Artificial Intelligence for Engineering Design, Analysis and Manufacturing (2011), *25*, 1–13.

doi:10.1017/S0890060409990199

Evaluating an evolutionary method of design style imitation

ANDRÉS GÓMEZ DE SILVA GARZA AND ARÁM ZAMORA LORES
Instituto Tecnológico Autónomo de México, Mexico City, Mexico

(Received October 8, 2008; Accepted May 19, 2009)

Abstract

We propose a computational method for producing novel constructs that fall within an existing design or artistic style. The method is based on evolutionary algorithms, and we discuss related knowledge representation issues. We then present an implementation of this method that we used in order to imitate the style of the Dutch painter Mondrian. Finally, we explain and give the results of a cognitive experiment designed to determine the effectiveness of the method, and provide a discussion of these results.

Keywords: Design Representations; Design Styles; Evolutionary Algorithms

1. INTRODUCTION

Using computers to generate artwork is not a new idea, and neither is using ideas from evolutionary algorithms (Mitchell, 1998) for this task. Some examples of design systems that are inspired in evolution are presented and discussed, for example, in Bentley (1999) and Bentley and Corne (2002). Most of the systems that generate artwork that are described in these two books (e.g., Rowbottom, 1999; Todd & Latham, 1999; Witbrock & Neil-Reilley, 1999; Eiben et al., 2002; Hancock & Frowd, 2002; Pagliarini & Lund, 2002; Rooke, 2002) do so by using the evolutionary operators of crossover and mutation to propose new paintings. However, all of these systems leave it to the users to decide, during the evaluation phase of the evolutionary algorithm, which of the new paintings, or which of the features of the new paintings, to keep for future evolutionary cycles, and/or how to rank the new paintings according to whatever the users' subjective, and probably unconscious, aesthetic criteria might be. Thus, the decisions on what is aesthetic or interesting are not made by the systems, and therefore one of the benefits of evolutionary algorithms (iterating rapidly and autonomously through a propose–evaluate–discard cycle) is not taken advantage of.

The above approach assumes that one is interested in producing new artwork, with the aid of the computer, that is deemed by people to be pleasing, and that there are enough people available to provide feedback to the evolutionary process to generate such artwork. In contrast, we are more interested in formalizing and furthering our understanding of design, and in the computational and algorithmic aspects of artwork generation rather than the artistic value of the final product. We would like to create fully autonomous systems that require no user feedback as their evolutionary or other algorithms proceed. The point of our approach is to explore, and perhaps push, the limits of what computers are capable of doing by themselves.

More specifically, what we are interested in is to have computer systems that can imitate an artistic style given examples of that style. This contrasts with systems that are programmed with their own, completely new, artistic style, for instance, by generating paintings that follow the patterns defined by some preprogrammed mathematical equations, or that, because of the way they are programmed, emulate the artistic style of their users or programmers, such as Harold Cohen's AARON (McCorduck, 1991). Imitating a style involves achieving a trade-off between making sure that the new designs or paintings that are created are novel, rather than simply copying previously existing ones, but at the same time making sure that the new products do not differ in significant ways from the original ones in order not to go beyond the limits of the style that is being imitated.

For various reasons given shortly, we believe that evolutionary algorithms provide us with a computational framework that allows us to achieve our goal. Work into capturing the style of particular artists or designers in the computer has often focused on shape grammars (e.g., Cha & Gero, 1999) or semantic networks (e.g., Gero & Jupp, 2003), although some work has also used evolutionary algorithms to explore style. For instance, Ding and Gero (2001) describe a system that

Reprint requests to: Andrés Gómez de Silva Garza, Río Hondo #1, Colonia Progreso-Tizapán, Mexico City 01080, Mexico. E-mail: agomez@itam.mx

generates traditional Chinese architectural facades after it infers a representation of their style. This inference and subsequent learning is done by using an evolutionary algorithm that produces a hierarchical genotype that represents a particular style according to the exemplars that have been shown to it. Recognizing whether a new design matches a particular style involves matching its features with the style representation embodied in the hierarchical genotype. However, these new designs are not generated by this system by using evolutionary algorithms.

In contrast, in our work we believe that the power of the evolutionary algorithm is in the generation of new artwork. In Section 2 we describe our method for generating imitations of style based on an evolutionary algorithm, and explain our reasoning behind this choice of computational method. In Section 3 we present a particular artistic style, Mondrian's, which we tried to imitate by implementing our process model, and some knowledge representation issues related to capturing the essence of Mondrian paintings in the computer. In Section 4 we explain a set of experiments we performed for evaluating our implementation and give the results of these experiments. Finally, in Section 5 we discuss our findings and future work to be done.

One project that has in the past been used to learn (evolve) generic genotype descriptions of patterns used by Mondrian in his paintings, using a method similar to Ding and Gero (2001), is described in Schnier and Gero (1997, 1998). The new evolved descriptions were then used in an evolutionary algorithm to generate new paintings that use some of the learned patterns. However, the resulting paintings were never evaluated to determine their "Mondrian-ness" either computationally or by showing them to human subjects or experts. The purpose of the project was to explore learning of patterns in evolved genotypes, and also style combination (some of the learned patterns were combined in one evolutionary algorithm with patterns that were learned from Frank Lloyd Wright window designs), rather than style imitation. Even though the resulting paintings shown in the two papers cited above are clearly inspired by Mondrian's style, many of them would not be mistaken for Mondrian paintings by anyone familiar with his paintings. In addition, the representation used in that project forces the entire space on either side of a black line in a painting to be of a uniform given color, but many Mondrian paintings have colored areas that are not bounded on all four sides by a black line (or the edge of a canvas). This can be seen in some of the examples shown in supplementary Figures S.1 and S.2 (online only)[1] and in our project because we have taken care to be flexible enough to allow it in one of the rules (rule 6, Subsection 3.2) that embody our description of Mondrian's style. Therefore, despite some superficial similarities (the use of Mondrian exemplars and the use of evolutionary algorithms), there are several significant differences between that project and ours, in both purposes and implementations.

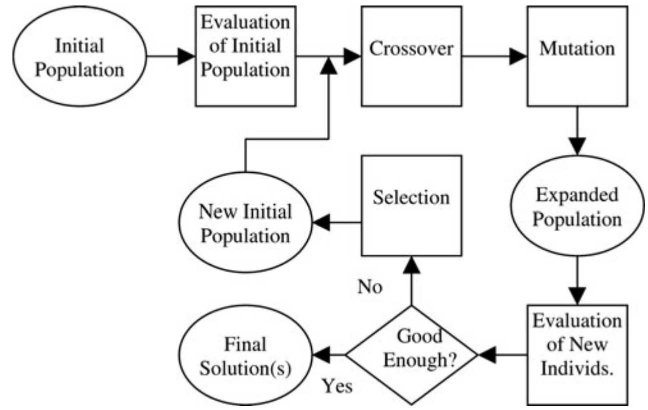

Fig. 1. The evolutionary method of style imitation.

2. EVOLUTIONARY METHOD OF STYLE IMITATION

In this section we describe in generic terms our evolutionary method of style imitation, some of its features, and the reasons behind them, leaving out the details and decisions related to the method that depend on a particular application to the next section of the paper, in which we describe our implementation domain. Figure 1 shows the flow of tasks in our evolutionary method of style imitation.

