

Evolving Swarms that Build 3D Structures

Sebastian von Mammen

Lehrstuhl für Programmiersprachen
Institut für Informatik
Friedrich-Alexander-Universität
Erlangen-Nürnberg
Martenstrasse 3, 91058 Erlangen
sebastianvonmammen@web.de

Christian Jacob

Evolutionary & Swarm Design Group
Dept. of Computer Science
University of Calgary
2500 University Drive NW
Calgary, AB, Canada T2N1N4
jacob@cpsc.ucalgary.ca

Gabriella Kókai

Lehrstuhl für Programmiersprachen
Institut für Informatik
Friedrich-Alexander-Universität
Erlangen-Nürnberg
Martenstrasse 3, 91058 Erlangen
kokai@informatik.uni-erlangen.de

Abstract- The complex interactions of natural swarms, formed by some social insects for example, are difficult to comprehend. Considering tasks such as nest-building, the necessary underlying communication presumably happens indirectly by changing and reacting on the environment. This paper presents an overall approach to interactively evolve rule-based swarms that create three-dimensional structures in continuous space. The approach comprises the design of the swarm individual, details about the breeding process and first results. A swarm is determined by a set of flocking parameters and a set of constructional rules that allow the individuals to react on and thereby change their local environment. The center or focus of the swarm's endeavour may be shifted either on a swarm individual or on a fixed point in space. Swarms have successfully been bred which orientate their constructions towards pre-defined 3D structures. The alteration of the supplied 3D structure during the course of evolution enables an external supervisor to interactively guide the swarm's development.

1 Introduction

Craig Reynolds' flocks [rey87] provide us with the possibilities to simulate an appropriately coordinated swarm movement. Based on this model Kwong and Jacob have discovered several flight behaviours that are induced by different weightings of a swarm individual's steering urges [jac03]. Besides the coordination of movement social insect swarms also accomplish complex tasks such as nest-building [bon99]. There is only little knowledge about insect swarms' ways of creating 3D structures. Before new insights can come in handy, for example in the domain of self-assembly processes, swarm models and behaviours have to be designed and their capabilities have to be investigated.

Contrary to anthropomorphic approaches of building nests which basically means that each individual follows a blueprint of an architecture an insect's behaviour depends on local information and is determined by a probabilistic stimulus-response scheme. Communication between the swarm individuals happens as an important by-product of this behaviour one individual alters the environment and another one reacts to these changes. It is assumed that qualitative stigmergy, where an individual reacts on a discrete occurrence of some sort, plays the dominant role in build-

ing complex nests. Discrete stimuli can be the construction elements that have already been built or the environmental structures that are provided by nature. Any occurrence that initiates and guides a certain building behaviour is called a template.

Bonabeau et al. have designed rule-based lattice swarms approximate the nest-building of eusocial wasps [bon99]. The individual's behaviour is determined by a set of rules whose preconditions consider the agent's local environment. Pilat [pil04] reproduced these results and found additional rule sets that lead to other forms of wasp nest constructions and totally new architectures. His discoveries were accomplished by interactive evolution of swarms that built promising structures.

This paper presents an approach of breeding similar rule-based swarms that act in a continuous world. The evolutionary process happens automatically by comparison between the swarms' construction and a pre-defined 3D structure. Another important new aspect is the parallel development of the swarm's flocking behaviour along with its set of constructional rules.

As in [pil04] the combination of an evolutionary algorithm with interactive evaluation is an appropriate choice to find new swarm behaviours that result in the construction of 3D structures. Interactive evaluation often leads to success, if there is only a "vague notion" of the concept's objective, the creative potential of the system is not yet fully investigated or the emergent mechanisms of the resulting system are not easily reducible ([daw87], [sim91], [whi01], [tho02], [kwo03] and [jac01]). This paper is organised as follows.

The next section presents the swarm individuals' functionality (genotype), comprising its flocking and construction behaviour, along with the simulation runs that compute a swarm's construction (phenotype). Section 3 describes the genetic operators that work on the swarm's representation and how a creative swarm emerges. Examples bred by swarms and a case study are shown in Section 4. Section 5 gives some ideas on extensions of further approaches.

2 The Swarm Simulation

One swarm consists of a number (by default 25) of uniform individuals. In order to find a swarm which produces interesting 3D structures, we evolve a population of swarms. First, a random flocking and construction behaviour are as-

signed to each member of this population. Then an appropriate swarm is bred by repeatedly computing the swarms' fitnesses and the consequent generation of a new swarm population.

