

Arts, News, and Poetry — The Art of Framing

Oskar Gross and Jukka M. Toivanen and Sandra Lääne and Hannu Toivonen

Department of Computer Science and Helsinki Institute for Information Technology HIIT
University of Helsinki, Finland

oskar.gross@cs.helsinki.fi, jukka.toivanen@cs.helsinki.fi, sandralaane@gmail.com, hannu.toivonen@cs.helsinki.fi

Abstract

This paper presents an art project which combines computational and human creativity. The paintings created during the project visualize a process of generating computational poetry from daily news stories. We describe how the computational processes of generating poetry were visualized and then turned into paintings by an artist. The project has been exhibited in Finland and Estonia. The feedback collected during the exhibition in Finland is also included in the paper.

Introduction

In this paper we introduce the art project *Arts, News, and Poetry* which combines human and computer creativity in a novel way. First, the computer carries out a creative process of poetry writing and produces an abstract image based on the process. The human artist then takes the poems and images as inspiration and paints them. Our motivation is to direct the audience's attention to the possibility that the inner workings of computers could be visualized and presented in some meaningful and aesthetically pleasing way.

From a computational creativity perspective we aim to introduce the possibility to use the computational processes to provide framing information for creative artefacts. The framing information is often presented to the art consumer in a natural language. In this paper, we explore an alternative approach where the information is expressed in a way which is more natural for computers.

In the rest of the paper, we first extend on the ideas behind framing, then give a brief overview of the art project as a whole. Next, we describe the poetry generation, process visualization, painting of the images and exhibitions. Then we present the related work and conclude the paper.

Framing

The way how artists explain their work has a very large influence on how the audience perceives them. A work of art might even have a completely different interpretation if we change parts of the framing information. For instance, Salvador Dalí's works with phallic symbols might have different meanings if he lived in different cultural context. Charnley et al. (2012) define the framing information as the "motivation, intention and processes involved in creating a

work". Currently, the framing information produced in computational creativity tasks incorporates information which is very humane. For instance, the Full-FACE poetry generator tells which news stories it analysed, what kind of affective words it found from there, and how it influenced the output (Colton, Goodwin, and Veale 2012).

Computers have an inhumane ability to memorize every step they make to reach a solution. We argue that this inherent feature could be taken advantage of by the computers in order to provide framing information. In this paper, we have provided a very simplistic (or even naïve) approach for solving this problem. In the ideal case, wouldn't it be interesting if the computer could visualize solving an optimization problem illustrating the drama of constantly reaching a local optimum, no matter how hard it tries?

Overview of the Art Project

The end result is a series of hand-crafted paintings, each visualizing the poem writing process of a computer and exhibited together with the computer-written poem. The topics of the poems were chosen to be based on news stories, so they could be seen as commentary to the events of the world.

The art project consists of the following steps. Steps 3–6 are further elaborated on later in the paper.

1. From 1 to 31 December 2012 we collected news stories from *BBC*, *CNN*, *Reuters*, *ABC News*, *CBS News* and *The Guardian* by automated crawling.
2. The news stories of each day were automatically clustered into 50 different topics. For clustering we used the *gensim* (Řehůřek and Sojka 2010) implementation of LDA (Latent Dirichlet Allocation) (Blei, Ng, and Jordan 2003).
3. For each topic the computer generated a topic-related poem using the methods proposed by Toivanen et al (2012; 2014) (Section *Corpus-Based Poetry Generation*).
4. For each topic an additional abstract image was created by analysing the poetry generation process (Section *Process Visualization*).
5. The abstract images and the associated poems were then presented to the artist Sandra Lääne. She hand-picked 12 image/poem pairs, and painted 12 paintings. (Section *From Abstract Images to Paintings*).

- The paintings were exhibited accompanied with the respective poems (Section *Exhibitions*).

