balance, (IV) the negative-space fraction, (V) the composition security, and (VI) the font pairing. The overall aesthetics measure is the weighted arithmetic mean of these attributes ranged from 0 (bad) to 1 (good). The weight of each measure, as well as the optimal thresholds and offsets of certain measures, are defined in the design guidelines sheet.

The alignment measures how much the edges of the content share similar horizontal positions on a poster. This measure depends on the distance between neighbouring text boxes. Thus, the module creates a histogram with the current positions of the left edges of the text boxes and after compares the distance between the values on adjacent positions. The closer the distance values are, the higher is the alignment score. The overall alignment measure is the arithmetic mean of all distances.

The regularity measures how much regular is the placement of text boxes in vertical on a poster, *i.e.* if the heights of text boxes are similar. The calculation of the regularity is a process similar to the calculation of alignment. However, it uses a histogram with the positions of the top edges instead of the left edges.

The balance measures how much the posters are centrally balanced. The centre balance measure of a composition is the difference between the centre of the visual weight and the visual centre. The centre of the visual weight is calculated based on the visual weights and centres of the text boxes. The visual weight of a text box is defined by its area times its optical density. The optical density of a text box is the 10 logarithm of the average normalised of its Luminance (Y). The Y is calculated through the weighted average formula of $Y = 0.2125R + 0.7152G + 0.0722B$ where R, G, B are the red, green, and blue channels from the average pixel, respectively. The position of the visual centre of a text box is the horizontal geometric centre and the vertical geometric centre with a small offset towards the top of the page. The poster's centre of the visual weight and the overall measure of balance is calculated according to the method presented by Harrington et al. [18].

The negative-space fraction measures if the space coloured with the background colour is balanced in the poster. This measure is the distance of the current percentage of pixels coloured with the background colour to a certain optimal percentage threshold.

The composition security measures if the text boxes positioned near the edges of the poster are secure, *i.e.* small text boxes when placed near the composition edges appear to fall off. The security of each text box is the minimum between the top edge and bottom edge. The overall value is the minimum between the values of all text boxes.

The font pairing measures if the employed typefaces pair well between them. The pairing value of each text box is the arithmetic mean of pairing scores between its typeface with other typefaces used on the poster. The overall value is the arithmetic mean of the value of all text boxes.

Semantics The semantics objective measures how much the posters' visual characteristics convey the semantic meaning of its content. This way, the most important parts of a poster's content should be emphasised over the less important ones, *i.e.* they should be composed on higher text boxes. Also, the background colour and the typography on the poster should transmit the semantic meaning of the content. The semantics are evaluated, then, according to (I) the background colour, (II) the layout of the text boxes, and (III) the typefaces employed. The overall semantic measure is the weighted arithmetic mean of these attributes ranged from 0 (bad) to 1 (good). The weight of each measure in the mean is defined in the design guidelines sheet.

The background colour of a poster conveys the semantic meaning of its content when it is aligned with the annotated map that describes the relationship between the colours and the content defined before. As before mentioned, the colours in this map are sorted by intensity. This way, the more intense is the

relation of one colour with the content, the fewer should be the percentage of tint employed and *vice versa*. Every colour in the map obtains a good assignment when used on output with the proper tint value. The optimal tint value of a colour is the quotient of the number of available percentages of the tint of this colour by the intensity of the relationship with the content (*i.e.* the position of the colour on the map). The overall measure of the appropriateness of the background colour is the normalised distance of the optimal tint value to the current value.

The layout of a poster should emphasise the most important parts of the content by assigning higher text boxes to them. We consider that the more important text boxes are those with a higher amount of emotions recognised in the local sentimental and emotional analysis performed before. Each recognised emotion has a certain weight on the layout. This weight is the division of the total number of emotions recognised in all text boxes by the number of rows on the grid. The optimal text box height is the product of the number of emotion present in its content by the weight of each emotion. The appropriateness of each text box's layout is the normalised distance between its current height and its optimal height. The overall value is the normalised arithmetic mean of all distances.