2.1. Evolution

Briefly, a population of potential solutions (e.g., paintings, in case we want to imitate an artistic style) is kept throughout the process. The makeup of the population changes because of the evolutionary process, in which new potential solutions are generated through the application of the genetic operators of crossover and mutation, which at random combine and modify the features of old potential solutions that are already in the population, respectively. This is in contrast with the use of shape grammars or other approaches that generate new potential solutions based on preprogrammed rules that embody expert knowledge about a particular design domain.

Returning to our process model, a temporarily expanded population is created by adding the new potential solutions to the original population. Each new potential solution is evaluated and a fitness value assigned to it. The fitness value is a measure of how close or how far each potential solution is from the style that we are trying to imitate. If one or more (depending on what is desired) of the new potential solutions is already perfect (i.e., already fits the desired style) according to the evaluation procedure, the evolutionary process stops (or if the search has gone on for too long without successfully producing any imitations). Otherwise, a selection procedure sorts the potential solutions in the temporarily expanded population according to fitness value, keeps the best of them,

[1] At the request of the Mondrian Trust, the printed version of this article does not include the images of Mondrian's paintings. However, they do appear in color as supplementary figures in the online version of the article (http://journals.cambridge.org/aie).

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

and discards the rest. The individuals that are kept become the initial population for the next evolutionary cycle. This new population may include both old and new potential solutions (i.e., some carried over from previous generations and some newly generated ones), something known as *elitism* in the evolutionary algorithms community, and the process repeats itself.

2.2. Crossover and mutation

As many people have remarked (e.g., Boden, 2003), creative solutions generated by people are hardly ever completely novel with respect to what has come before; they tend to combine and/or tweak certain aspects of previously existing solutions in a creative way. Although there is a general agreement that, in retrospect, the features of new solutions can often be traced back to previous solutions that were combined and/or modified, there is less certainty about how to achieve these two types of transformation of known solutions to achieve new, creative ones. One way in which these two desired effects, combination and tweaking or modification, can be directly achieved algorithmically is by using the two previously mentioned genetic operators of crossover and mutation, respectively.

Crossover and mutation both operate on linear data structures that are considered to be genotypes (i.e., a linear description of an individual instance of the type of species that is being processed by an evolutionary algorithm). The most basic form of crossover takes two parent genotypes as input, randomly chooses a crossover point for them, splits them into two pieces at that point, and produces two offspring genotypes as output by splicing together the opposite pieces of the two parents. That is, one offspring consists of the first part of the first parent concatenated with the second part of the second parent, whereas the other offspring consists of the first part of the second parent concatenated with the second part of the first parent. Figure 2 illustrates the functioning of the crossover operation graphically.

The most basic form of mutation takes one parent genotype as input and changes one of its genetic characteristics by randomly deciding which one to change and randomly choosing a new value for it, and produces the resulting offspring genotype as output. Figure 3 illustrates the functioning of the mutation operator graphically. Note that both crossover and mutation operate "blindly," without paying attention to the semantics of the genotypes they operate on, and thus are very generic operators that do not depend on the implementation domain, that is, on the type of individual that is described by the genotypes in an evolutionary algorithm's population. The individuals in the population may represent paintings, car designs, airplane wing profiles, or footwear sole cross sections, just to give a few examples of artistic or engineering domains for which an evolutionary method of design imitation like ours may be used.

The number of times that crossover and mutation are to be performed in each evolutionary cycle depend on the number of new individuals that need to be produced at each generation and on the relative percentage of the desired origin for these new individuals (i.e., how many of them we want to originate from a combination operation and how many from a modification operation). In addition, in some domains it may be necessary or useful to have the evolutionary algorithm manipulate nonlinear genotypes, such as those embodied in hierarchical (tree-shaped) data structures or matrix-shaped data structures. The corresponding crossover and mutation operators, while retaining their effects of combining or modifying aspects of previous solutions, respectively, would have to be implemented in ways that take into account the nonlinear nature of the genotypes, such as those explained in Koza (1992) and Kane and Schoenauer (1996). These are all decisions that may vary depending on the implementation domain, and because our evolutionary method is generic, it does not make any commitments in this regard.

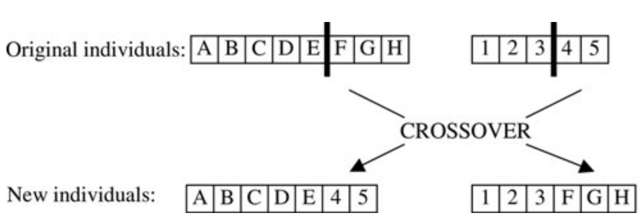

Fig. 2. A graphical illustration of the genetic crossover operator.

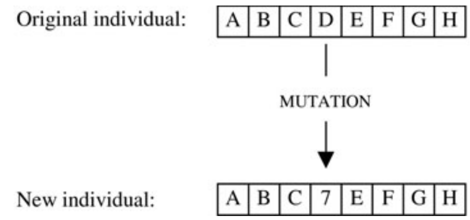

Fig. 3. A graphical illustration of the genetic mutation operator.

2.3. Makeup of the population and evaluation

When we analyze Figure 1 to determine what it is that is being combined and/or modified in our process model, we note that it is a population (of potential solutions) that is operated on by crossover and mutation. However, what exactly does this population consist of? Eventually over time, the population that is processed by the evolutionary algorithm consists at least partly of potential designs or paintings, represented as genotypes, that were generated by the evolutionary algorithm in previous generations. However, what is the initial contents of the population? If we want to combine and modify aspects of previously existing solutions, this would seem to indicate that the initial population should consist of known solutions. Fortunately, to achieve our task of style imitation we already have known solutions to draw from: exemplars of the style that we would like to imitate. For other design-related tasks, depending on how creative one wants to be, it might not always be so easy to have access to previously known solutions as for style imitation. In contrast, if the initial population consists exclusively of exemplars, then producing any new solutions that do not consist solely (and boringly) of some combination of the features of the already existing solutions would take a very long time and would depend on the efficiency and effectiveness of the mutation operator's use of randomness. Thus, in addition to exemplars, we recommend placing some randomly generated potential solutions into the initial population to add some variety into the population, a popular trait used in most evolutionary algorithms. This raises a new question: how much of each (exemplars and random "solutions") to use in the initial population. In Section 3 we discuss this issue with respect to our implementation domain.

The other aspect that needs to be discussed is evaluation. As can be seen in Figure 1, our process model involves the evaluation of the solutions in the initial population and of each new potential solution generated and placed into the population during the evolutionary algorithm. There is also the matter of selecting the best individuals in the population before starting each new evolutionary cycle, which depends on the results of the aforementioned evaluation. The following points have to be taken into account with respect to evaluation:

1. In the initial population those individuals that are real exemplars (as opposed to randomly generated solutions) of the style to be imitated should be awarded a fitness of 100% by the evaluation procedure, because they all belong to (fit within) the desired style.
2. The randomly generated initial individuals and the potential solutions that are newly generated during each evolutionary cycle should be awarded a degree of fitness that depends on how closely they fit within the style that is being imitated.
3. The evolutionary process has successfully imitated the style (i.e., reached convergence) when an individual in the population that was not there initially gets awarded a fitness of 100%.