Corpus-Based Poetry Generation

In this project, we used the poetry generation machine by Toivanen et al. (2012; 2014). The grammar, including the syntax and morphology of the generated poetry, is obtained in an instance-based manner from English poetry in Project Gutenberg as described by Toivanen et al. (2012). Thus, instead of explicitly representing a generative grammar of the output language, we copy a concrete instance from an existing text and substitute the contents by new words from the document specific associations. In contrast to the original poetry writing method (Toivanen et al. 2012), here we use a specific document or a set of documents as an input to the automatic poetry composition system. In this work, we use news stories as input documents for the poetry composition process. The topics of generated poetry are controlled by using the document specific associations as described in parallel paper by Toivanen et al (2014).

Given a document (or a set of documents), the general outline of the method is the following:

1. Calculate document specific associations by contrasting document associations to English Wikipedia as the background;
2. Choose a poem template from the poetry corpus;
3. Substitute the words in the poem template with words from the document specific associations.

Process Visualization

The abstract images were generated by using two different aspects of the poetry generation process. The geometrical composition of the images was determined by the inputs and outputs of the document specific association generation. The colour palette of the final image was defined by using colors associated to the representative words of the poem.

For generating the geometry of the composition we calculated a transformation matrix between the input text and the document specific associations produced in the poetry writing process. We consider sentences in the input text as bags-of-words. The whole text can then be represented as a binary matrix $I^{m \times n}$, where n rows correspond to words and m columns to sentences of the input text: the value of $I_{i,j}$ is 1 if the word j appears in sentence i and 0 if it does not. The matrix of document specific associations $O^{m \times m}$ is, in turn, defined as a square matrix where $O_{i,j}$ is the association strength between the words i and j . If there is no association between words i and j , the value is 0. We model the term association extraction process as a transformation matrix $P^{n \times m}$ which is obtained as a linear approximation from the following equation

$$I \times P = O.$$

Due to the sparsity of the matrix P , we reduced its dimensionality by principal component analysis and only use the top 15 principal components for aesthetic considerations.

The colour palette of the image was determined by making a Google Image Search with the 3 most important words of the news story, selected based on their sum of association strength to other words (see Toivanen et al. (2014) for details of the weight computation). The 3 words were used as the query to Google Image Search, and from the results, first 3 images were retrieved. The images were concatenated together and the Colorific tool (Hotson and Yencken 2012) was used for extracting their joint colour palette. Finally, the process matrix P was visualized using matplotlib package (Hunter 2007) and the respective colour palette.

An example image can be seen on the left in Figure 1.

From Abstract Images to Paintings

The artist hand-picked 12 computer-generated image/poem pairs to be painted on canvas with acrylics. The artist decided to mainly choose the images by their visual aesthetics (colours and patterns) and less by the associated poems. The artist knew that the images are representations of computational processes, but did not explicit generation method.

Before seeing the computer-generated abstract images, the artist imagined that they contain clear and monotonous lines, opposite to the actually generated images. This gave her the idea to use the computer images as inspiration to create paintings similar to the ones she had imagined before.

The artist then developed a painting technique that involved paper tape to ensure a clinical accuracy of the painted lines, in contrast to the more gradient transitions from one colour to another in the computer-generated images.

A photograph of a final painting can be seen on the right in Figure 1, next to the computer-generated original image on the left. The poem accompanying this image is:

I am gotten like a firm
Al-essawi and thither.
The ban about me
Serves itself into gun of total plans,
York and ordering,
Minimises stoped by the weapon.

Date: 21 December 2012

This was the first time for the artist to work in collaboration with a computer. In her opinion the process was inspiring and interesting. What made it different from her previous experience was that she had to work in a certain framework – provided by the computer, which led her to the idea of using more accurate lines than the computer.

Exhibitions

The work has been exhibited in three venues:

- June 5 - June 30, 2013, Art Museum of Tartu, Estonia
- August 1 - August 30, 2013, Culture Center of Jõgeva, Estonia
- October 24 - November 14, 2013 “Art Corridor” of the Exactum building of the University of Helsinki, Finland

The first exhibition got media coverage in Estonian national newspaper Eesti Ekspress¹ and also in local newspapers.