The typefaces employed on a poster convey the semantic meaning of its content when their shape and weight reflects the sentiments and emotions present in the content. This measure is the arithmetic mean of the other 3 measures: (I) global appropriateness; (II) local appropriateness; and (III) font's weight appropriateness. The global appropriateness measures if the typefaces employed on poster convey the results of the global sentimental and emotional analysis performed before. Thus, the global appropriateness of a typeface in a poster is the arithmetic mean of the differences between the intensities of the emotions and sentiments present on the content and the intensities that this typeface conveys these emotions and sentiments. The overall value is the arithmetic mean of the value of all typefaces used on the poster. The local appropriateness measures if the typefaces employed on the poster convey the results of the local sentimental and emotional analysis performed before. The calculation of this measure is similar to the global measure, however, performed at the text boxes level. Thus, the local appropriateness of a typeface in a text box is the arithmetic mean of the difference of the intensities that this typeface conveys the emotions and sentiments present on the text box. The overall value is the arithmetic mean of the value of all text boxes. The font's weight appropriateness measures how the weight of the typeface employed in a text box conveys the results of the local sentimental and emotional analysis. This measure was created because certain emotions are easier conveyed by the typeface weight than by typeface design [19, 22]. The font's weight appropriateness of a text box is the distance between the current font's weight and an optimal range. The optimal range is the average of the optimal ranges of all emotions present on the text box. The overall value is the arithmetic mean of the values of all text boxes.

4 Experiments

We conducted experiments to study and analyse the possibilities of the system, by evolving posters for contents with different lengths and textual purposes. The experimental parameters used in these experiments are defined by empirical exploration and are summarised in table 1. The weights of the parameters on the fitness assignment and on the evaluation of the aesthetics and semantics objectives were also defined by empirical exploration and are summarised in table 2.

Table 1. Experimental parameters.

Parameter	Value
Generations	500
Population size	30
Elite size	1
Selection	Tournament
Tournament size	2
Mutation probability	0.7
Phenotype size	298×420
Margin size	15px
Grid	26×1
Maximum line length (in characters)	40
Optimal line length (in characters)	[10–30]
Visual centre vertical offset	1/12
Optimal percent of negative space	50%

Table 2. Experimental weights on evaluation components.

Evaluation	Parameter	Weight
Fitness	Legibility	90 %
	Aesthetics/Semantics	10 %
Semantics	Background Colour	25 %
	Layout	25 %
	Typographic choices	50 %
Aesthetics	Alignment	10 %
	Regularity	10 %
	Balance	15 %
	Negative-Space Fraction	10 %
	Page Security	05 %
	Font pairing	50 %

The experiments were conducted using 15 typographic families. The type-faces, their pairing scores as well as their emotional and sentimental score, were

empirically defined by us in the design guidelines sheet. The typographic families were dynamically loaded using the Adobe Typekit Webfonts service [1]. All the typefaces in the typographic family are loaded, except the italic and slanted ones. We defined the emotional and sentimental score of each typeface based on the work of Koch [22] and Hyndman [19]. The pairing scores were defined based on the guidelines presented by Bringhurst [5] and Lupton [28]. One may analyse the scores of each typeface in the design guidelines sheet example used on these experiments and available for download at <https://cdv.dei.uc.pt/evoposter/>. The typeface families selected are the following: (i) *LTC Globe Gothic* designed by Colin Kahn, Frederic W. Goudy, and Morris Fuller Benton (Lanston Type, 1897); (ii) *News Gothic* designed by Morris Fuller Benton (Adobe, 2007); (iii) *Futura PT* designed by Isabella Chaeva, Paul Renner, Vladimir Andrich, and Vladimir Yefimov (ParaType, 2007); (iv) *Century Old Style* designed by Morris Fuller Benton (Adobe, 2007); (v) *Azo Sans* designed by Rui Abreu (R-Typography, 2013); (vi) *Franklin Gothic* designed by Morris Fuller Benton (URW, 2002); (vii) *Clone Rounded* designed by Lasko Dzurovski (Rosetta Type Foundry, 2016); (viii) *Bureau Grot* designed by David Berlow (Font Bureau, 1989); (vii) *Titling Gothic FB* designed by David Berlow (Font Bureau, 2005); (ix) *Benton Modern Display* designed by David Berlow, Dyana Weissman, Richard Lipton and Tobias Frere-Jones (Font Bureau, 2008); (x) *Miller Display* designed by Matthew Carter (Carter & Cone, 1997); (xi) *Whitman Display* designed by Kent Lew (Font Bureau, 2008); (xii) *Bodoni FB* by Richard Lipton (Font Bureau, 1989); (xiii) *Trade Gothic Next* designed by Akira Kobayashi, Tom Grace and Jackson Burke (Linotype, 2008); (xiv) *Zeitung Pro* designed by Underware (2017); and (xv) *Zeitung Mono* designed by Underware (2017).