Because of this there are some major, although perhaps subtle, differences between our process model and standard evolutionary algorithms, which generally begin with a completely random initial population, none of whose individuals have a 100% fitness, or anywhere close to that, and where convergence usually takes place the first time that any individual's fitness reaches or approaches 100%.

How to determine the degree of fit of a given design to a given style is the key factor in the evaluation procedure. We propose that the same exemplars that are used to seed the initial population of the evolutionary algorithm can be used to generate a generic description of the style that is being imitated, which can then be used during the evaluation phase of the evolutionary algorithm. Whether this generic description is obtained through a process of knowledge engineering, data mining, training a neural network to recognize the style, or some other means depends on the appropriateness of each of these possible methods for each design domain. In the following section we also discuss this issue with respect to our implementation domain.

3. IMPLEMENTATION DOMAIN AND KNOWLEDGE REPRESENTATION

3.1. Mondrian

The Dutch painter Piet Mondrian was born in 1872 and died in 1944. He was active mainly in the first half of the 20th century and, like many other modern painters, started his career painting landscapes, human figures, and other realistic subjects. Eventually, however, he developed his own distinctive and abstract style. Many people call Mondrian's style simply *de stijl*, which is Dutch for "the style." However, *De Stijl* was the name of a journal that was founded by Mondrian and Theo van Doesberg, to which Mondrian contributed many articles (and the artists attracted to it were called the De Stijl movement). Mondrian named his style *de neue Beelding* (in English, Neoplasticism). Paintings in Mondrian's style typically include vertical and horizontal black lines painted over a white background, with some or all of the primary colors (blue, red, and yellow) and/or black filling in some of the square or rectangular regions (or parts of the regions) that are separated from the background by the black lines. To imitate this style, we implemented a system that follows our evolutionary method.

In the Mondrian-Imitating Computer Artist (MONICA) System, which is written in C++ and for the graphical aspects OpenGL, we were able to store 55 paintings by Mondrian that we use as exemplars of his style. These 55 paintings do not include his early nonabstract work, his lozenges (following the style described above, but painted on diamond- rather than rectangle-shaped canvases), or his later work (in which he started to use colored lines, sometimes even multicolored lines, rather than just black lines, to separate the white or colored regions in his paintings). Supplementary Figures S.1 and S.2 show some of the 55 exemplars we used in MONICA, which were obtained from several Websites or scanned from one of two books: Deicher (1999) or Bax (2001).

In MONICA we chose to duplicate the size of the evolutionary algorithm's population temporarily (i.e., when creating the temporary expanded population shown in Fig. 1) at each evolutionary cycle, and of the new genotypes created in each generation, 80% are created through the use of the crossover operator and the other 20% through mutation. Through a series of experiments described in the following paragraph, we determined that, for our domain and representation, the optimal mixture of individuals in the initial population of the evolutionary algorithm to ensure quicker convergence is to use 60% exemplars and 40% randomly generated individuals. Therefore, MONICA's evolutionary algorithm operates on populations of 92 individuals, which consist initially of our 55 exemplars plus 37 randomly generated ones. The 60/40 split was determined as follows.

We initially knew that we wanted to combine exemplars and randomly generated paintings in the initial population, for the reasons given in Subsection 2.3 above, but did not know in which proportion, or if there would be any difference in outcome depending on the relative proportions of each. We therefore set up an experiment to test out different proportions. For this experiment we limited the total size of the population to 55. We first tried out the evolutionary algorithm using a population consisting exclusively of exemplars and without random paintings in a 55/0 or 100%/0% split. We then tried out the evolutionary algorithm using a 90%/10% split, an 80%/20% split, and so forth until 10%/90%. We ran each of these different splits 10 times, each time measuring the convergence time in CPU time, where convergence was determined to occur the first time a new painting fitting Mondrian's style, according to the evaluation rules discussed below, was produced. Figure 4 shows the results of this experiment.

As can be seen in Figure 4, for our particular knowledge representation and domain, and using speed of convergence as the criterion for comparison, it seems that any percentage higher than or equal to 30% of exemplars in the initial population is best, with 60% being up to 30% faster than some of the other possibilities. Our interpretation of this result is that if there are not enough exemplars in the initial population to give the system a head start, it may be a long time before Mondrian-like paintings start resulting from the evolutionary algorithm's operations. For our system this cutoff point, below which there are not enough exemplars in the population to guide the algorithm to a quick convergence, seems to be somewhere between 20% and 30%. The slight deterioration in convergence speed when the percentage of exemplars is above 60% may be due to the diversity factor: having less than 40% random individuals in the initial population might not allow the algorithm to produce new Mondrian-like paintings as quickly as when there is a larger amount of randomness, and therefore diversity, in the initial population.

For the experiments, the random solutions included in the initial population were created by randomly deciding how many colored regions a solution is going to have, and randomly choosing values for the color, dimensions, and positions of each one of them. More details on the experiments can be consulted in Gómez de Silva Garza and Zamora Lores (2005).

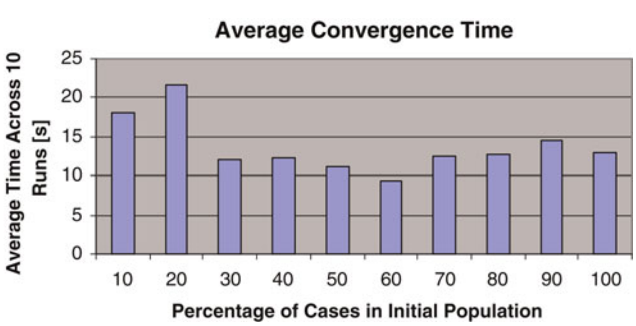

Fig. 4. The amount of CPU time for different combinations of exemplars (cases) and random individuals in the evolutionary algorithm's initial population. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

3.2. Evaluation rules

For the Mondrian domain we figured it would be relatively simple to come up with a series of evaluation rules based on our own observations of the patterns that we found to be present in the 55 exemplars. The series of evaluation rules we implemented pay attention to such factors as the total number of vertical or horizontal lines present in a painting, the total number of colored regions present in a painting, the locations of the colored regions with respect to the black lines and/or the edge of the canvas, and the thickness of the black lines. Judging from the 55 exemplars, it seems to us that Mondrian was also taking into account these constraints while producing new paintings in his style, at least subconsciously. The implemented rules that are used to evaluate the new individuals produced by our evolutionary algorithm (plus those in the initial population) are the following:

1. EvaluateColor: Each rectangular region that is contained in an individual must have one of the four valid colors (red, green, blue, black).
2. EvaluateCoordinates: The height, width, x coordinate, and y coordinate of each rectangular region in an individual must all fall between within the limits admitted by OpenGL.
3. EvaluateLineThickness: Up to two black rectangular regions are allowed in an individual that are not thin, but all other black regions in the individual must be either vertically or horizontally thin and thus must visually represent a line rather than a rectangular region. The definition of "thin" used here is having a width of at most 2 pixels.
4. EvaluateNumberOfVerticalLines: A minimum of 1 and a maximum of 10 vertical lines must be present in an individual.
5. EvaluateNumberOfHorizontalLines: A minimum of 2 and a maximum of 10 horizontal lines must be present in an individual.
6. EvaluateLimits: Each colored rectangular region in an individual must be adjacent vertically, horizontally, or both to rectangular regions that represent lines or to the edge of the canvas. A vertical adjacency here means that the colored region must touch, both above and below it, either a horizontal line or the top or bottom edge of the canvas. A horizontal adjacency means that the colored region must touch, both to its left and to its right, either a vertical line or the left or right edge of the canvas.
7. EvaluateFrame: All rectangular regions in an individual must fall within the coordinates of the frame/canvas, whose background color is white by default, whose center represents the origin, and whose dimensions are such that +4.0 and –4.0 are the maximum and minimum x- and y-coordinate values, as per OpenGL specifications.