¹<http://ekspress.delfi.ee/archive/article.php?id=66524456>




Figure 1: A process visualization image generated by the computer (left) and a photograph of the corresponding acrylic painting (right). For the respective poem, see the text.



Figure 2: Image from the exhibition in Tartu, Estonia

During the exhibition in Helsinki we collected feedback from the audience. Beside the artworks, we placed feedback forms and a box for slipping them in. In order to make giving feedback easy, the feedback form contained two questions:

1. What do you think of this exhibition of "Arts, News & Poetry"?
2. Could you please circle a number below to give a score to the exhibition (1-worst, 5-best)?

The questions were designed this way for three reasons: 1) our goal was to keep people open to giving their ideas; 2) we wanted to avoid giving any sort of bias in any direction; 3) we found it to be more likely to get feedback if the forms are short and easy to fill in.

We received a total of 24 feedback forms from which we removed 4 of the forms which had unreadable gibberish or unrelated comments. 7 of the forms had comments which tended to be negative or sceptical, 10 forms had comments which could be considered positive, and 3 forms had general comments, e.g. "the exhibition raises interesting thoughts".

The positive comments tended to be longer than negative ones. One of the visitors proposed using the technique for encrypting messages. Interestingly, one person found a connection between the exhibition and the computer game Minecraft. One of the longer comments stated:

"Raises interesting thoughts about what art is. The poems and paintings are seemingly meaningless and will cause thoughts and feelings with the probability of a wall, forest or just about anything [sic!]. Yet there is artists experience involved. I do not perceive any interesting experience from the exhibition apart from these meta-thoughts. All in all, the exhibition feels like random data (which raises thoughts :)"

Some of the negative comments stated that the results is "just noise", or

"[The exhibition is] very boring, no artistic value, creative, maybe, but dull, monotonic and lacking depth. No serendipity!"

In total we got 19 scores from the feedback forms, with 3.13 as an average and slightly skewed towards positive end.

Related Work

Computer-Human Collaboration in Arts. Our project seems rather unique in the sense that it creates artwork (paintings, in human-computer collaboration) about another creative process (computational poetry generation). This could be classified as conceptual art, claiming that it is the idea and process that constitute the artwork, not alone the resulting paintings and poems. There are numerous conceptual works of art using computational or mechanistic generation of artefacts, and given the richness and variety of the field, we would be surprised if there are no others that take this to the metalevel like we have done.

Even though we are not aware of other art projects addressing exactly the same aspects, the general idea of artistic collaboration between computers and humans is of course not a new one. For instance, the *biomorphs* of Richard Dawkins (1986) have inspired at least Machado and Cardoso (2000) and Sims (1991). In their systems visual art is generated by genetic algorithms but at least partially guided by their users, so that the end result is a mix of computational creativity and human aesthetics.

Computational Creativity Theory. In the field of computational creativity research, a concept related to our work is that of *framing*, i.e., (computer-generated) commentary that adds value to the generated artwork e.g. by describing the underlying processes (Colton, Charnley, and Pease 2011). Process visualization could obviously be considered as a kind of framing for the poetry, providing an (abstract) image of the generation process. However, in our case, the roles are mutual: the paintings clearly take the role of the primary results, and the poems become part of the commentary for the paintings.

Process Visualization. We based our visual artwork on process visualization. An overview of different approaches to program visualization is given by Roman & Cox (1992). In general, the goal of program visualization is to take advantage of humans' high bandwidth of visual system and possibly give another way for people to analyse and understand algorithms (Roman and Cox 1992). They described

examining program's input and output actions and treating the program as a "black box" which transforms the inputs to the outputs as a method which has important theoretical implications, but is not very informative to get insight into the algorithms and is not applicable to all programs of interest.

The rest of the related work tends to be more practical in nature, for instance there is research in algorithm animation (Brown and Sedgewick 1984), visual programming (Myers 1990), and data structure visualizations (Hendrix, Cross II, and Barowski 2004).