In our simulations the swarm individuals, by default 25 in number, are represented as small spheres¹. A swarmette (= swarm agent) decides on a specific action whenever it collides with a cubic construction element. Each computation of a swarm's fitness takes 500 simulated seconds. The proportions of the simulated world and the occurring objects are listed in Table 1.

Spatial dimensions	$5 \times 10 \times 5$
Min. distance of two construction elements	0.001
Block edge size	0.15
Agent sphere radius	0.05

Table 1: Proportions of the simulation

To initiate the building process, at least one construction element has to be provided. Figure 1 draws a scheme of the constructional process initiated by the collision between an agent and a construction element. An inherited behavioural rule checks the local environment for some structural characteristics. If they apply, which means that there are construction elements at certain positions relative to the collision's location, the rule's consequent action is executed.




Figure 1: An agent, represented by a sphere, is going to collide with the grey cubic construction element (the agent's velocity vector is indicated by the thick arrow). The agent's behaviour depends on the structural configuration that surrounds the collision partner. In the illustrated case thin vectors point to the positions that will be checked by the agent for whether construction elements are present.

The set of possible actions after an agent-block collision is listed in Table 2. A construction element can be built at a location relative to the spatial coordinates of the construction element the agent has collided with. This relation is stated by a direction vector \vec{d} . There is a set D of default direction vectors according to the basic points of the compass (in detail these are \vec{d}_{North} , \vec{d}_{South} , \vec{d}_{East} , \vec{d}_{West} , \vec{d}_{above} , \vec{d}_{below} and \vec{d}_{here}). The absolute values of these standard direction vectors are normalised to the construction element

size, so that direction and distance of \vec{d}_{above} point from one construction element exactly to its upper neighbour's center.

There are two types of destruction methods. The agent may destroy the construction element it has collided with. Another method is to destroy a construction element at a relative distance. The location of the construction element is computed by addition of a vector $\vec{d}_{destroy}$ to the collision construction element's coordinates \vec{p}_c . Of course, if there is no construction element at $\vec{p}_c + \vec{d}_{destroy}$, destruction does not take place.

The last class of actions refers to a center of the swarm which is not implicitly given but has to be explicitly announced. An individual can declare a specific location (including its own) as center of the swarm. If the urge towards the center is sufficiently high, a tendency towards the selected goal will appear. In this way all the agents have a common leader, if an agent has declared itself as center of the swarm. Alternatively, the swarm can focus on a specific region. If no center is defined, the swarm's flight is not influenced by the center urge. In addition to its declaration an agent may recant the current swarm center.

Action	Parameter
Create a new construction element	Vector \vec{d}
Destroy the collision's construction element	None
Destroy a remote construction element	Vector \vec{d}
Declare itself as the swarm center	None
Set the swarm center to a specific location	Vector \vec{d}
Recant the swarm center	None

Table 2: The rule-based agent's actuators. \vec{d} is a location vector that starts in the center of the collision's construction element and points to an arbitrary direction.

On top of these collision dependent actions, the swarm individuals continually alter their velocity according to several differently weighted steering urges. According to Reynolds' "boids" flocking model [rey87] the agents are equipped with visual senses to perceive their neighbours. Assume \vec{d}_{si} as the vector between a swarmette s and another individual i of the swarm. Every individual i is in the neighbourhood N_s of a swarmette s , if the absolute value of \vec{d}_{si} is within the swarmette's radius of perception r and the angle α_{si} between the direction of s and the location of i is within some range $[0, 2]$ degrees radians.

$$\forall i \in N_s (|\vec{d}_{si}| \leq r) \wedge (\alpha_{si} < 2) \quad (1)$$

In each time step of the simulation, the swarmette's velocity \vec{V}_{vel} is updated with an acceleration vector \vec{V}_{acc} :

$$\vec{V}_{acc} = \sum_{j=0}^5 w_j \vec{V}_j \quad (2)$$

w_j , $j \in [0..5]$ are the weights of the distinct tendencies during the flight. The six urges are:

Center, \vec{V}_0 : Considers the vector towards a fixed location within the simulation environment or to an agent's location. If no center is defined, \vec{V}_0 is zero.

¹ We used the VIGO swarm simulation library [bur04].

