A Meme-Based Architecture for Modeling Creativity

Shinji Ogawa
Nagoya, Aichi, Japan
perfectworld@nyc.odn.ne.jp

Bipin Indurkhya and **Aleksander Byrski**
AGH University of Science and Technology, Cracow, Poland
{bipin, olekb}@agh.edu.pl

Abstract

This research is a collaborative work between a visual artist, a computer scientist, and a cognitive scientist, and focuses on the creative process involved in connecting two pictures by painting another picture in the middle. This technique was involved in four *Infinite Landscape* workshops conducted at Art Museums in Japan and Europe over the last five years. Based on the artist's verbal recollection of the ideas that occurred to him as he drew each of the connecting pictures, we identify the micro-processes underlying these ideas, and propose a meme-based, evolutionary-inspired architecture for modeling them.

Introduction

Research in recent years has revealed that though creativity may involve an *aha* moment with a gestalt shift or a sudden perceptual or conceptual reorganization, it is typically preceded and followed by several micro-processes that play an equally important role as the *aha* moment itself (Dunbar 1997; Sawyer 2006). These micro-processes can occur within a cognitive agent itself, or in different agents within a group or society. Our goal in this research is to study and model these micro-processes.

Infinite Landscape Workshops

This research is a collaborative effort between a visual artist [henceforth *the Artist*], a computer scientist and a cognitive scientist. Over the last five years, the Artist conducted four workshops at art museums in Japan and in Europe with the common theme *connecting different spaces*. In each workshop, there were 15-19 participants, all children (8-14 years) except in one workshop there were six adults. All the workshops followed the following modus operandi.

In the first step, the children were shown about 20 photographs of scenery from around the world, and then they were asked to draw imaginary landscapes using the building, people, animals etc. in these pictures as they liked. In the second step, the Artist brought the children's imaginary landscapes to his studio, and then he drew one picture to be inserted between every two children's pictures, so that all three pictures form a seamless scene. One such trio of pictures is shown in Fig. 1: scenes 9 and 10 were drawn by participants, and the Artist drew S9 to connect the two.

Figure 1: First strip

In the third and final step, all the pictures were connected in a ring without a beginning and an end, and the completed ring was suspended from the ceiling of the museum where the workshop was held. The ring was placed with the paintings on the inner side, so that the viewer is surrounded by the work while viewing it.

Overview of the Project and Methodology

Specifically, our goal in this project is to model the micro-processes involved in creating the connecting picture. Our methodology is as follows. In the first step, the Artist has recorded various ideas that occurred to him as he drew each of the connecting pictures. In the second step, we analyze these steps to identify and classify underlying processes. In the third step, we outline a model for implementing these processes. Finally, we would like to do experiments with the implemented system and evaluate the results.

In the current paper, we report our observations from analyzing the data from the workshop conducted at the Meguro Museum of Art, Tokyo (Japan) on 2 August 2007. The Meguro workshop was different from the other three workshops in that the participants were given only pencil and paper; there was no color, so the focus was on forms, shapes and space. Also, this workshop included six adults among nineteen participants; the remaining 13 were children (8-14) years. Based on our observations, we identify various micro-processes and how they interacted with each other to create the macro-level connecting pictures. Finally, we propose a meme-based, multi-agent architecture for modeling the underlying cognitive process, and discuss future research directions.

Observations on the 'Connecting' Process

We analyzed data from ten connecting pictures that the Artist drew for this workshop. Here we present the Artist's self-reflection on the genesis of ideas that led to the creation of connecting pictures. We include here seven of the more interesting cases. (The original comments were in Japanese. Translation and slight editing is by one of the other authors of this paper.)

We start with the Artist's observations on connecting 9 and 10 (Fig. 1): "These two had completely different atmosphere from each other. Sketch 9, drawn by an adult participant, is a scene set at dusk; a person looking at the artist is drawn wearing a sad expression. Sketch 10 has a bright atmosphere with flowers, fountains, buildings on a hill, and a horse. Moreover, each picture had an important character in the bottom left. The idea for connecting these sketches came to me while looking at the wonderful horse in 10. I thought of putting a parent horse running nearby. Because the background color of 9 and the body color of the horse in 10 was the same, I transformed the background of 9 into the parent horse in S9, which became a nested image structure. Then I extended the baby horse and the hill with the buildings."

On connecting 11 and 12 (Fig. 2): "There was the ground and the sky in the left one-third of 11, but the sea covered the remaining part on the right. In 12, a vast meadow was drawn with rich pictorial details. Here my attention was drawn to the connection between the color of the giant bridge in 11 and the color of the sky in 12. In S11 I drew the enlarged bridge of 11 and connected it with the picture on 12, which resulted in a nested image structure."