Fig. 5. Typical outputs generated by the system using contents gathered from tweets posted online between 14/04 and 29/04/2020 by the users @elporrote, @jessphillips, @joy, @latimes, @mypaws, @realDonaldTrump and @RealKunalMC. More example outputs are accessible at <https://cdv.dei.uc.pt/evoposter/>.

Figure 5 display several obtained results. One may see more examples of the system outputs and demo video of the system in <https://cdv.dei.uc.pt/evoposter/>. Visually observing the results, one can conclude that the system often prefers to use typefaces from the same typographic family. This way, more extensive typographic families, such as *Bureau Grot* or *Titling Gothic FB*, are more observable in the results. The reason is simple: the extension of these typographic families ensures that often exists a proper typeface for each text box size as well as achieving good scores of pairing by using them together, since they are part of the same family and share common traits. Concerning the background colours, it is possible to see that for some contents, the system tends to select always the same set of colours over different runs. Also, we observed that the system evolves well-evaluated outputs faster when the content is lengthier, since the number of possible solutions is minor and, so, it needs to perform less operations to achieve a good result. Nevertheless, the lengthier the content, the more similar will the achieved outputs be over different runs.




Fig. 6. Evolution of the posters’ fitness (top) and objectives evaluation (bottom) over the generations. In the figure above, the solid line displays the fitness’ best individual. The semi-transparent line displays the average fitness of the population. In the figure below, the blue, green, and purple lines display the legibility, aesthetics, and semantics objectives, respectively. The solid lines display the fitness’ best individual. The semi-transparent lines display the average fitness of the population. The visualised data is the average of 30 runs.

Figure 6 displays the evolution of the fitness, the evolution of the objectives of the best individual and the average of the population, over the generations. The displayed data is the average of 30 runs, using 6 different contents with an average of 83 characters by content. The fitness of the individuals (the top of Figure 6) is a qualitative value since the relation between the aesthetics and semantics objectives is not a quantitative value, *i.e.* it corresponds to the front that the individual belongs. This way, to study the system’s behaviour, this

value should be observed in conjunction with the evaluation of the objectives on the population (the bottom of Figure 6). This way, one can observe that high fitness values are attained in a few generations and this value is strictly related to the evaluation of the legibility objective. It is also observable that, in the early stages, the evolution is faster than in the later stages. This occurs because, in the earliest stages, the evolution is mostly focused on generating a good layout and in the later ones it is mostly focused on performing minor adjustments to increase the fitness of the best individual and, so, of the entire population.

One can also observe that in the earlier stages of evolution, the evaluation of aesthetics and semantics achieves higher values than in the later ones and, often, these values decrease over the generations. Also, it is observable that in these stages the average evaluation of these objectives in the population is often above to their evaluation of the best individual. This occurs because the legibility acts like a constraint and in the earlier stages of evolution, posters have poor legibility evaluations. Nevertheless, the values of the evaluation of aesthetics and semantics objectives are always directly related to the characteristics of the content.