Separation, \vec{V}_1 : The direction away from all neighbours. As with the next two urges, a neighbour i 's influence on agent s ' flight decreases with a growing distance $\|\vec{d}_{si}\|$. The division by the number of neighbours $|N_s|$ normalises the resulting vector:

$$\vec{V}_1 = -\frac{1}{|N_s|} \sum_{i \in N_s} \frac{\vec{d}_{si}}{\|\vec{d}_{si}\|^2} \quad (3)$$

Alignment, \vec{V}_2 : Adjusting the agent's velocity to the average of the neighbours' velocities. The tendency is computed by subtraction of the swarm agent s ' original velocity \vec{v}_s . Referring to each neighbour's velocity \vec{v}_i :

$$\vec{V}_2 = -\vec{v}_s + \frac{1}{|N_s|} \sum_{i \in N_s} \frac{\vec{v}_i}{\|\vec{d}_{si}\|^2} \quad (4)$$

Cohesion, \vec{V}_3 : The center of gravity of the agent's neighbours. The tendency is computed by subtraction of the swarm agent s ' original location \vec{l}_s . Referring to each neighbour's location \vec{l}_i :

$$\vec{V}_3 = -\vec{l}_s + \frac{1}{|N_s|} \sum_{i \in N_s} \frac{\vec{l}_i}{\|\vec{d}_{si}\|^2} \quad (5)$$

Ground, \vec{V}_4 : Is responsible for the general tendency towards the ground. $height_s$ entitles the swarmette s ' height and $height_{max}$ is the maximum height of the simulation environment.

$$\vec{V}_4 = \left(0, \frac{1}{2} height_s / height_{max}, 0\right)^T \quad (6)$$

Random, \vec{V}_5 : A normalized random vector. Depending on this vector's coefficient (w_5), an unpredictability is introduced to the swarmettes' flight.

3 Evolution of a Creative Swarm

The basic functionality presented in the previous chapter allows a swarm to create three-dimensional structures. Specific building tasks can be evolved by supplying a 3D structure composed of construction elements, towards which the swarm is supposed to orientate its construction.

3.1 Representation

The mentioned weights w_0 to w_5 and maximum values for acceleration A_{max} and velocity V_{max} result in a swarm's flocking behaviour. A set of rules R determine its constructional abilities. Each rule $r \in R$ has the form:

$$c_0 \wedge c_1 \wedge \dots \wedge c_n \rightarrow a, \quad (7)$$

where each condition c_i is internally represented as a vector \vec{p}_i . A condition c_i is fulfilled, if a construction element is found at $\vec{p}_i + \vec{p}_c$ (where \vec{p}_c denotes the location of a construction element the swarmette has collided). The rule consequence a is represented by the number of a specific action

(one of Table 2) and a 3D vector as the action's parameter. Once a collision occurs, each rule $r \in R$ is tested and if applicable its consequent action is executed.

3.2 Genetic Operators

Although a whole set of crossover modes is implemented (random, equal, one point and multi-point crossover masks, as described by Mitchell [mit97]), it is sufficient to confine to one of them. A two point crossover mask is the default choice, so that each of the two offspring owns one part of one and two parts of the other ancestor. If there exist any dependencies within the agent's genotype (e. g., a constructional rule that makes only sense with another one within the same set), its partitioning into three parts very likely conserves them.

The offspring inherit flocking parameters and constructional rules according to the given crossover mask. The offspring's number of rules is limited to the smallest order of the ancestors' rule sets. Hence, it may happen that the average size of the swarms' rule sets decreases in the course of evolution. A small set of rules and still good performance corresponds with Occam's razor which states that "the simplest hypothesis is the best".

Mutation (see Algorithm 1) is applied on every genotype of the new generation. At first the mutation operator checks whether it should alter the given value (v) or not according to a mutation rate (mr). If the decision is made in favour of alteration, an update value (Δv), smaller than a given mutation distance (md), is chosen and added to or subtracted from the original value. Whenever the resulting values leave an interval (defined by a lower bound lb and an upper bound ub), they are trimmed to the next boundary. The mentioned boundaries ensure that the evolved parameters make sense and the mutation distance defines the procedure's maximum effect. Table 3 shows the simulation's default parameters.