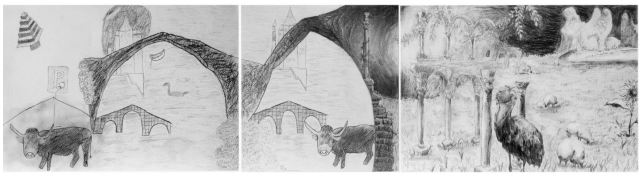

Figure 2: Second strip

On connecting 12 and 13 (Fig. 3): "I felt these two could not be connected with the techniques I had used so far. Then I noticed the wall on the top-right corner of 12 and the curved ledge surrounding the fountain in 13. Using these two curves, I drew a large Mobius strip in S12. As this Mobius strip divided S12 into four sections, in each section I extended the adjacent scenery. It felt like pouring in the scenery. Accordingly, I was able to connect them without blending, and this became the first work with this technique."

On connecting 7 and 8 (Fig. 4): "Because 8 was a richly detailed realistic presentation, to contrast it with the presentation in 7, I decided to stress dimensionality in the connection. The realistic rocks and the bridge in 8 were rendered in 3-d and were connected with the bridge in 7 that was extended in 2-d. To make this connection smoother and give an accent to the picture, I drew 3 Russian onion domes from 7 into S7."

Figure 3: Third strip

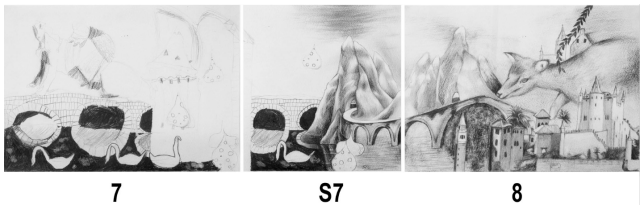

Figure 4: Fourth strip

On connecting 5 and 6 (Fig. 5): "Sketches 5 and 6 could be naturally connected. However, I had decided to refuse ordinary, conventional way of connecting things. I got the hint from the composition of 6. Oddly, on the right of 6, everything is drawn tilted towards the bottom left along a vector, but in the middle part, another horizontal vector appears. As a result, the horizon is split into two: one horizontal and another pointing to bottom left. I further emphasized this split of horizon, and drew a horizon pointing to the sky where the cow is, and another horizon that is sinking down where the buildings are."

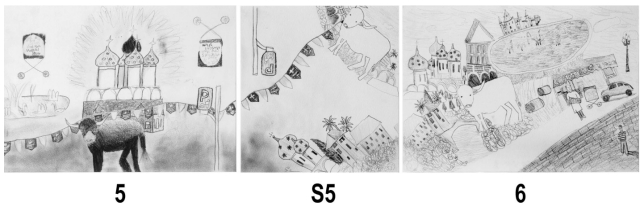

Figure 5: Fifth strip

On connecting 4 and 5 (Fig. 6): "Connect 5 on the right of 4. I was very interested in the row of flags that was hanging in 5 from left to right. On the right edge of 4, there is an upside-down building. What a challenge! I took that challenge and extended the gate of that fort-like building, and turned the top-right part of it into water surface. I extended that dark water surface to the right, making it narrower, and connected it with the contour of the lake in 5. On top of it, I placed the swans and plants from 5. I left the top-right part of the picture white in order create a contrast effect with the black space that is extended to the left. In the bottom right, I extended the flags."

On connecting 3 and 4 (Fig. 7): "I had a strong impression that the participants were expressing their own images instead of sketching by sampling from the photographs of the scenery I had shown. An extreme case of this is 4.

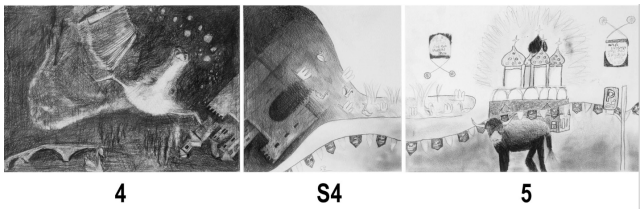

Figure 6: Sixth strip

At first the sketch was filled-in completely black, and then brightened by the eraser. It had no earth and sky, but an ambiguous space from a dark fantasy I decided to connect this dark picture with 3, which had a child-like pictorial space. However, it would be impossible to connect the two in an ordinary way. Here, I decided to ignore all the meaning in these pictures, but to focus on the pattern of light and dark instead. I said to myself, 'it is just a blotch'. The only connecting point in both pictures was the street in 3 and the bridge on the bottom left of 4. I could connect this street and the bridge. Luckily, bottom left of 4 looked like the sea, and bottom right of 3 also looked like a body of water. In S3, I extended the road in 3 in S-shaped curve and connected it with the bridge in 4. Continuing, I also extended the sea The problem was what to do on top of this. On the left part of S3, the only possibility was to extend the street-side houses on 3, so I did that in the same touch. Then I gradually changed the color of houses from gray to black, while introducing spatial distortion, and changing them from solid to liquid. I floated a swan in the dark pond that the buildings were turned into."