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

Fig. 5. The subdivision of a painting into 20 parts, one for each possible colored region in a Mondrian painting.

8. EvaluateNumberOfColoredRegions: There must be at least 1 colored rectangular region in an individual, and at most 13, not counting the thin, black "rectangular regions" that are actually lines, but at most 20 in total if we do count the lines.

One could wonder whether a different set of rules may have arisen if a different set of people other than us had been the ones to observe the 55 Mondrian exemplars looking for patterns. Although this is true, we do not think that too many variables different from and complementary to the ones listed in the rules above (i.e., variables missed by us) would have been found by others to be relevant or important to such a degree that they would have produced radically different results. To support this claim, we performed an experiment (Gómez de Silva Garza & Zamora Lores, 2004) that shows the high degree of agreement that occurs between the ranking of a preselected set of potential solutions when evaluated using the rules implemented in MONICA and the ranking of the same potential solutions when evaluated by human subjects, where the ranking was from the least to the most Mondrian-like. It is also worth noting that our set of rules does assign a 100% fitness value to each of the 55 Mondrian exemplars stored in MONICA, which is a further sign that it is an appropriate set of rules.

The series of evaluation rules we implemented in MONICA, as mentioned above, take into account what seem to us, after a careful analysis of the exemplars, to be the general characteristics and limits of Mondrian's style. For other implementation domains in which there are many more or many less exemplars available, or the exemplars are very complex, perhaps it would have been necessary to use neural networks or data mining techniques, rather than doing the necessary knowledge engineering manually. It is also not necessary for rules to be used as the generic representation scheme for evaluating the style to be imitated; design prototypes (Gero, 1990), genotypes as in Ding and Gero (2001), or other knowledge representation schemes can be used instead. Our process model is flexible enough to allow these alternatives.

3.3. Representing Mondrian paintings

The variables of importance and patterns we observed in the Mondrian exemplars, and the evaluation rules listed above, led us to adopt the following internal genotype representation scheme for the 55 exemplars and for all subsequent potential solutions generated by our evolutionary algorithm in MONICA. First, each painting is subdivided into 20 parts, one for each of the 20 possible colored regions in the painting (see rule 8 above). Figure 5 shows this representation scheme for Mondrian paintings at the highest level.

In a Mondrian painting, each rectangular region can be described by its color, its dimensions (width and height), and its position (the x and y coordinates of its center with respect to the center of the entire canvas). The position is the center point of the painting because this is the standard way in which OpenGL (Schreiner et al., 2007) represents positions. Figure 6 shows how each of the 20 subparts into which a painting is subdivided is represented in this way.

Taking into account that each rectangular region in a Mondrian painting can only take one of four colors, OpenGL places a limit of ± 4.0 as the value of any x or y coordinate, and each painting in MONICA at the lowest level is represented as a binary number, we can determine how many bits are needed to represent each of the five variables used to describe each of the 20 regions into which a painting is subdivided. This lowest level representation of each subpart of the description of a Mondrian painting in our system is shown in Figure 7.

Figure 7 shows 3 bits used to represent a color because initially we were considering the possibility of having to represent white explicitly and perhaps having to represent gray (a few of Mondrian's paintings use gray as well as the primary colors), although in the end, after we had decided on the representation scheme for paintings, we discarded these two options. Thus, only 2 of the 3 available bits are used to represent colors in the implementation. The figure also shows the integer and fraction part of only one of the four coordinates and dimensions that go into representing each colored region, as shown in Figure 6, because the other three are represented in the same way.

In summary, each of the 55 paintings in MONICA is represented by following the scheme described above and shown in Figures 5, 6, and 7, making MONICA require 1660 bits to store the features of each painting, much less than if we had stored the paintings in BMP, TIFF, JPG, or other similar graphical formats. If a given Mondrian painting has less than 20 colored regions (including lines), we fill the unneeded parts of the bit-level representation of the painting with zeros. To ensure that offspring genotypes produced by the crossover operator in MONICA will be of the same length as their parents (i.e., will also consist exactly of 20 parts, whether filled with zeros or not), the implementation of that operator chooses only one crossover point (which it applies to both parent genotypes). For a different implementation domain in which there is no upper bound on the size of a represented design, a different scheme permitting variable-length genotypes (instead of filling some parts of them with zeros, as needed) could

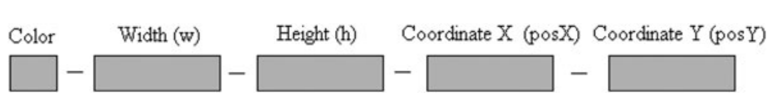

Fig. 6. The internal description of each of the 20 subparts of a painting.

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

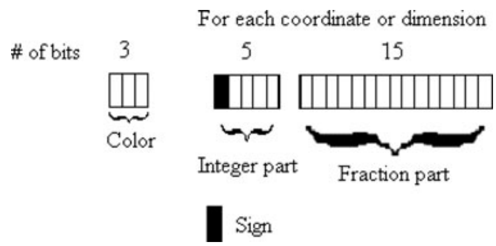

Fig. 7. The bit-level representation of the subparts of a painting.

be used. As a consequence, a different crossover point could be chosen for each parent genotype, thus making it possible to create offspring genotypes that have lengths different from those of both their parent genotypes, as shown in the example of crossover in Figure 2, which illustrates the general case.

4. EXPERIMENTAL SETUP AND RESULTS

To test the effectiveness of MONICA we designed an experiment in which we would be able to determine whether people in general are able to distinguish between paintings produced by MONICA and original paintings produced by Mondrian. The experiment was performed in two phases for reasons that will be explained below. In both phases we used the five Mondrian-style paintings produced by MONICA that are shown in Figure 8.

In both phases we tried two experimental setups. In the first setup we put each one of the five MONICA-produced paintings shown in Figure 8 next to the nine Mondrian paintings that seemed to us to be closest to the set of five MONICA paintings. These Mondrian paintings were probably the direct genetic predecessors, which produced the MONICA paintings shown in Figure 8 after a small number of evolutionary cycles, and are shown in supplementary Figure S.1 (online only). The paintings in the figure are shown with "MondrianX" (where X is a number) labels in order to be able to refer to them easily in the text of Section 4, but supplementary Table S.1 (online only) provides the reference numbers for these paintings given in the *Catalogue Raisonné* for Mondrian (Welsh & Joosten, 1998). The idea in this setup was to make it as difficult as possible for the experimental subjects to identify each MONICA-produced painting by putting it next to some very similar Mondrian paintings.

In the second setup we put each one of the five MONICA-produced paintings shown in Figure 8 next to the nine Mondrian paintings that seemed to us to be furthest from the five MONICA paintings. These Mondrian paintings are shown in supplementary Figure S.2 (online only). Supplementary Table S.2 (online only) provides the reference numbers in the *Catalogue Raisonné* for Mondrian (Welsh & Joosten, 1998) for the paintings shown in the figure. The idea behind this setup was to serve as a control for the first setup by making it easier for the subjects to identify the MONICA-produced painting, because it would be maximally different than the Mondrian paintings shown to them, while in theory still fitting within the same style.

