Developing Aesthetic Computer Generated Drawings through Artificial Evolution

Kevin Moynihan
Department of Computer Science, University of Miami
P.O. Box 248154, Coral Gables, FL 33124, USA
Email: ktmoyn@hotmail.com

Abstract

This paper discusses the production of visually appealing computer-generated images, through the emulation of human drawing techniques and artificial evolution. Employing a database of line drawn objects from user-input, the Darwinian process of evolution produces original compositions with increasing levels of user satisfaction.

Introduction

The production of aesthetic (e.g., (Machado & Cardoso 1998)) computer-generated images by LASCAUX, a computer programmed artistic agent, begins with a database populated by *genotypes*. Genotypes contain the information required to recreate two-dimensional line drawings. Genotypes are initially formed by recording object representations drawn by humans. As soon as one genotype has been entered into the database, LASCAUX may begin to populate the database with genotypes of LASCAUX generated compositions, via artificial evolution. LASCAUX populates a database with fresh *offspring genotypes* by using artificial evolution to introduce random variations and combinations on any number of *parent genotypes* from within the database.

The parent genotypes selected by LASCAUX, and the variations applied through artificial evolution to produce offspring deemed aesthetically successful by human users, are encouraged in analogous future situations. Offspring genotypes indicated to be more aesthetically fit than their parents have a greater likelihood of being selected to reproduce in the future. The parent objects selected and the variations applied during the generation of aesthetically unfit offspring are used to lower the probability of subsequent automated design choices developing the same adverse results in similar future situations.

Artificial evolution proceeds in a simple cycle, compelling the database genotypes considered aesthetically fit by human users to be chosen by LASCAUX as parent genotypes. The recurring cycle of human users inputting compositions into LASCAUX, and the application of artificial evolution, progressively enhances LASCAUX's decision-making process, the population of line drawn compositions as a whole, and most importantly, LASCAUX's ability to produce aesthetic imagery.

Previous Work

In 1973, the computer programmed artistic agent AARON began as Harold Cohen's attempt to answer a simple question. "What is the minimum condition under which a set of marks functions as an image?" Cohen proposed that if a spectator believed some marks had resulted from a purposeful human, or human-like, act, the marks would indeed serve as imagery. Cohen's definition of "human-like" was that a program would need to exhibit cognitive capabilities quite like the ones we use ourselves to make and to understand images (Cohen 1994).

AARON performs as humans do, in a feedback mode: all AARON's decisions about how to proceed with a drawing, from the lowest level of constructing individual lines and shapes to higher-level issues of composition, are made by considering what it wants to do in relation to what it has already done and the current "state" of the drawing (Cohen 1994). A pleasant side effect of using this feedback-oriented drawing strategy is the distinctly freehand, ad-hoc appearance that AARON's drawings possess.

AARON has verified that aesthetically pleasing images can result from sparse information, given proper interplay between two main types of interacting stored knowledge: specific knowledge about how to draw certain classes of objects (e.g., rocks, plants, human figures, etc.); and a general knowledge of representation (e.g., how to depict solidity or occlusion).

In 1991, Karl Sims presented his research (Sims 1991) on artificial evolution in imagery production, based largely on the use of *genetic algorithms* - robust search techniques over collections of test points evolved by random alterations and combinations. Genetic algorithms attempt to recognize and utilize tendencies towards optimal solutions as guides in a search. Sims employs genetic algorithms to derive imagery that consists of initially simple, randomly generated shapes. Given a population of nine initially random figures, a human user selects the two that appear to be the most interesting. The symbolic expressions that were employed to generate the two images become parent entities, and are combined numerous times in various ways to produce a subsequent population of nine new images displaying attributes of both parent images. From this new population of offspring images, two parent images are selected and the cycle continues. Through this iterated process of interactive selection, the images are advanced under the guidance of the human user to a level considered aesthetically pleasing.

In Sims' system, the user must remain an active participant throughout the entire evolution, continuously selecting two images to be the parents of the subsequent generation's population. However, Sims has demonstrated artificial evolution to be a potentially powerful tool for the creation of procedurally generated structures, textures, and motions. Reproduction with random variations and survival of the visually interesting can lead to useful results.