5 Conclusions and Future Work

We have described a system that generates letterpress-inspired posters, using tweets posted online as content. The content is dynamically gathered based on a textual input given by the user. Outputs are generated through 3 modules: (I) Input Processing; (II) Evolution; and (III) Evaluation. In the Input Processing module, sentiments, emotions and colours are recognised on the content using an NLU classifier and lexicon-based approaches. The Evolution module employs a GA to generate and evolve a population of posters. The Evaluation module assesses the outputs according to 3 objectives: (I) legibility; (II) aesthetics; and (III) semantics. The merit of each poster is assigned by a hybrid fitness scheme that combines a hardwired fitness function part (the legibility) with a multi-objective optimisation part (relation between the aesthetics and semantics). We experimented with the system to perceive its behaviour and its ability to evolve posters. The system outputs achieves high levels of legibility and diversity. Also, we observed that the legibility objective works as a constraint, allowing the system to balance the relation between aesthetics and semantics to find the best relationship between these two objectives, without decreasing the legibility.

Future work will focus on (I) to explore different fitness assignment schemes (which may promote diversity and visual novelty in the generated outputs), (II) to design and to develop an interface that allows a parametric definition of several system's variables and to lock several visual properties of the outputs during the evolution, (III) to automate the gathering of the typefaces and the definition of their pairing, emotional and sentimental scores (using *e.g.* Adobe Typekit Webfonts or Google Fonts APIs), (IV) to include images and illustrations in the posters, and (V) to create real letterpress posters, based on the outputs generated by the system.

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: Adobe typekit web fonts (2020), <https://fonts.adobe.com/typekit/>
2. Armstrong, H., Stojmirovic, Z.: Participate: Designing with User-Generated Content. Princeton Architectural Press, New York, NY (2011)
3. AXA Group Operations Spain S.A.: Nlp.js (April 2020), <https://github.com/axa-group/nlp.js/>
4. Blauvelt, A.: The persistence of posters. In: Blauvelt, A., Lupton, E. (eds.) *Graphic Design: Now in Production*, chap. 11, pp. 92–111. Walker Art Center, Minneapolis, MN (2011)
5. Bringhurst, R.: *The Elements of Typographic Style*. Hartley & Marks, Vancouver, Canada, 3rd edn. (2005)
6. Cleveland, P.: Style based automated graphic layouts. *Design Studies* **31**(1), 3–25 (January 2010)
7. Cooper, M.: Computers and design. *Design Quarterly* **1**(142), 1–31 (1989)
8. Cunha, J.M., Lourenço, N., Correia, J., Martins, P., Machado, P.: Emojinating: Evolving emoji blends. In: Ekárt, A., Liapis, A., Castro Pena, M. (eds.) *EvoMUSART 2019: Computational Intelligence in Music, Sound, Art and Design*. pp. 110–126. Springer-Verlag, Leipzig, Germany (24–26 April 2019 2019)
9. Damera-Venkata, N., Bento, J., O’Brien-Strain, E.: Probabilistic document model for automated document composition. In: *Proceedings of the 11th ACM Symposium on Document Engineering*. pp. 3–12. ACM, Mountain View, CA (September, 2011 2011)
10. De Bleser, F.: *Generative Design: The NodeBox Perspective*. Ph.D. thesis, University of Antwerp, Antwerp, Belgium (2016)
11. Dorris, N., Carnahan, B., Orsini, L., Kuntz, L.A.: Interactive evolutionary design of anthropomorphic symbols. In: *Proceedings of the 2004 Congress on Evolutionary Computation*. pp. 433–440. IEEE, Portland, OR (19-23 June 2004 2004)
12. Eiben, A.E., Smith, J.E.: *Introduction to Evolutionary Computing*. Springer-Verlag, Heidelberg, Germany, 2nd edn. (2015)
13. Frascara, J.: Graphic design: Fine art or social science? *Design Issues* **5**(1), 18–29 (Autumn 1988)
14. Gatarski, R.: Breed better banners: Design automation through on-line interaction. *Journal of Interactive Marketing* **16**(1), 2–13 (November 2002)
15. Goldenberg, E.: *Automatic Layout of Variable-content Print Data*. Master’s thesis, University of Sussex, Brighton, United Kingdom (2002)
16. Groß, B., Laub, J.: *Diploma – generative systeme posters* (2007), <https://benedikt-gross.de/projects/diploma-generative-systeme-posters/>
17. Guffey, E.E.: *Posters: A Global History*. Reaktion Books, London, United Kingdom (2014)
18. Harrington, S.J., Naveda, J.F., Jones, R.P., Roetling, P., Thakkar, N.: Aesthetic measures for automated document layout. In: *Proceedings of the 2004 ACM Symposium on Document Engineering*. pp. 109–111. ACM, Milwaukee, WI (October 2004 2004)