Trying the two experimental setups allowed us to avoid "choosing sides" in a debate that still exists in situations such as police lineups in which a set of people together with a suspect are presented to a witness. If the set of people is chosen to look too much like the suspect then the witness might not be able to clearly identify the suspect (or would do so with approximately the same probability as identifying anyone else in the lineup). If the set of people is chosen to consist of people that look too different from the suspect then the police can be accused of leading the witness by making the suspect stand out. What we have done with our two setups in our experiments is to try to find the middle ground by trying to cancel out the disadvantages of the two approaches.

In all of the experiments that we performed, our subjects were given a sheet of paper on which we presented 10 paintings, 9 painted by Mondrian and 1 produced by MONICA, in two rows of five, the layout of which was different for each phase. The sheet of paper contained the following text: "One of the following paintings was not painted by the same artist as all of the others. Can you indicate which one?" The experimental subjects were undergraduate students (mainly engineering, mathematics, actuarial science, economics, and accounting) who participated voluntarily, and they were told that they could take as much time as they wanted to answer the survey. If the MONICA-produced paintings are virtually indistinguishable from the Mondrian paintings, and if none of the Mondrian paintings stand out compared to the others, we would expect each painting on the sheet of paper that was handed out to each participant to be chosen approximately 10% of the time, because there are 10 of them.

4.1. Phase 1

The layout of the paintings that we presented to our experimental subjects during phase 1 is shown in Figure 9.

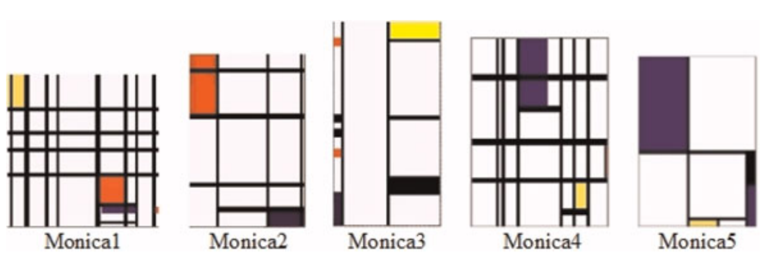

Fig. 8. Five Mondrian-style paintings produced by MONICA. Piet Mondrian 1872–1944. Reproduced with permission of the Mondrian/Holtzman Trust, c/o HCR International, Warrenton, VA, USA. Copyright 2010 Mondrian/Holtzman Trust. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

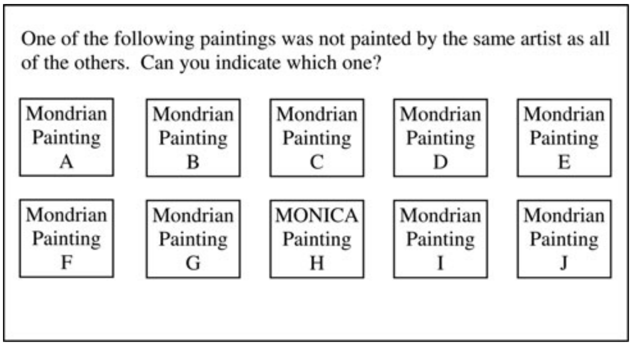

Fig. 9. The layout of the paintings on the paper given to each participant in phase 1.

Phase 1 was split into two experiments. In the first experiment, the nine Mondrian paintings in positions A–G, I, and J from Figure 9 are the ones shown in supplementary Figure S.1, the ones that are very similar ones to the MONICA paintings, using an ordering that we chose at random but that remained fixed across the experiments. In the second experiment, the nine paintings in positions A–G, I, and J from Figure 9 were the ones shown in online supplementary Figure S.2, the ones that are relatively dissimilar to the MONICA paintings, again choosing a random ordering for the Mondrian paintings that then remained fixed.

Each of the two experiments had five variants, one for each of the MONICA paintings shown in Figure 8. These MONICA paintings were always shown in position H, the middle position of the lower row of paintings on the paper handed out to the participants, to avoid having the results of the experiments vary according to the physical locations of the paintings on the page. Thus, we numbered the experiments 1-1, 1-2, . . . , 1-5 and 2-1, 2-2, . . . , 2-5. In total we had 282 respondents whose answers were not cancelled, spread equally among the different variants of the experiments, so there were nearly 30 volunteers that performed each variant of the experiments.

Figure 10 shows the results (the percentage of time that each of the ten paintings shown on the handout was chosen) for experiment 1, Figure 11 the results for experiment 2, and Figure 12 the results for both experiments combined.

As can be seen in Figure 10, even though experiment 1 was purposefully designed to be as difficult as possible for the participants, because the Mondrian paintings that were chosen to be shown next to the MONICA ones were as close as possible to the MONICAs, the correct answer was chosen 37% of the time, at least 10% more frequently than any other painting shown. However, there was also one Mondrian painting (in position J, corresponding in this experiment to the painting labeled Mondrian5 in online supplementary Figure S.1) chosen with very high frequency (26%) and therefore deemed by many people to not fit Mondrian's style (or at least not to fit the style as much as the other paintings shown).

As can be seen in Figure 11, despite experiment 2 being designed to be as "easy" as possible by choosing the Mondrian paintings that were furthest from the MONICAs displayed, in this experiment there were three Mondrian paintings that were chosen more or as frequently as the MONICA painting, and by a large difference. The MONICA paintings, in other words, were on average deemed to be more Mondrian-like than two of Mondrian's own paintings. This is probably because there is more variance among the Mondrian paintings that were presented to people in experiment 2, and thus less opportunity for all our experimental subjects to come up with the same intuitive mental model of what the 10 paintings shown to them all have in common. Thus, in both experiments 1 and 2 we observed the opposite effect of what we had expected as far as the difficulty or ease of the identification.

As can be seen in Figure 12, after putting together all the experiments from phase 1, the correct painting was identified a higher percentage of the time (23%) than any other. We can also see from the figure that each of the five paintings on the bottom row of the paper handed out to the participants was chosen more frequently than any of the paintings from the top row. This led us to believe there to be a bias toward choos-

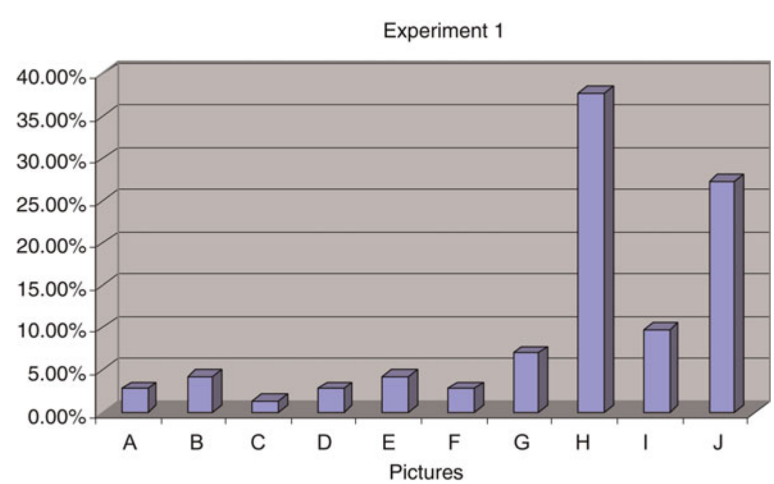

Fig. 10. The results of experiment 1. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