Algorithm 1 Mutation Procedure

Returns: An unchanged or mutated value v

Generate a random value $t \in [0, 1]$

if $t > mr$ **then**

Set Δv to a random value between $-md$ and md

$v \leftarrow v + \Delta v$

Return $\min(\max(v, lb), up)$

else

Return v

end if

3.3 Guidance with a Given 3D Structure

Each run of a simulation that is guided by interactive evaluation builds upon a certain idea. One might, for example, try to achieve a construction that reaches very high. During the course of evolution the breeder might run into an unforeseen though interesting structure that inspires his/her objective's notion. However, when starting the simulation, the supervisor must have in mind a reasonably explicit conception. If one is able to map some attributes of the imagined structure

onto a three dimensional construction (such attributes can be height or a general shape, etc), the provision of orientation towards an object turns out to be beneficial.

<i>Evolution parameters</i>	
Population size	20 swarms
Number k of best genotypes	10
Crossover rate	0.4
<i>Mutation Rate</i>	
In general	0.2
On rule actions	0.1
<i>Mutation Distance</i>	
Flocking parameters	0.05
Rule conditions	0.2
<i>Alleles' Boundaries</i>	
Flocking parameters	[0, 2]
Rule vectors	The world size
Maximum number of rules	20
Maximum number of conditions	5

Table 3: Settings of the evolutionary process

In order to guide the search, the difference between a swarm's construction and a pre-defined three dimensional structure is used as the swarm's fitness rating. To measure the difference of the two 3D objects the set of built construction elements is tested for intersection against the set of given cubes. We define the following measures:

Covering Volume, C : Represents the summed intersections of built and pre-defined construction elements.

Non Covering Volume, \bar{C} : Is the volume of the swarm's construction that does not intersect with the pre-defined structure.

Fitness Object Volume, F : Is the total volume of the pre-defined structure.

The fitness function for a swarm is then:

$$fitness_{swarm} = \frac{C}{F} - \frac{\bar{C}}{F} \quad (8)$$

This function reaches its maximum when $C = F$, and $\bar{C} = 0$ which is the case if the whole given structure is rebuilt by the swarm. To ensure that the swarm's construction is as near to the given 3D objects as possible, any out-growth is rated negatively by the second term of Equation 8, $\frac{\bar{C}}{F}$. The more construction elements are built that do not contribute to the given structure's approximation, the lower is the fitness of the swarm. Instead of direct subtraction of the volume of the misplaced construction elements, a softer penalty is imposed. Consequently, the punishment is small as long as the swarm creates fewer construction elements than necessary to fill up the given structure. Additionally, this gives the swarm an impetus to construction in general. However, if the built structure grows rampantly, the penalty will decrease the swarm's fitness.

3.4 Initial Settings

The maximum order of the rule set as well as the maximum length of a single rule are given. The higher the number of conditions, the lower is the chance that a rule is complied with. In order to leave the rules easily applicable, the number of conditions is limited to five. The maximum number of rules is set to 20.

Each condition of a rule is randomly initialized to a vector of the set of basic directions D , introduced in Section 2. The same holds for the vector that is part of the rule's consequence. The action itself can be any of those presented in Table 2, whereas the probabilities are equal that it is about changing the swarm center, creation of a new or destruction of an old construction element.

Triggering the swarm's creative behaviour depends on at least one template construction element. Due to this fact, a "seed"-file is read, which tells the program where to place some initial construction elements before a simulation is run.

The structure which guides the evolutionary process is also read from a file. The genotype of each swarm along with its achieved fitness are stored in a protocol file. The evolutionary process can be resumed, if it had to be stopped. The k fittest genotypes are saved automatically in a separate file.

4 Evolutionary Discovery of Swarms that Build 3D Structures

In this chapter some interesting structures built by rule-based swarms are presented. Furthermore a case-study is conducted. The case-study comprises the swarm's evolutionary development, attributes of its created structure and characteristics of its flocking behaviour. Finally a few examples show, how an already evolved swarm is interactively guided by providing derivations of the originally given 3D structure.

4.1 First Example Structures

Many swarms have been evolved that create interesting structures. Figures 2, 3, 4, 5 and 6 show the given 3D structure along with resulting swarm constructions. Except for the swarms of Figures 3 and 4, which have both a set of 18 rules, all the underlying swarms have 10 constructional rules and make use of maximally 4 to 5 conditions.




Figure 2: (a) The given 3D object that guides the evolutionary search. (b) An evolved swarm builds lines starting from any provided seed construction element. It occurred in generation 15 and achieves a fitness of 0.0064.