19. Hyndman, S.: *Why Fonts Matter*. Virgin Books, London, United Kingdom (2016)
20. Kitamura, S., Kanoh, H.: Developing support system for making posters with interactive evolutionary computation. In: 2011 Fourth International Symposium on Computational Intelligence and Design. pp. 48–51. IEEE, Hangzhou, China (28-30 October 2011 2011)
21. Klein, D.: Crossing, mixing, mutation (2012), <http://www.gutenberg-intermedia.de/wissenschaft-gestaltung/denis-klein-crossing-mixing-mutation/>
22. Koch, B.E.: Emotion in typographic design: An empirical examination. *Visible Language* **46**(3), 206–227 (2012)
23. Leavitt, R.: *Artist and Computer*. Harmony Books, New York, NY (1976)
24. LESS: Evolving layout (2016), <http://www.evolvinglayout.com/>
25. Levin, G., Feinberg, J., Curtis, C.: Alphabet synthesis machine (2002), <http://web.archive.org/web/20080513044335/http://www.alphabetsynthesis.com/>
26. Lewis, M.: Evolutionary visual art and design. In: Romero, J., Machado, P. (eds.) *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, chap. 1, pp. 3–37. Springer-Verlag, Heidelberg, Germany (2008)
27. Li, J., Yang, J., Hertzmann, A., Zhang, J., Xu, T.: Layoutgan: Generating graphic layouts with wireframe discriminators. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (January 2020)
28. Lupton, E.: *Thinking With Type: A Critical Guide for Designers, Writers, Editors, & Students*. Princeton Architectural Press, New York, NY, 2nd edn. (2010)
29. LUST: Graphic design museum: Poster wall for the 21st century (2008), <https://lust.nl/{#}projects-3041/>
30. LUST: Camera postura (2014), <http://lust.nl/{#}projects-5939/>
31. Maeda, J.: *Maeda@Media*. Thames & Hudson, London, United Kingdom (2000)
32. Martins, T., Correia, J.a., Costa, E., Machado, P.: Evolving stencils for typefaces: Combining machine learning, user’s preferences and novelty. *Complexity* **2019** (March 2019)
33. McCarthy, L., Processing Foundation, NYU ITP: P5.js (2020), <https://p5js.org/>
34. Meggs, P.B., Purvis, A.W.: *Meggs’ History of Graphic Design*. John Wiley & Sons, New York, NY, 6th edn. (2016)
35. Mohammad, S.M.: Colourful language: Measuring word-colour associations. In: *Proceedings of the Second Workshop on Cognitive Modeling and Computational Linguistics*. pp. 97–106. ACL, Portland, OR (June 2011 2011)
36. Mohammad, S.M., Bravo-Marquez, F.: Emotion intensities in tweets. In: *Proceedings of the Sixth Joint Conference on Lexical and Computational Semantics (*Sem)*. pp. 65–77. ACL, Vancouver, Canada (August, 2017 2017)
37. Mohammad, S.M., Turney, P.D.: Crowdsourcing a word-emotion association lexicon. *Computational Intelligence* **29**(3), 436–465 (September 2012)
38. Morcillo, C.G., Martin, V.J., Fernandez, D.V., Sanchez, J.J.C., Albusac, J.A.: Gaudii: An automated graphic design expert system. In: *Proceedings of the Twenty-Second Conference on Innovative Applications of Artificial Intelligence*. pp. 1775–1780. AAAI, Atlanta, GA (11–15 July 2010 2010)
39. Müller, B.: *Poetry on the road 2002–2013* (2002), <https://www.esono.com/boris/projects/poetry02/>
40. O’Donovan, P., Agarwala, A., Hertzmann, A.: Learning layouts for single-page graphic designs. *IEEE Transactions on Visualization and Computer Graphics* **20**(8), 1200–1213 (August 2014)