Artificial Evolution

Artificial evolution (Sims 1991) could be considered a system for helping a human user with creative explorations, or it might be considered as a system that attempts to "learn" about human aesthetics from a human user. In either case, it allows a user and a computer to interactively work together and produce results that neither could easily produce alone. Bringing the expertise of art and technology together has usually been the achievement of one person working alone. As we consider more recent digital art, increasing collaboration occurs between people from different disciplines with different skills. The paradigm for digital art seems to be shifting toward collaborative practice as a norm (Edmonds & Candy 2002).

Analogous to the field of biology, artificial evolution includes the basic concepts of *genotype* and *phenotype*, as well as the processes of *selection* and *variant reproduction*. The genotype is the genetic makeup of an individual. In biological systems, genotypes are strands of DNA, whilst in artificial evolution, genotypes are represented in numerous ways, e.g., strings of binary digits, sets of procedural parameters, or symbolic expressions. The phenotype is the form that results from a set of specialized situational developmental rules and a genotype. In LASCAUX, the phenotype is a composition generated from a genotype in LASCAUX's database.

Fitness is the ability of an organism to survive and reproduce. The likelihood of survival and the number of new offspring an individual generates is directly proportional to its fitness measure. Selection is the process by which the fitness of a phenotype is determined. In artificial evolution, fitness can be either calculated by an explicitly defined fitness evaluation function, or through human feedback (as in LASCAUX and Sims' work). When sets of genotypes are found containing structures with different successful features, it is often desirable to combine the successful features into a single genotype through reproduction, known as the act of *recombination*.

Inducing reproduction by simply splicing fit entity sets together puts an upper bound on the number of possible offspring; an N-dimensional genetic space will always remain N-dimensional. To surpass this limitation, some degree of variation by recombination and *mutation* must be incorporated into the new genotypes. Mutation is accomplished with a *mutation operator* that selects a node within a genotype, and randomly changes its contents. In most cases, only a slight modification of genotype expressions is encouraged. Small mutations over long periods of time have a much greater chance of being eliminated if deemed unfit. Large random mutations in genotype expressions usually result in large jumps in phenotype that are less likely to be improvements, due to the unguided nature of the alteration.

LASCAUX

Overall Algorithm and Process

LASCAUX operates in two basic modes, an input mode with evaluation, and an output mode with evaluation. In input mode, LASCAUX offers a drawing *canvas* and four *primitives* to develop *compositions*. The canvas is a window on the screen. A primitive is a basic two-dimensional shape. LASCAUX's four primitives are (i) a rectangle, (ii) a circle, (iii) a polyline (an open form comprised of connected line segments, i.e. the first endpoint of the first line segment does not have to be the final endpoint of the final line segment), and (iv) a polygon (a closed form comprised of connected line segments, i.e. the first endpoint of the first line segment is always the final endpoint of the final line segment). A composition is a tree of *objects*. The root of the tree is the top-level object. An object contains primitives and/or *sub-objects*. One composition is stored as one genotype entry within LASCAUX's database in the *guidance file*. A composition's genotype stores the information necessary to render the composition. The (sub-)objects within a composition have corresponding (sub-)genotypes within the composition's genotype. Any object (or sub-object) can be extracted from a guidance file genotype and used as a new genotype. Each genotype in LASCAUX's guidance file contains a top-level composition title, an overall compositional rating, each object's title in the composition, a selection method for each object in the composition, and an object specific rating for each sub-object. Each part of a genotype is described below.

Input Mode Before using primitives to draw a composition, a human user enters a *composition title* to describe what will be drawn. The user also provides a title to identify each object in a composition before representing that object on the canvas. Each object may be represented with primitives and/or sub-objects. There are no guiding principles or boundaries to the type of subject matter a user may choose to draw for LASCAUX. All decisions, from the selection of a subject to the design aesthetics, are restricted only by the imagination of the user.

Evaluation Once a composition has been completed on the canvas, LASCAUX prompts the user to provide an *overall compositional rating*, and *object specific ratings*. The overall compositional rating is based on the human user's impression of the aesthetic success of the graphical representation of the intended subject matter in the composition.