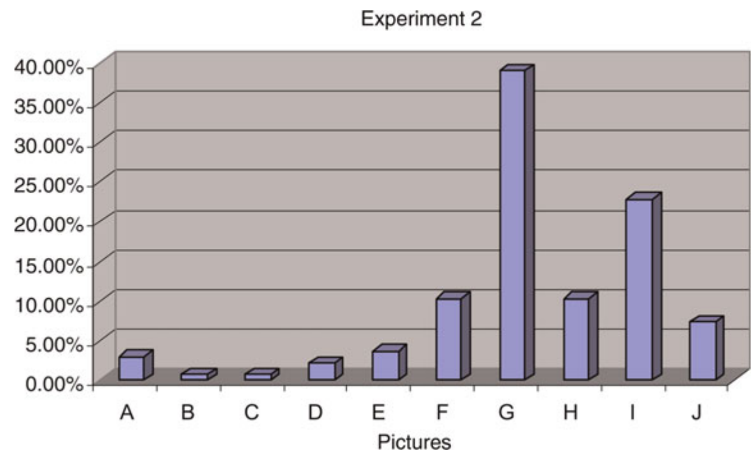

Fig. 11. The results of experiment 2. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

ing a painting on the bottom row, on which people were perhaps focusing their eyes more for some reason. This is the reason we decided to perform phase 2 of the experiment.

4.2. Phase 2

In phase 2 we decided to flip the top and bottom rows of paintings in the papers handed out to the experimental subjects to eliminate the bias toward choosing one of the bottom-row paintings. Although it was not printed on the handout, we also asked the subjects to explain their reasons for choosing a given painting over the others, if they could, to have additional information. The layout of the paintings that we presented to our experimental subjects during phase 2 is thus shown in Figure 13.

Like phase 1, and for the same reasons, phase 2 was divided into two experiments, each of which had five variants, labeled experiments 3-1, 3-2, . . . , 3-5 and 4-1, 4-2, . . . , 4-5. In phase 2 we had a total of 145 respondents whose answers were not cancelled (i.e., just under 15 for each variant of the experiment), perhaps a low number, but these experiments were complementary to the ones from phase 1, giving us further information on the same things, rather than designed to substitute them.

Figure 14 shows the results (the percentage of time that each of the ten paintings shown on the handout was chosen) for experiment 3, Figure 15 the results for experiment 4, and Figure 16 the results for both experiments combined.

As can be seen in Figure 14, picture E produced by Mondrian was chosen more frequently (33% of the time) than picture C produced by MONICA (which was chosen 25% of the time) as not fitting into the same style as the rest (picture E in this experiment corresponds to the picture labeled Mondrian5 in online supplementary Figure S.1, just as in experiment 1). However, the MONICA painting was still identified more than the 10% of the time that would be expected if it were completely indistinguishable from the others.

Figure 15 demonstrates that pictures D and B produced by Mondrian (labeled Mondrian17 and Mondrian16, respectively, in online supplementary Figure S.2) were chosen more frequently (22 and 20% of the time, respectively) than picture C produced by MONICA (which was chosen 18%

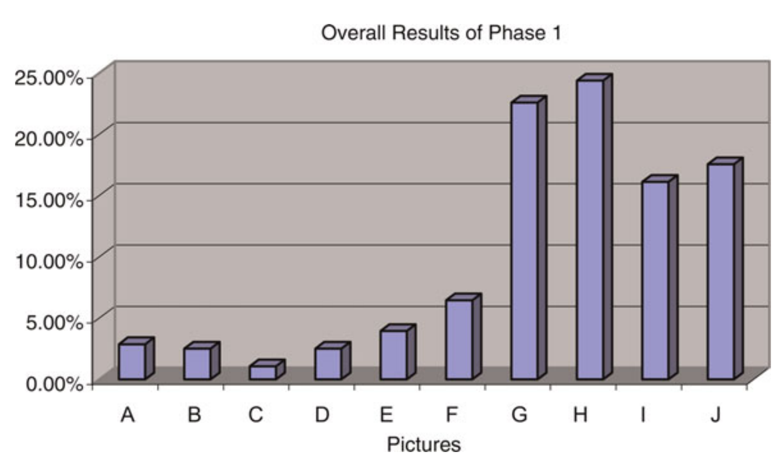

Fig. 12. The overall results of phase 1 (experiments 1 and 2 combined). [A color version of this figure can be viewed online at journals.cambridge.org/aie]

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

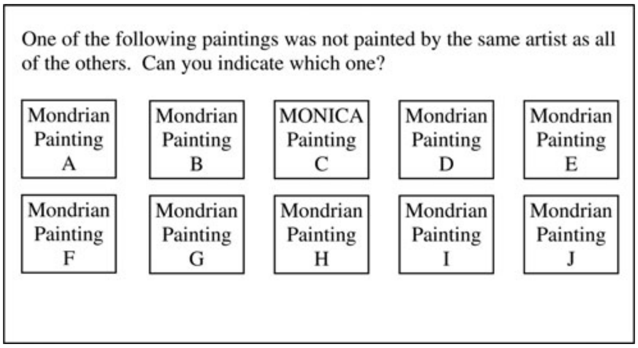

Fig. 13. The layout of the paintings on the paper given to each participant in phase 2.

of the time) as not fitting into the same style as the rest. However, the MONICA painting was still identified more than the 10% of the time that would be expected if it were completely indistinguishable from the others.

What we can observe from the overall results of phase 2 shown in Figure 16 is that the bias we observed in phase 1 toward choosing one of the paintings from the bottom row is not present anymore. It is just the opposite: people were now choosing paintings from the top row most of the time. Therefore, it appears that that "bias" was not because of the positions of the paintings on the handout, but because of something more inherent to that particular set of 5 paintings (actually 10 paintings, 5 used in experiments 1 and 3, and 5 different ones used in experiments 2 and 4), which were on the bottom row in phase 1 and on the top row in phase 2, which continued to be chosen more frequently no matter their position on the page with respect to the others.

Of the 32 respondents (22%) in phase 2 who chose the MONICA-produced painting, most of them (17) gave reasons caused by them picking up on unfortunate and unintended visual cues whose existence we did not realize until after the experiment was performed. These visual cues were artifacts of the scanning process that we used to obtain most Mondrian paintings before printing the handout that was given to the experimental subjects, such as the "white" background of most paintings scanned from books came out looking slightly grayish in the digitized image (thus making the actual use of gray in a few paintings difficult to determine, which is why we ended up not representing gray explicitly in the exemplars we stored in MONICA, as mentioned above). This contrasts quite clearly with the pure white background of all the paintings produced by MONICA. Thus, these 17 respondents gave reasons such as "different colors and clarity," "sharper image," "more brilliant colors," "different tones," "different shade of blue/violet" for choosing the MONICA painting. A further 10 respondents gave what we would call nonsense reasons, or at least nonunderstandable reasons, for choosing the MONICA painting, such as (our comments on each reason are given in brackets) "the lines are thicker than in the other paintings shown [which they're not . . . many of the Mondrian paintings shown together with the MONICA one had some lines of the same thickness as the MONICA one]," "the painting looks too exact, without variations [what does this mean?]," or "the color [which one? There were several in the painting this particular person chose] does not correspond/belong [to what/why?]." If we filter out these two sets of experimental subjects, we are left with very few people that chose the MONICA painting because of noticing stylistic differences that were not caused by the scanning process. Thus, we can conclude that our implementation of our evolutionary method for style imitation manages to imitate Mondrian's style quite well.