Figure 3: (a) Given: one plane. (b) A swarm has been evolved that achieves a fan-like construction by starting from a single seed block at a small distance from the ground in the simulation world center (generation 506, fitness 0.0058).

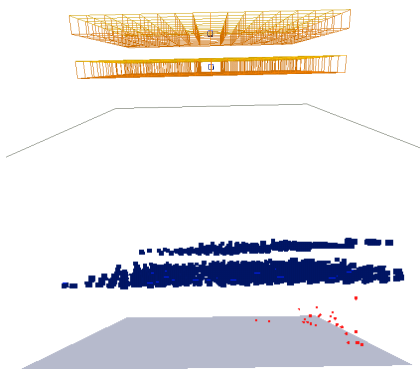


Figure 4: (a) Given are two planes. (b) A swarm of generation 314 creates “two-level flats” starting with two seed blocks at the corresponding heights (fitness 0.0122).




Figure 5: (a) An unsymmetrical 3D structure to guide the search. (b) An asymmetric, organic looking shape (7th generation, fitness is 0.1305).




Figure 6: (a) A simple tower to guide the evolutionary search yields: (b) An interesting shape, reminding of a statue (35th generation, 0.1970 fitness). (c) A swarm that is divided into two flocks at an early stage. Both flocks loop back and forth from their sides to the construction. The many holes in the structure make it look like a bush (1143rd generation, 0.0066 fitness).

4.2 A Case-study: Approximation of a Tower

The given tower structure, displayed in Figure 6, has also been basis of the development of the swarm which is examined in this subsection. In general the swarm in question has similar attributes as the ones presented in Subsection 4.1: Its constructional behaviour is determined by a set of ten rules, whereas the maximum amount of conditions of these rules is five. It took 143 generations for the considered structure to occur and it achieves a fitness of 0.1926.

But the swarm's overall behaviour is exceptional. First it builds up a tower which is very close to the given structure in no time. Then it is fleeing from the world center to avoid further construction which might lead to fitness penalty (Figure 7). The whole construction procedure takes approximately 200 simulated seconds, but the biggest part is already built after 80 seconds have passed. At the simulation start and as long as it takes to create the tower, the swarm stays around the center. Afterwards it is divided into three to four flocks which head as far away from the center as possible, lingering in the simulation world corners.




Figure 7: (a) and (b) show intermediate states of the building process of a tower. The given structure consisting of the bigger cubes, elucidates the degree of the swarm's approximation. (c) After the approximation of the tower the individuals flee into the simulation world corners (seen from above).

The construction can be characterised as follows.

Compactness: The construction elements are built close to each other, therefore the construction is compact.

Structure: Although an iterative construction process can be observed (Figure 7), patterns or modules of construction elements are not apparent. However, the fact that the construction elements are built on top of each other to gain height and next to each other to approximate the given shape's breadth is sufficient to outline

a structure.

Coordination: The construction steps of Figure 7 clearly show that several agents build at different locations without endangering the total construction. Therefore the construction is coordinated.

Another question arises when looking at the single images of Figure 7: How does the swarm change its flocking behaviour? To answer this question the agents' genotype has to be investigated. It consists of the following rules:

- an unconditional rule that recants the declaration of the swarm center,
- five rules adding a new construction element,
- two rules destroying a block,
- two rules rearranging the swarm center.

The succession of the rules plays an enormous role. Emptying the swarm from its center is only realized, if no subsequent rule defines the swarm center anew. Since the rules that redefine the swarm center ask for certain structural configurations around the agent to come into effect, the swarm is focused on the building as long as these conditions are fulfilled. Through steady alteration of the built structure it might occur that these conditions are not satisfied for agents at a certain location. As a consequence the swarm's intrinsic flocking behaviour could come into action urging the swarm to the simulated world corners. This reasoning conforms to the seen phenomenon and is supported by the swarm's genotype.

4.3 Guiding the Search with Diversified 3D Objects

Based on the perfect rule-based tower building swarm, presented in Section 4.2), we further evolved variations of the original tower. Three different ideas for the tower shape changes are discussed (Figure 8): its height extension, an additional branching "crown" and stairs on top of the original tower.




Figure 8: An extension of a smaller tower's height (a), an additional branching "crown" (b) and stairs on top (c).

It took 184 generations until the tower building swarm has adopted to the elongated shape of the fitness structure. In general the same course of events takes place as seen in Figures 7: Now the agents build a somewhat higher tower before they flee from the world center.