Object specific ratings are based on the human user's assessment of the successfulness of an object's graphical representation, and an object's contribution to the overall composition. Both ratings are based on a scale from 0 to 100 (100 being the best). All user ratings are stored within the composition's genotype in the guidance file.

Output Mode Output mode is initiated when a human user enters a *compositional title request* for LASCAUX to draw. LASCAUX generates an original composition by applying the process of artificial evolution on the most aesthetically fit genotypes or sub-genotypes in the guidance file. Any genotype in the guidance file may be selected for the root object of the composition if it is *similarly titled* to the compositional title request. Similarly titled indicates that the compositional title request shares the same title with the genotype's title, or the compositional title request is contained within the genotype's title. The newly created genotype's compositional title is the user's compositional title request that initiated LASCAUX to generate that composition.

During the process of artificial evolution, the combinations and *selection methods* applied to parent genotypes to produce offspring genotypes are based largely on past decisions deemed aesthetically successful by human users. Selection methods describe the reason an object appears in a composition. In compositions produced by a human user, the selection method of each object is attributed to the user. In compositions produced by LASCAUX, an object can be selected by one of the following methods:

- Copied from another similarly titled genotype in the guidance file without variation
- Replaced by a similarly titled object in the guidance file
- Replaced by an object that contains a similarly titled subobject in the guidance file
- Replaced by an object that contains a similarly titled parent object in the guidance file
- Replaced by a completely random object in the guidance file

While a human user titles each object before it is drawn, in genotypes produced by LASCAUX, object titles are copied from the selected parent genotypes. Once LASCAUX completes a composition, a human user is prompted to provide an overall compositional rating and object specific ratings. The user feedback is stored within the offspring genotype that represents the newly created composition in LASCAUX's guidance file.

The Representation of Compositions

LASCAUX has a two-part database. The first part of LASCAUX's database is the guidance file, which contains all the genotypes. The genotypes represent a composition as a group of one or more objects. The second part of LASCAUX's database is the set of *object files*. Object files graphically define individual objects in compositions.

Genotypes and The Guidance File Genotypes are designed so LASCAUX can compare object relationships within a genotype to object relationships within other genotypes. In order to evolve aesthetically pleasing imagery, it is essential for LASCAUX to determine the relationships, selection methods, and genotype combinations that produce the most aesthetically successful compositions.

The objects in LASCAUX's genotypes are arranged in a *nested list* from left to right, beginning with the compositional top-level object and concluding with the last object drawn in the composition. The nested list is an organizational technique that reflects the hierarchical structure and interrelationships of the objects within a genotype. If an object has any sub-objects, the sub-objects are listed in `()` parentheses after the parent object, separated with a comma, in the order that they were drawn.

Although many objects may share the same title, each object represented in LASCAUX's guidance file is a single unique entity and needs to be treated as such. An object's user rating, position within a genotype, and selection method are all applied in the process of artificial evolution in response to a user's compositional title request. The disambiguation of similarly titled objects is vital to a productive learning environment for LASCAUX. Each time a human user enters an object title, LASCAUX scans all of the object files in its database to check if the title has already been used. Each object title is appended with a *unique identifying number* that sequentially follows the unique identifying number of the most recent previously represented object sharing that same title. If an object's title has not been used before, the appended unique identifying number is zero.

In the example below, the genotype represents a car. The compositional title is `car`, and `89` is the overall compositional rating. The top-level object is `car0`; `0` is a unique identifying number, signifying it was the first time `car` was used as a title. `USER` is the selection method for `car0`, signifying it was entered by a human user. `90` is the object specific rating of `car0` in the composition. `car0` has two sub-objects, `door1` and `wheel0`, listed in `()` parentheses after the parent object. The unique identifying number `1` in `door1` signifies that `door` was used as a title once before.

```
car<89>   car0<USER 90>(
door1<USER 78>, wheel0<USER 85>));
```

Object Files Object file data is generated through the observation of human users' graphic depictions of objects on LASCAUX's canvas. Object files are only created for objects that are learned through interactions with human users. When LASCAUX generates original compositions, object files are emulated, not created.