5. DISCUSSION

In this paper we presented a computational method based on evolutionary algorithms whose task is to imitate an existing design style. A key aspect of the method is the existence of

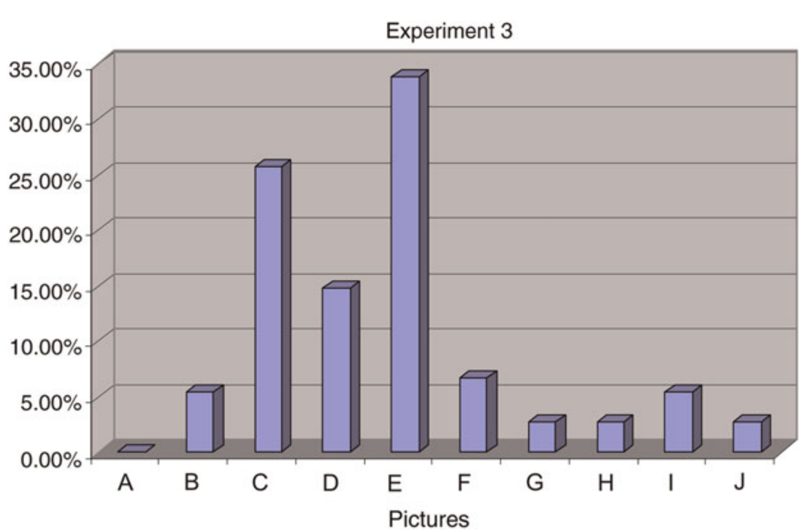

Fig. 14. The results of experiment 3. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

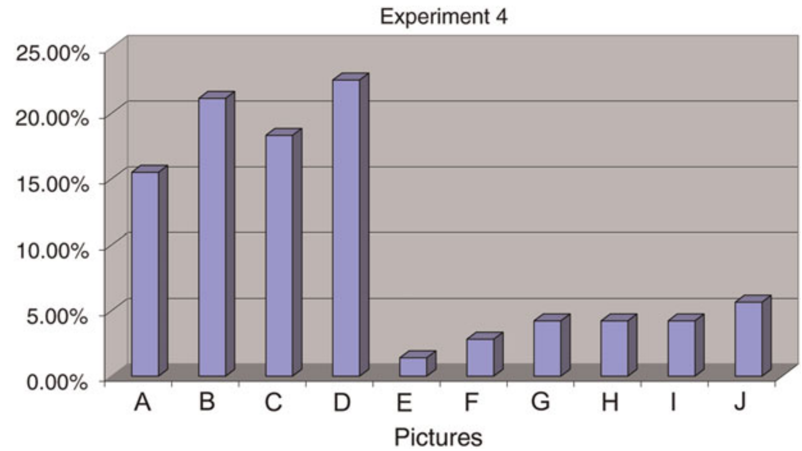

Fig. 15. The results of experiment 4. [A color version of this figure can be viewed online at journals.cambridge.org/aie]

exemplars of the style, which serve both as an important component of the initial population of the algorithm and as the source of knowledge needed for being able to perform the evaluation phase of the algorithm effectively. Another key aspect of this method is that by using an evolutionary algorithm to generate potential solutions we have shifted the need to program domain knowledge to the evaluation (recognition) phase rather than the generation phase. This is important because related research that has tried to imitate style, for instance, using shape grammars or semantic networks require investing a lot of time in performing the knowledge engineering needed to come up with the shape grammar or semantic network for a particular domain. The shape grammars or semantic networks embody generative domain knowledge. Our intuition is that shifting the need for domain knowledge to the recognition phase rather than the generative phase reduces the time required for knowledge engineering both because recognition is less complicated than generation and because, as a consequence, domain experts would find it easier to articulate recognition knowledge than generation knowledge. This intuition can best be understood through an example: it is simpler (both faster and requiring less cognitive effort) for someone that is learning to speak Igbo, for instance, to recognize when Igbo is being spoken by strangers encountered in the street, and to even recognize some of the words being spoken and perhaps even understand their meaning, than it is for the same Igbo learner to generate correct and complete sentences in Igbo, or even carry on an understandable conversation in Igbo.

After presenting the evolutionary method for design style imitation we then described a particular style that we wanted to imitate, the Dutch painter Mondrian's, and a system called MONICA, which implements our evolutionary method for the Mondrian domain. The evaluation rules that were implemented and the representation scheme used for the genotypes in our system were also presented.

We then discussed and gave the results of a series of experiments that were designed to test the effectiveness of MONICA. Our final conclusion is that MONICA manages to imitate Mondrian's style quite well. In addition, the results of the experiments also permit us to make several observations with respect to the notion of style in general and the attempt to imitate it computationally in particular.

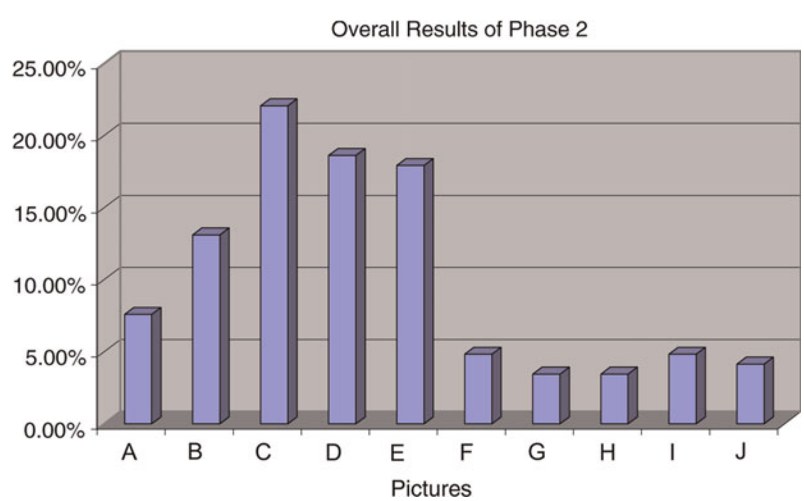

Fig. 16. The overall results of phase 2 (experiments 3 and 4 combined). [A color version of this figure can be viewed online at journals.cambridge.org/aie]

https://doi.org/10.1017/S0890060409990199 Published online by Cambridge University Press

In some experiments Mondrian paintings were chosen by people as being less Mondrian-like than the MONICA-produced paintings, despite them being specifically chosen to be as close as possible to the MONICA paintings, so the concept of style in general, and any given style in particular, seems to be a radial category, as defined by Rosch (1988). These are categories that are easier to describe to others by showing one or more prototypical examples than by explicitly articulating a generic linguistic definition of the concept. Implicit in this kind of category is that different objects will have different degrees of pertaining to the category, depending on their distance from the prototypical exemplar(s). It is also implicit that the boundaries of the category are in general undefined or fuzzy. Hence, people's impressions on the differing degrees of Mondrian-ness of even the Mondrian-produced paintings, not to mention the MONICA paintings, are shown by the results of the experiments (to paraphrase Orwell, all Mondrians are Mondrians, but some are more Mondrian-like than others).