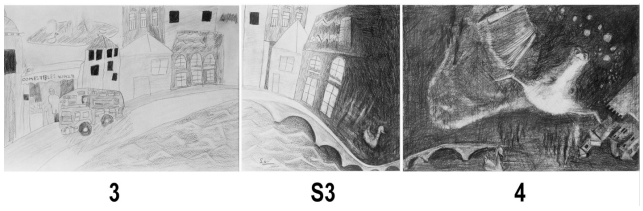

Figure 7: Seventh strip

Identifying Micro-Processes in 'Connecting'

Carefully going through all these comments, as well as examining the trio of pictures ourselves, we came with the following list of micro-processes that played a role in the genesis of connecting two pictures:

Copy elements This was by far the most common operation. Elements were copied from both the left and right pictures and incorporated in the connecting pictures just like that. One can see examples of this in almost every instance of connection. Among the examples presented above, one can see that swan is copied from 4 to S3, flags, plants and swan from 5 to S4, onion domes and swan from 7 to S7, small bridge and bull from 11 to S11, and so on.

Copy elements and transform This is similar to the above except that the element gets transformed while copying. For example, the rocky peak and the bridge are rendered in 3-d while they are copied from 8 into S7, and parking sign is turned around as it is copied from 5 into S5.

Copy elements and swap attributes Here elements that are being copied interact during copying and swap attributes. One example is provided in 6-S6-7 (Fig. 8), where two people are copied from 6 into S6, but their poses and the object one of them is holding are swapped.

Extend elements An element is continued in the adjacent picture; for example, the sea from 3 into S3, the masonry from 6 into S6, and the meadow from 12 into S12.

Same form (shape, shade,...) → search for meaning This is illustrated by 9-S9-10 (Fig. 1), where the same shading for the horse's body in 10 and the background in 9 led to the idea that the background in 9 can be morphed in the mother horse in S9. This process can also be evidenced between S11 and 12 (Fig. 2).

Similar form and semantic association → morph forms This is evidenced in 3-S3-4 (Fig. 7), where a semantic association between the road and the bridge, and similar forms (notice that they are similar but not the same) led to the idea that they can be joined by morphing one into the other.

Form-based continuation This is different from the extend element above in that the continuation is based on the shape and shade only, and does not involve meaning. This is seen in S3 and 4 (Fig. 7).

Form-contrast → concept-contrast This is illustrated by 7-S7-8 (Fig. 4). The contrast between a richly detailed sketch (8) and a plain sketch (7) suggested a 3-d vs. 2-d contrast.

Form-similarity → unifying concept In 12-S12-13 (Fig. 3) form-similarity between the wall on the top right of 12 and the ledge around the fountain on the bottom left of 13 suggested the idea of a Möbius strip.

Emphasize concept In 5-S5-6 (Fig. 5), different planes (horizons) in 6 were incorporated in S5 and emphasized. This is similar to *copy element* and *transform* except that the element is a concept rather then a concrete object.

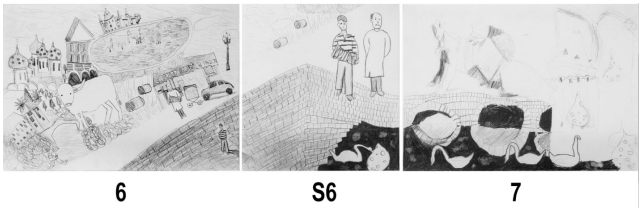

Figure 8: Eight strip

Meme: A Representation for Ideas

In order to represent all these micro-processes, we propose to use the formalism of *meme,* which was popular-

ized by Richard Dawkins in his celebrated *The Selfish Gene* (Dawkins 1989). Memes are cultural counterpart of genes, and represent ideas that can be generated, be passed on, get transformed, be combined with each other, and die out. As we observed many of the similar operations and interactions among the micro-processes in connecting two pictures, we chose meme as a unit of representation for modeling.