The object titles listed in a genotype are used to locate the object files in LASCAUX's database. When used together with the guidance file, object files provide LASCAUX with the size, shape, and location of every object in a composition, relative to each other object in that composition.

LASCAUX's object files organize the representation of an object into three sections. The first section contains the object title, object dimensions, and object canvas location. The second section is the graphical representation of the object. The graphical representation of an object is a list of primitives and their associated locations. Each primitive is chronicled as the minimum collection of points necessary to regenerate that primitive. The third section of an object file is a list of any sub-object titles, sub-object relative locations, and sub-object dimensions. The sub-objects' titles are used to locate their associated object files in LASCAUX's database.

LASCAUX's Database Design In addition to making LASCAUX's database an efficient storage unit and a flexible framework for the generation of compositions, the database was designed to provide insight into humans' attitudes and reactions to the semantic quality of various objects' graphic representations. For this reason, each part of LASCAUX's two-part knowledge base is written in plain English.

A representational language could have more compactly represented the same information than LASCAUX's plain English database does, but without giving human users the ability to read genotypes and easily comprehend past users' partiality towards object representing graphics, and the success of object representing graphics in different compositional contexts.

LASCAUX's entire database, including the UNIX based directories of object files, can be read, understood, and navigated by a human user with rudimentary computer knowledge and little or no prior knowledge of LASCAUX, or the procedures LASCAUX employs to depict compositions on the screen.

The Production of a LASCAUX Composition

LASCAUX develops original compositions by using the process of artificial evolution in combination with its database of genotypes, object files, and user ratings. The artificial evolution process of recombining genotypes can occur through any of the five previously described selection methods. LASCAUX's selection methods allow any part of a genotype to be inserted into any part of another genotype, permitting parts of even dissimilar genotypes to be combined. If an object selected to replace another object has sub-objects, those sub-objects are inserted into the genotype of the composition being generated, as well. The newly introduced sub-objects are also subject to recombination.

Before generating a composition, LASCAUX scans the guidance file to calculate a *success value* for each selection method. The success values are probabilistic guides that assist LASCAUX in choosing a method to recombine genotypes, with increasingly successful results. The calculation of a selection method's success value begins with the average object specific rating for each object that was chosen by the selection method in the guidance file. Each selection method's average value is then slightly weighted by the average overall compositional rating of each genotype that contains an object chosen by the selection method. Finally, the weighted averages of all selection methods are normalized.

If a user enters a compositional title request that has no similarly titled object in LASCAUX's database, LASCAUX will switch to input mode and ask the user to draw a composition of the requested title. Once the user has completed the requested title composition, LASCAUX switches back to the output mode and generates a new original composition based on the recently entered information.

If a compositional title request has a similarly titled object in LASCAUX's database, an *initial genotype* is selected from the guidance file to guide LASCAUX during the generation of the new composition. The initial genotype is selected based on a combination of the highest overall compositional rating and the highest object specific ratings for the requested object and associated sub-objects within a genotype. The initial genotype can either be an entire existing genotype, or an object and its sub-objects within an existing genotype.

The initial genotype is transferred into RAM as nodes in a linked list tree data structure. Beginning at the root node, LASCAUX traverses the tree and exposes each node to the possibility of recombination through substitution based upon LASCAUX's selection methods' success values, and the object specific rating of the node's associated object. Selection methods that have a high success value are more likely to be employed in the recombination of a genotype than selection methods with a low success value. Likewise, objects that have a low rating are more likely to be recombined than objects with a high rating.

LASCAUX generates the new composition after the entire tree data structure has been completely traversed and recombined. Each node's object title is used to locate the associated object file. Scalar, dimensional, and positional variations are applied to each object before it is drawn in the composition to ensure it roughly occupies the space allocated for the sub-object of the previous object.

Once LASCAUX completes a composition, the user is queried for feedback. When each node in the tree data structure has been assigned an object specific rating, and the user has provided an overall compositional rating, the information is written out to the guidance file as a new genotype.