Results of the tower extension experiment are displayed in Figure 9. The images represent the results of two separate evolutionary runs, both starting with the genotype of the tower building swarm as discussed above. It is obvious that there is a general tendency of the swarms to fulfill the new requirements. The extension of the original tower's shape gives a strong impetus to build higher. The evolutionary process accommodates the shape's variance and thereby generalises beyond the new given structure.




Figure 9: (a) Pillars arise after 1000 generations (b) Extension of the original "fitness structure" results in endless efforts to gain height (644. generation) Both images clearly show the vertical line flight formation that contributes to the swarm's construction

The branching crown on top of the tower, has been approximated within 1341 generations (Figure 10).



Figure 10: After 1341 generations the original tower building swarm has adjusted to the new challenge: Now it contributes to the branching structure by extending its construction's diameter with growing height.

Figure 11 shows the approximation of the third shape variation. It seems to be hard to cope with the stairs on top of the tower. However, after only two generations the hill-climber strategy yields a genotype with a fitness of 0.2087 which produces structure bent into the direction of the stairs.

A simulation run with no specific starting point has not been able to achieve equally good results within 1000 gen-

erations. Its best result has occurred in generation 925 with a slight tendency of building towards the stairs and a fitness of 0.0401.




Figure 11: (a) The result of an evolutionary run with an optimized tower building swarm as basis. (b) An outcome without a specific initialization (925. generation).

5 Conclusion and Future Work

Normally the evolutionary search can be guided by the supervisor with simple selection or rating of a current population's phenotypes. A method is introduced that allows to direct the course of evolution by iteratively changing its objective.

Adjustments of the swarm model have to be made, too. Currently, the agent's genotype is limited to maximally five conditions per rule. So far artificial evolution has originated mostly swarms that obtain rules with maximally four to five conditions, which clearly exhausts the given limit of conditions. Therefore the effect of an increased number of allowed conditions should be analysed.

Each of the presented swarms has one genotype which holds for all its individuals. It has been shown that the presented swarm model works in general and yields some interesting results. However, more complex structures might arise by assigning individual genotypes to each swarm agent. Since a whole set of distinct swarms is already found, it would be interesting to merge some of their individuals to let them create totally new structures.

Bibliography

- [bon99] E. Bonabeau et al. (1999) "Swarm Intelligence: From Natural to Artificial Systems", Oxford University Press, New York, Oxford.
- [bur04] I. Burleigh (2004) "Master Thesis: Agent-based model of Lactose Operon", <http://pages.cpsc.ucalgary.ca/~burleigh/0412221200/>.
- [daw87] R. Dawkins (1987) "The Blind Watchmaker", Longman Scientific and Technical, Harlow.
- [jac01] C. Jacob (2001) "Illustrating Evolutionary Computation with Mathematica", Morgan Kaufmann, San Francisco.
- [jac03] H. Kwong and C. Jacob (2003) "Evolutionary Exploration of Dynamic Swarm Behaviour", Congress on Evolutionary Computation, Canberra, Australia.
- [kwo03] H. Kwong (2003) "Evolutionary Design of Implicit Surfaces and Swarm Dynamics Master Thesis", University of Calgary.
- [mit97] T. M. Mitchell (1997) "Machine Learning", McGraw Hill, Boston, Massachusetts.
- [pil04] M. L. Pilat (2004) "Wasp-Inspired Construction Algorithms", <http://www.pilat.org/>, University of Calgary.
- [rey87] C. W. Reynolds (1987) "Flocks, Herds, and Schools: A Distributed Behavioral Model", Computer Graphics, 21(4) (SIGGRAPH '87 Conference Proceedings) pages 25-34.
- [sim91] K. Sims (1991) "Artificial Evolution for Computer Graphics", Proceedings of the 18th annual conference on Computer graphics and interactive techniques, pages 319-328, ACM Press, New York, USA.
- [tho02] D. Thomas (2002) "Aesthetic Selection of Developmental Art Forms", Artificial Life VIII, The 8th International Conference on the Simulation and Synthesis of Living Systems, Sydney, Australia, pages 157-163, MIT Press, Cambridge.
- [whi01] T. Whitelaw (2001) "Breeding Aesthetic Objects: Art and Artificial Evolution", Creative evolutionary systems, Section: Evolutionary creativity, pages 129-145, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.