This observation about style as a radial category is also reinforced by the relatively large percentage of the people who could not coherently articulate their reasons for choosing one painting over another that we described in phase 2 of the experiment. It also confirms the appropriateness of several aspects of the design of our process model, for instance, that it is based on using exemplars as starting points for coming up with new potential solutions, its functioning is based on assigning a fitness value according to the degree of Mondrian-ness of the individuals in the evolutionary algorithm population, and this value is calculated based on measuring the distance of each individual to the exemplars (i.e., the degree of fit of the individual according to the rules and constraints derived from the exemplars). In the future, we want to continue using our evolutionary method to explore its effectiveness in imitating more styles, both visual (other painters) and nonvisual (e.g., musical). For styles that are more complex than Mondrian's, it may not be feasible to come up with a set of rules that describe and constrain the style (which can be seen as an explicit linguistic definition of the style, as the rules make reference to certain types of descriptive features and their values), and some other form of knowledge representation that instead embodies prototypical examples of the style may have to be used.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the permission granted for the reproduction of the Mondrian paintings shown in online only supplementary Figures S.1 and S.2 by the Mondrian/Holtzman Trust, c/o HCR International, 85 Waterloo Street, Warrenton, VA 20186, USA, especially Hilary Richardson.

SUPPLEMENTARY MATERIALS

The following supplementary materials can be found online:

1. Figure S.1
2. Figure S.2
3. Table S.1
4. Table S.2

REFERENCES

Bax, M. (2001). *Complete Mondrian*. Aldershot: Lund Humphries/Ashgate Publishing.

Bentley, P., Ed. (1999). *Evolutionary Design by Computers*. San Francisco, CA: Morgan Kaufmann.

Bentley, P., & Corne, D.W., Eds. (2002). *Creative Evolutionary Systems*. San Francisco, CA: Morgan Kaufmann.

Boden, M. (2003). *The Creative Mind: Myths and Mechanisms*. London: Routledge.

Cha, M.-Y., & Gero, J.S. (1999). Style learning: inductive generalisation of architectural shape patterns. In *Architectural Computing from Turing to 2000 (eCAADe)* (Brown, A., Knight, M., & Berridge, P., Eds.), pp. 629–644. University of Liverpool.

Deicher, S. (1999). *Mondrian*. Cologne: Benedikt Taschen Verlag.

Ding, L., & Gero, J.S. (2001). The emergence of the representation of style in design. *Environment and Planning B: Planning and Design 28(5)*, 707–731.

Eiben, A.E., Nabuurs, R., & Booij, I. (2002). The Escher Evolver: evolution to the people. In *Creative Evolutionary Systems* (Bentley, P., & Corne, D.W., Eds.), pp. 425–439. San Francisco, CA: Morgan Kaufmann.

Gero, J.S. (1990). Design prototypes: a knowledge representation schema for design. *AI Magazine 11(4)*, 26–36.

Gero, J.S., & Jupp, J.R. (2003). Feature based qualitative representation of architectural plans. *Proc. Computer-Aided Architectural Design Research in Asia Conf. (CAADRIA-03)*.

Gómez de Silva Garza, A., & Zamora Lores, A. (2004). A cognitive evaluation of a computer system for generating Mondrian-like artwork. In *Design Computing and Cognition DCC'04* (Gero, J.S., Ed.), pp. 79–96. Dordrecht: Kluwer Academic.

Gómez de Silva Garza, A., & Zamora Lores, A. (2005). Case-based art. In *Case-Based Reasoning Research and Development (ICCBR'05)* (Muñoz-Ávila, H., & Ricci, F., Eds.), pp. 237–251. Berlin: Springer–Verlag.

Hancock, P.J.B., & Frowd, C.D. (2002). Evolutionary generation of faces. In *Creative Evolutionary Systems* (Bentley, P., & Corne, D.W., Eds.), pp. 409–424. San Francisco, CA: Morgan Kaufmann.

Kane, C., & Schoenauer, M. (1996). Topological optimum design using genetic algorithms. *Control and Cybernetics 25(5)*, 1–25.

Koza, J. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.

McCorduck, P. (1991). *Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. New York: W.H. Freeman.

Mitchell, M. (1998). *An Introduction to Genetic Algorithms (Complex Adaptive Systems Series)*. Cambridge, MA: MIT Press.

Pagliarini, L., & Lund, H.H. (2002). Art, robots, and evolution as a tool for creativity. In *Creative Evolutionary Systems* (Bentley, P., & Corne, D.W., Eds.), pp. 367–385. San Francisco, CA: Morgan Kaufmann.

Rosch, E. (1988). Principles of categorization. In *Readings in Cognitive Science* (Collins, A.C., & Smith, E.E., Eds.), pp. 312–322. San Mateo, CA: Morgan Kaufmann.

Rooke, S. (2002). Eons of genetically evolved algorithmic images. In *Creative Evolutionary Systems* (Bentley, P., & Corne, D.W., Eds.), pp. 339–365. San Francisco, CA: Morgan Kaufmann.

Rowbottom, A. (1999). Evolutionary art and form. In *Evolutionary Design by Computers* (Bentley, P., Ed.), pp. 261–277. San Francisco, CA: Morgan Kaufmann.

Schnier, T., & Gero, J. (1997). Dominant and recessive genes in evolutionary systems applied to spatial reasoning. In *Advanced Topics in Artificial Intelligence: 10th Joint Australian Conf. Artificial Intelligence* (Sattar, A., Ed.), pp. 127–136. Heidelberg: Springer–Verlag.

Schnier, T., & Gero, J.S. (1998). From Mondrian to Frank Lloyd Wright: transforming evolving representations. In *Adaptive Computing in Design and Manufacture III* (Parmee, I., Ed.), pp. 207–219. London: Springer–Verlag.

Schreiner, D., Woo, M., Neider, J., & Davis, T. (2007). *OpenGL Programming Guide: The Official Guide to Learning OpenGL*. Reading, MA: Addison–Wesley.

Todd, S., & Latham, W. (1999). The mutation and growth of art by computers. In *Evolutionary Design by Computers* (Bentley, P., Ed.), pp. 221–250. San Francisco, CA: Morgan Kaufmann.

Welsh, R.P., & Joosten, J.M. (1998). *Piet Mondrian Catalogue Raisonné* (Vols. 1 & 2). New York: Harry N. Abrams.
Witbrock, T., & Neil-Reilly, S. (1999). Evolving genetic art. In *Evolutionary Design by Computers* (Bentley, P., Ed.), pp. 251–259. San Francisco, CA: Morgan Kaufmann.

Andrés Gómez de Silva Garza is an Associate Professor in the Computer Engineering Department at the Instituto Tecnológico Autónomo de México (Mexican Autonomous Technology Institute, ITAM) in Mexico City. He received a PhD in 2000 from the University of Sydney in Australia, a Master's degree in computer science in 1994 from the Georgia Institute of Technology in Atlanta, and an undergraduate degree in computer engineering in 1991 from the Universidad Nacional Autónoma de México (National Autonomous University of Mexico).

Arám Zamora Lores has a Master's degree in software engineering from the University of Stuttgart in Germany and an undergraduate degree in computer engineering from ITAM in 2004.