The Discussion of a LASCAUX Original

Figure 1 contains three seperate user-input compositions from LASCAUX's database. Moving clockwise from the top left hand corner of Figure 1, the graphical representations of a cow, a rabbit and an eye are portrayed.
The genotype for the cow composition in Figure 1 is:

```
cow<77>   cow0<USER 82>(
head2<USER 86>(eye2<USER 79>,
horns0<USER 88>), legs3<USER 89>,
udder0<USER 85>,
butt0<USER 71>(tail1<USER 84>));
```

The genotype for the rabbit composition in Figure 1 is:

```
rabbit<92>   rabbit0<USER 90>(
head1<USER 86>(ears0<USER 91>,
eyes1<USER 83>, mouth0<USER 72>,
nose0<USER 75>(whiskers0<USER 78>)),
arms1<USER 85>(hands0<USER 79>),
legs2<USER 75>(feet1<USER 76>));
```

The genotype for the eye composition in Figure 1 is:

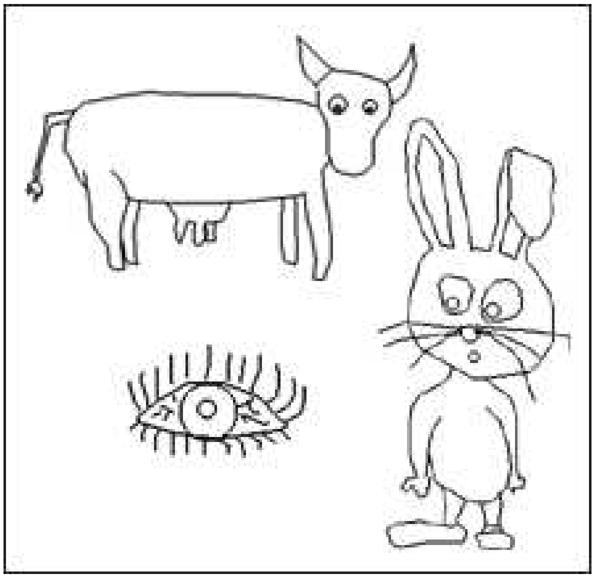

Figure 1: Three seperate user-input compositions from LASCAUX's database. Clockwise from top left: a cow, a rabbit, and an eye.

```
eye<75>   eye0<USER 69>(
eye_lid0<USER 71>(eye_lashes0<USER 82>),
pupil0<USER 77>, veins0<USER 60>);
```

Figure 2 contains an original LASCUAX composition, generated in response to a user's compositional title request for a cow. The genotype for LASCAUX's cow composition in Figure 2 is:

```
cow<85>   cow0<INITIAL 80>(
eye_lashes0<CHILDSUB@head2 86>,
arms1<RANDOM@legs3 89>(
hands0<NOSUB 75>), udder0<NOSUB 82>,
butt0<NOSUB 75>(tail1<NOSUB 75>));
```

The first object in the above genotype is cow<85>, which designates cow as the human user's compositional title request. Within the <> brackets is the human user's overall compositional rating of 85. Reading to the right, the first sub-object is cow0<INITIAL 80>. INITIAL means that cow0 is the root object in the genotype used by LASCAUX in the generation of the new composition in Figure 2 (the initial genotype in this instance is the genotype that produced the cow in Figure 1). The object file for cow0 holds the graphical representation of the polyline that represents the back of the cow, with an object specific rating of 80 in this composition. The rest of the sub-objects in the newly generated genotype are all enclosed in-between a set of ( ) parentheses after cow0, signifying that all subsequent sub-objects are sub-objects of cow0.

The next sub-object is eye_lashes0<CHILDSUB@head2 86>. In the initial genotype that LASCAUX used for guidance, head2 was the first sub-object of cow0. In this new composition, LASCAUX has substituted head2 with eye_lashes0, using a CHILDSUB (child substitution). In the initial genotype, the cow's head possessed an eye as a sub-object. LASCAUX selected another eye object in the guidance file (from the genotype for the eye in Figure 1), and recombined the newly selected eye's sub-object eye_lashes0 into the new composition, with an object rating of 86.