In our particular domain, a meme can be an element like a swan, a horse, or a building. It is a particular element, so it carries specific attributes. In other words, the horse meme that plays a key role in 9-S9-10 (Fig. 1) is not the general concept of a horse, but carries concrete attributes like the shade and the shape of the horse that was drawn in 10. There can also be conceptual memes, for example, 'horizon tilted to top-right', '3-d rendering', or 'dark shade'. Such memes represent specific operations or attributes that can be imparted to an element or a scene.

It is possible to have generalized memes and to organize them in a hierarchy. So, for example, there can be a 'horse' meme of which the horse meme of 10 (Fig. 1) would be an instance; or there can be a 'tilted horizon' meme, which would be a parent of the 'horizon tilted to the top-right' meme. But for the time being we are not considering such general memes.

Following actions can be carried out on individual memes:

Copy or replicate In this case the element is copied as it is, or the concept is applied as it is. So a swan is copied with all its attributes intact, or the horizon can be tilted toward the top-right corner of the pictures for a part of the scene that is selected.

Copy with transformation In this case, the element is copied but one or more of its attributes are changed along the same dimension. For example, its size can be made bigger or smaller, its color or shade can be changed, its orientation can be changed, and so on. For a concept meme, some of its parameters are changed during application; for example 'horizon tilted to top-right' can change to 'horizon tilted to top-left'.

Two memes can also interact with each other and we specify the following four modes of interaction:

Swap attributes Two memes can swap attributes of each other. We saw an example of this in Fig. 8 where the pose and the "object-held' attributes of two people were exchanged.

Overwrite attribute In this case the attribute of one meme overwrites the attribute of the other meme. So, for example, the size or the color of one meme can be rendered according to the other meme. This is illustrated by an instance in the Osaka workshop, where the silhouettes of cliffs were made to conform to the silhouettes of buildings.

Unify This allows two memes to bond together and act as one meme. Any common attributes of the two become the attributes of the unified meme, and in addition some extra attributes may be created based on the spatial or other relationships between the two. This is similar to the grouping operation in many graphic editors.

Create a new meme This allows creation of a new meme with attributes inherited from each of the parent memes.

There are a number of other features that we are not considering at the moment. For example, it may be possible for a meme to activate another meme. We saw an example of it in our observation above when the Mobius strip idea was suggested by the similarity in form between the wall and the fountain ledge 3. However, in order to model this mechanism, we need to have some kind of global associative knowledge network.

A Memetic Architecture

We are implementing a meme-based system to model the process of creating the intermediate picture. In particular, our system incorporates the following features: 1) modeling of visual attention to identify prominent elements or areas in the neighboring pictures; 2) specifying memes for spatial relationships among the picture elements; 3) specifying memes for general techniques like extension and continuation; and 4) various heuristics for choosing among competing memes.

For a lack of space, and also as our system is currently being implemented, we limit ourselves to only pointing out that we are exploring two approaches to generating the connecting image for a given pair of images:

- Evolutionary algorithm: the two images should be digitalized, and potential solution generated stochastically from them with the use of crossover and mutation (Michalewicz 1998). Formulation of the fitness function should take into consideration the similarity of the potential solution to both of images. We also plan to incorporate aesthetic criteria in the fitness function (Norton, Heath, and Ventura 2010).
- Agent-based approach: some complex approaches utilizing multi-agent notions (Byrski and Kisiel-Dorohinicki 2005)) bring interesting additions to the process, as autonomous individuals, as agents are, may utilize other means to evaluate the resulting images, and may choose different crossover and mutation operators in an intelligent way to apply them to the current solution.

Both approaches may leverage concepts well-known from the memetic computation—local search (Moscato and Cotta 2010)—thus applying a number of mutation operators (instead of only one) before final evaluation.

Relation with Previous Research

Needless to say, the ideas and the architecture presented here are based on a number of existing and past research efforts to model different aspects of creativity. The origin of the parallel, competition-cooperation architecture can be traced back to Selfridge (Selfridge 1959). Subsequently, Lesser et al. (Lesser, Fennell, and Reddy 1975) formalized it as *blackboard* architecture and used it for speech recognition; and in our earlier research (Indurkhya 1997) we used a similar approach to model creativity in legal reasoning. Hofstadter

and his colleagues (Hofstadter 1995) proposed a *parallel terraced scan* architecture for modeling creativity in analogical thinking and our approach outlined above is heavily influenced by their work. One key point of difference is that a meme is more like an agent that carries its own data with it, unlike a knowledge source in the blackboard architecture or a codelet in Hofstadter's architecture.

The system proposed above also draws from the *meme media* architecture of Tanaka, Fujima and Kuwahara (Tanaka and Kuwahara 2008). They have developed the C3W wrapper framework that allows the user to open a web application page, clip some input and output portions as pads, and link them with pads clipped from other web applications.