41. Oliver, A., Monmarché, N., Venturini, G.: Interactive design of web sites with a genetic algorithm. In: Isaías, P. (ed.) *Proceedings of the IADIS International Conference WWW/INTERNET*. pp. 355–362. Lisbon, Portugal (13–15 November 2002 2002)
42. Önduygu, D.C.: *Graphagos: Evolutionary Algorithm as a Model for the Creative Process and as a Tool to Create Graphic Design Products*. Master’s thesis, Sabanci University (2010)
43. Plutchik, R.: A general psychoevolutionary theory of emotion. In: Plutchik, R., Kellerman, H. (eds.) *Theories of Emotion*, chap. 1, pp. 3–33. Academic Press, Cambridge, MA (1980)
44. Purvis, L., Harrington, S., O’Sullivan, B., Freuder, E.C.: Creating personalized documents: an optimization approach. In: *Proceedings of the 2003 ACM Symposium on Document Engineering*. pp. 68–77. ACM, San Jose, CA (29 September – 2 October 2003 2003)
45. Quiroz, J.C., Banerjee, A., Louis, S.J., Dascalu, S.M.: Document design with interactive evolution. In: Damiani, E., Jeong, J., Howlett, R.J., Jain, L.C. (eds.) *New Directions in Intelligent Interactive Multimedia Systems and Services - 2*, chap. 29, pp. 309–319. Springer-Verlag, Heidelberg, Germany (2009)
46. Rebelo, S., Bicker, J., Machado, P.: Evolutionary experiments in typesetting of letterpress-inspired posters. In: Cardoso, F.A., Machado, P., Veale, T., Cunha, J.M. (eds.) *Proceedings of the Eleventh International Conference on Computational Creativity*, September 7–11, 2020. pp. 110–114. Coimbra, Portugal (2020)
47. Rebelo, S., Fonseca, C.M., Bicker, J., Machado, P.: Evolutionary experiments in the development of typographical posters. In: Rangel, A., Ribas, L., Verdicchio, M., Carvalhais, M. (eds.) *6th Conference on Computation, Communication, Aesthetics & X (xCoAx 2018)*. pp. 65–75. Universidade do Porto, Madrid, Spain (2018)
48. Rebelo, S., Pires, C., Martins, P., Bicker, J., Machado, P.: Designing posters towards a seamless integration in urban surroundings: A computational approach. In: *ARTECH 2019: Proceedings of the Ninth International Conference on Digital and Interactive Arts*. p. Article 54. ACM, Braga, Portugal (2019)
49. Reynar, J.C., Ratnaparkhi, A.: A maximum entropy approach to identifying sentence boundaries. In: *Proceedings of the Fifth Conference on Applied Natural Language Processing*. pp. 16–19. ACL, Washington, DC (March 1997 1997)
50. Richardson, A.: *Data-Driven Graphic Design: Creative Coding for Visual Communication*. Bloomsbury Publishing, London, United Kingdom (2016)
51. Rodenbröker, T.: *Programming posters* (2018), <https://timrodenbroeker.de/projects/programming-posters/>
52. Srinivas, N., Deb, K.: Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation* **2**(3), 221–248 (September 1994)
53. Stephan, B., Haag, C.: Bash scripts for generative posters. *Libre Graphics Magazine* **3**(1), 30–34 (August 2011)
54. Syswerda, G.: Uniform crossover in genetic algorithms. In: *Proceedings of the Third International Conference on Genetic Algorithms*. pp. 2–9. Morgan Kaufmann Publishers Inc., Fairfax, VA (June 1989 1989)
55. Zhang, Y., Hu, K., Ren, P., Yang, C., Xu, W., Hua, X.S.: Layout style modeling for automating banner design. In: *Proceedings of the on Thematic Workshops of ACM Multimedia 2017*. p. 451–459. ACM, Mountain View, CA (October 2017 2017)
56. Zheng, X., Qiao, X., Cao, Y., Lau, R.W.H.: Content-aware generative modeling of graphic design layouts. *ACM Transactions on Graphics* **38**(4), Article 133 (July 2019)