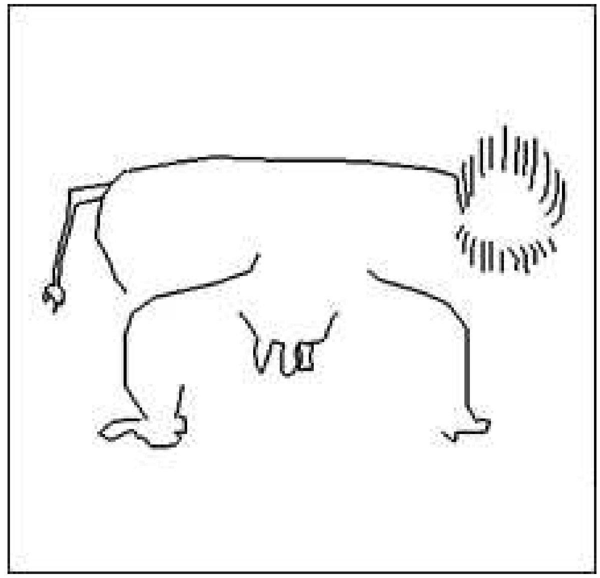

Figure 2: A LASCAUX generated composition of a cow.

The next object in the new genotype is arms1<RANDOM@legs3 89>. Once again, LASCAUX performed a substitution, this time randomly replacing the initial genotype's sub-object with another object from the guidance file. LASCAUX chose to replace legs3 from cow0 with arms from the rabbit's genotype in Figure 1. A human user gave the recombination, and the presence of arms1 in this composition a rating of 89. The next sub-object, a sub-object of arms1 (as noted by the set of ( ) parentheses), is hands0<NOSUB 75>. hands0 was a sub-object of arms1 in the rabbit's genotype and was combined into the new genotype without any substitution, as noted by NOSUB, with a rating of 75.

The remaining objects in the new LASCAUX produced genotype were not substituted from the initial genotype, indicated by the selection method NOSUB that each remaining sub-object possesses. The nested list structure of the genotype confirms that tail1 is a sub-object of butt0, before a semicolon designates the end of the genotype.

If LASCAUX was requested to generate another composition of a cow, all the information contained in the newly added cow genotype would be included in LASCAUX's decision making process. The user feedback from the newly added genotype would update the selection method success values, and the relational information about objects. There is even a chance that LASCAUX's newly entered cow genotype would be selected as the initial genotype for the next LASCAUX generated cow composition. Figure 3 contains 5 sequential LASCUAX compositions, generated in response to a user's compositional title request for a cow.

Conclusion

LASCAUX's decision-making process is initially entirely random. The introduction of human feedback establishes an evolving framework that LASCAUX relies upon for design decisions, through a better understanding of the choices that generate compositions humans find aesthetically fit. User feedback in the guidance file indicates that as the amount of human interaction with LASCAUX increases, the aesthetic success of LASCAUX's compositions ascend.

LASCAUX does not need to exercise any of the cognitive capabilities humans use to create and comprehend imagery in order to produce appealing compositions. Similarly, the user requires no understanding of the methods LASCAUX employs in the production of compositions to direct the progress of evolution and school the program on human preferences. The alliance of LASCAUX's analytical skill with humans' natural ability to judge collectively produces compositions with levels of complexity and expression that do not transpire in separate works created by either partaker.

Many graphically represented objects have either similarly titled, similarly shaped, or similarly located subobjects, which LASCAUX interchanges with interesting and unexpected results. Similarly, LASCAUX also recombines completely unrelated objects into compositions that are scored with successful aesthetic user ratings. User feedback in the guidance file reveals that human users not only accept, but highly rate unorthodox design decisions in some conditions. As Picasso insightfully observed, art is not the application of a canon of beauty but what the instinct and the brain can conceive beyond any canon. When we love a woman we don't start measuring her limbs (Warnche 1995).

In the future, it is hoped that LASCAUX will be able to exploit the mass of information that the Internet holds to better understand the semantics of object titles. Many objects in LASCAUX's knowledge base have either synonymous or related titles that are currently beyond LASCAUX's comprehension. The semantic understanding of object titles would signify a substantial step forward in the ability of LASCAUX to comprehend and demonstrate a more humanistic approach to the graphic representations and understanding of aesthetically pleasing objects and compositions.