A number of approaches have been developed for applying evolutionary algorithms to generate visual art (Sims 1991; Lewis 2007; Machado, Romero, and Manaris 2007), but their goal is to generate aesthetically pleasing visual objects. In the long run, it may be possible to use some of these techniques by incorporating constraints from the neighboring picture objects to generate novel but related picture objects for the connecting picture.

As for systems that generate constrained visual objects or scenes, there has been some research on automatic collage generation (Krzeczkowska 2009) and on completing a partially drawn picture in the intended style (Colton 2008), and some of the techniques developed therein can be exploited in our system as well.

Conclusions and Future Research

We analyzed data from the Artist's verbal recollection of his thoughts as he drew the middle pictures to connect pairs of pictures seamlessly. From this analysis, we identified a number of micro-processes that led to the big picture idea. We described a memetic approach to formalize these micro-processes, and outlined an evolutionary-inspired approach to support the process of generating the connecting picture.

We are also interested to study the cognitive processes of the viewers as they look at the trio of pictures. It has been noted in the past that surface-level perceptual similarities influence how viewers connect pairs of images and relate them conceptually (Indurkhya et al. 2008). It would be interesting to see how this process is affected when there is an intervening picture in the middle. We plan to conduct behavioral and eye-tracking experiments to measure the viewers' response and incorporate those observations in our model.

References

Byrski, A., and Kisiel-Dorohinicki, M. 2005. Immunological selection mechanism in agent-based evolutionary computation. In Klopotek, M. A.; Wierzchon, S. T.; and Trojanowski, K., eds., *Intelligent Information Processing and Web Mining : proceedings of the international IIS: IIPWM '05 conference : Gdansk, Poland*, Advances in Soft Computing, 411–415. Springer Verlag.

Colton, S. 2008. Experiments in constraint-based automated scene generation. In *Proc. of the 5th International Joint Workshop on Computational Creativity*.

Dawkins, R. 1989. *The Selfish Gene (2nd ed.)*. Oxford University Press.

Dunbar, K. 1997. How scientists think: On-line creativity and conceptual change in science. In Ward, T.; Smith, S.; and Vaid, J., eds., *Creative thought: An investigation of conceptual structures and processes*. American Psychological Association.

Hofstadter, D. 1995. *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. New York: Basic Books.

Indurkhya, B.; Kattalay, K.; Ojha, A.; and Tandon, P. 2008. Experiments with a creativity-support system based on perceptual similarity. In Fujita, H., and Zualkernan, I., eds., *New Trends in Software Methodologies, Tools and Techniques*. IOS Press: Amsterdam. 316–327.

Indurkhya, B. 1997. On modeling creativity in legal reasoning. In *Proceedings of the Sixth International Conference on AI and Law*, 180–189.

Krzeczkowska, A. 2009. Automated collage generation from text. Master's thesis, Imperial College, London.

Lesser, V.; Fennell, R.D. nad Erman, L.; and Reddy, D. 1975. Organization of the hearsay-ii speech understanding system. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 23(1):11–24.

Lewis, M. 2007. Evolutionary visual art and design. In Romero, and Machado, P., eds., *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, 3–37. Berlin: Springer-Verlag.

Machado, P.; Romero, J.; and Manaris, B. 2007. Experiments in computational aesthetics: An iterative approach to stylistic change in evolutionary art. In Romero, and Machado, P., eds., *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, 381–415. Berlin: Springer-Verlag.

Michalewicz, Z. 1998. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer.

Moscato, P., and Cotta, C. 2010. A modern introduction to memetic algorithms. In Gendrau, M., and Potvin, J.-Y., eds., *Handbook of Metaheuristics*, volume 146 of *International Series in Operations Research and Management Science*. Springer, 2 edition. 141–183.

Norton, D.; Heath, D.; and Ventura, D. 2010. Establishing appreciation in a creative system. In *Proceedings of the International Conference on Computational Creativity: ICCC-X,*, 26–35.

Sawyer, K. 2006. *Explaining Creativity*. Oxford University Press.

Selfridge, O. 1959. Pandemonium: A paradigm for learning. In Blake, D., and Uttley, A., eds., *Proc. of the Symposium on Mechanisation of Thought Processes*, 511–529. H.M. Stationery Office: London.

Sims, K. 1991. Artificial evolution for computer graphics. *ACM Computer Graphics* 25:319–328.

Tanaka, Y., F. J., and Kuwahara, M. 2008. Meme media and knowledge federation. In Dengel, A., and et al., eds., *Proc. of KI2008: LNAI 5243*, 2–21. Springer Verlag.