Poetry Generation. The poetry generation methods of this work are based on the methods by Toivanen et al. (2012; 2014). A thorough review of different poetry generation methods is not in the scope of this paper as our emphasis here is the process visualization as a possible method of giving framing information, but, e.g. Colton et al. (2012) provide a good overview of the field.

Conclusions and Future Work

In this paper we have given an overview of an implemented and exhibited art project that combines both computational and human creativity in a rather novel way. We proposed an approach for extracting a visual abstraction of a process based on the input and the output of a system. We combined this together with a methodology for generating poems from a news story and used these pieces together for visualizing the abstraction of a process of generating respective poems. An artist then hand-picked some of the images and painted them in her chosen style. The paintings have been exhibited together with the associated poems.

There are many directions for future work. An interesting technical research problem would be developing more intelligent methods for extracting (aesthetic) abstractions of the process. In the best case, analysing such abstractions could be a way of getting insight into creative artefact generation.

An exciting creative possibility would be to make the process visualization interactive: allowing visual manipulation of the process matrix, and then repeating the creative process using the modified matrix to produce a modified poem as an output. If this approach works, it could unify verbal and visual arts in a most interesting way.

Acknowledgments

This work has been supported by the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 611733 (ConCreTe), the Academy of Finland (Algodan Centre of Excellence), and the Helsinki Doctoral Program in Computer Science (HECSE).

References

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.

Brown, M. H., and Sedgewick, R. 1984. *A system for algorithm animation*, volume 18. ACM.

Charnley, J.; Pease, A.; and Colton, S. 2012. On the notion of framing in computational creativity. In *Proceedings of the Third International Conference on Computational Creativity*.

Colton, S.; Charnley, J.; and Pease, A. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In *Proceedings of the 2nd International Conference on Computational Creativity*.

Colton, S.; Goodwin, J.; and Veale, T. 2012. Full FACE poetry generation. In *Proceedings of the Third International Conference on Computational Creativity*, 95–102.

Dawkins, R. 1986. The blind watchmaker: Why the evidence of evolution reveals a universe without design. *Penguin, London, UK. DE LA MARE WK (1997) Abrupt mid-twentieth-century decline in Antarctic sea-ice extent from whaling records. Nature* 389(6646):57–60.

Hendrix, T. D.; Cross II, J. H.; and Barowski, L. A. 2004. An extensible framework for providing dynamic data structure visualizations in a lightweight ide. *ACM SIGCSE Bulletin* 36(1):387–391.

Hotson, D., and Yencken, L. 2012. Extracting colors with colorific. Online; Last accessed 24-January-2014. <http://99designs.com/tech-blog/blog/2012/05/11/color-analysis/>.

Hunter, J. D. 2007. Matplotlib: A 2d graphics environment. *Computing In Science & Engineering* 9(3):90–95.

Machado, P., and Cardoso, A. 2000. Nevar—the assessment of an evolutionary art tool. In *Proceedings of the AISB00 Symposium on Creative & Cultural Aspects and Applications of AI & Cognitive Science, Birmingham, UK*, volume 456.

Myers, B. A. 1990. Taxonomies of visual programming and program visualization. *Journal of Visual Languages & Computing* 1(1):97–123.

Řehůřek, R., and Sojka, P. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50. Valletta, Malta: ELRA.

Roman, G.-C., and Cox, K. C. 1992. Program visualization: The art of mapping programs to pictures. In *Proceedings of the 14th International Conference on Software Engineering, ICSE '92*, 412–420. New York, NY, USA: ACM.

Sims, K. 1991. *Artificial evolution for computer graphics*, volume 25. ACM.

Toivanen, J. M.; Toivonen, H.; Valitutti, A.; and Gross, O. 2012. Corpus-based generation of content and form in poetry. In *International Conference on Computational Creativity*, 175–179.

Toivanen, J.; Gross, O.; and Toivonen, H. 2014. The officer is taller than you, who race yourself! Using document specific word associations in poetry generation. In *Proceedings of the 5th International Conference on Computational Creativity, Ljubljana, Slovenia*.