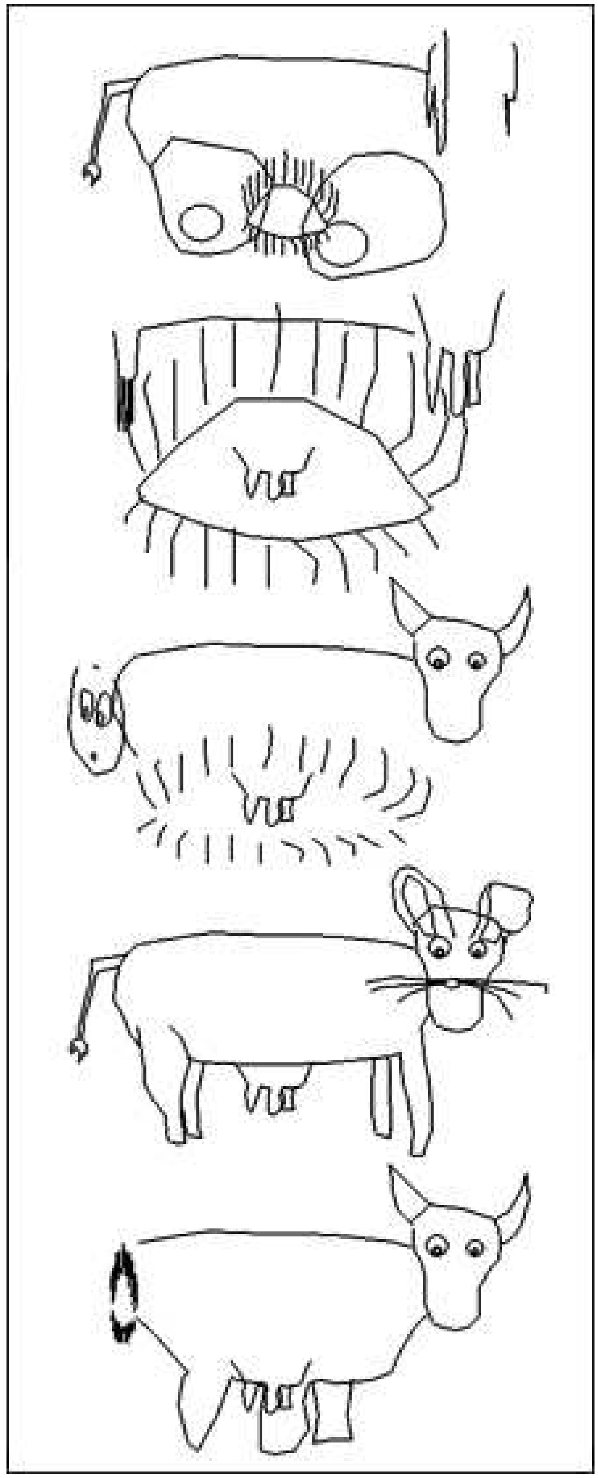

Figure 3: The evolution of a cow. 5 sequential LASCUAX generated compositions of a cow, in order from top to bottom. Each cow composition was generated from the three user-input compositions in Figure 1 and the LASCUAX generated compositions of the cow(s) that precede it. The overall compositional ratings, from top to bottom, are: 20, 5, 38, 92, 65.

References

Cohen, H. 1994. The further exploits of Aaron, painter. *Stanford Hum. Rev.* 4(2):141–158.

Edmonds, E. A., and Candy, L. 2002. Creativity, art practice, and knowledge. *Commun. ACM* 45(10):91–95.

Machado, P., and Cardoso, A. 1998. Computing aesthetics. *Lecture Notes in Computer Science* 1515:219–??

Sims, K. 1991. Artificial evolution for computer graphics. In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, 319–328. New York, NY, USA: ACM Press.

Warnche, C.-P. 1995. *Pablo Picasso 1881-1973*. Koln, Germany: Taschen.