Clonal Mosaic Model for the Synthesis of Mammalian Coat Patterns

Marcelo Walter[a] Alain Fournier[a] Mark Reimers[b]

[a]Department of Computer Science
University of British Columbia
[marcelow|fournier]@cs.ubc.ca

[b]Algorithmics Inc., Toronto, ON
mreimers@algorithmics.com

Abstract

We introduce the Clonal Mosaic (CM) model for the synthesis of mammalian coat patterns, present its implementation for modeling and display purposes, and give a few examples of generated patterns. The model is based on cell division and cell-to-cell interactions, and it can generate repeating spotted and striped patterns occurring in several species of mammals, especially the big cats and giraffes.

From a biological perspective, the model has a strong appeal in light of recent experimental evidence on pigment cells and other pigment related mechanisms; from a computer graphics perspective, the model can not only deliver many patterns which are visually similar to real patterns and can be used as textures, but it is also amenable to simulation on arbitrary surfaces.

Résumé

Nous présentons dans cet article un modèle de mosaïque clonale (CM) pour la synthèse des patrons du pelage de mammifères; nous présentons le but et les caractéristiques de ce modèle et son implantation et nous donnons quelques exemples des configurations produites. Le modèle est basé sur les interactions de cellules, et il peut produire des configurations répétées de taches, de rayures et de rosettes se produisant dans plusieurs espèces de mammifères, particulièrement les grands chats et les giraffes.

D'un point de vue biologique le modèle est très attirant à la lumière de l'évidence expérimentale récente avec les cellules pigmentaires; d'un point de vue infographique le modéle peut non seulement fournir de nombreuses configurations, qui sont visuellement semblables aux vraies configurations et qui peuvent être utilisées comme textures, mais elles sont également appropriées pour la simulation sur des surfaces de topologie arbitraire.

Keywords: natural phenomena, animals, pattern formation, mammalian coat patterns, texture mapping, texture maps, giraffe

Introduction

In computer graphics, the detailed visual information of a surface, such as colour, is usually integrated with the object's surface via texture mapping techniques [3, 14]. Within a texture mapping framework the visual information is maintained as an array — the texture map. There are basically three ways to obtain a texture map: scanning in a real-world pattern, either using a desktop scanner or a 3D digitizer, using a painting system to create an image which will be later used as a texture map (or even painting the image directly on the surface of the model [13]), and procedurally computing or synthesizing a texture.

Scanned real-world pictures are still the main source of texture maps for many texture mapping applications. For some classes of objects, however, scanning a real-world texture or painting an image is too cumbersome and time consuming. Besides, even if we consider that somehow a good texture map is available, problems intrinsic to the mapping technique still make the task of correctly mapping the available texture onto the final surface a difficult one.

Take the case of modeling a tiger and using texture mapping to generate the tiger stripes. If we decide to scan in a texture, we would need a full color image of a "good" and "stretched" tiger skin. On the other hand, we could decide to paint an image to mimic the tiger fur pattern. Even for a skilled user, both approaches would demand a great deal of effort and artistic abilities. The problem would be even worse if we wanted more than one tiger. Despite having roughly the same appearance, each tiger skin has its own characteristics and therefore we would need to build as many texture maps as tigers needed, or find a way to modify a given texture into another, a task almost as difficult as generating the texture in the first place. Alternatively we could define a procedure which would output a "good" tiger pattern. The model introduced in this paper provides such a procedure.

The animal patterns to which we will apply the model include members of the *Felidae* family and the giraffe. Examples of these patterns include the stripes on a tiger

and the rosettes on a leopard. The main idea of the model is that the patterns on these animals reflect a spatial arrangement – a mosaic – of epithelial cells which derive from a single progenitor, i.e. they are clones. Hence we use the name *Clonal Mosaic* (CM). Different hair colors result from different types of underlying cells. The model takes into account important and recent biological experimental data such as the migration of and interactions among epithelial cells [12]. The attractiveness of the model from a computer graphics point of view is that it can generate a large number of animal patterns with a relatively small number of parameters, and this can be done on surfaces of arbitrary shape and topology.

Previous Work

Pattern Formation Models in Biology

Although several models for mammalian pattern formation have been proposed, the actual mechanism responsible is still an open question in biology for most patterns. We can classify the existing approaches into 3 categories: reaction-diffusion, mechanochemical, and cellular automata.

Reaction-Diffusion (RD) was introduced by Turing [35] as a possibility for many pattern formation phenomena. The idea is that the chemical interaction of two substances, under some conditions, can produce stable spatial patterns. The pattern appears if we "visualize" the concentration of the substances. Turing's initial ideas were extended and elaborated into many different models. Murray [22, 20, 21], for example, proposes that the fur pattern reflects a *pre-pattern* established by an RD mechanism. The concentration of substances involved in the process would function as a switch to activate or not specialized pigment cells (melanocytes)[29] to produce one of the 2 types of melanin [15, 10].

The work by Bard [2] presented RD models for the more complex patterns such as the rosette and the markings of different giraffe species. He proposed two mechanisms: *cascade RD*, an idea explored later in graphics by Turk [36], where a sequence of RD processes would explain more complex patterns, and a *threshold interpretation mechanism* for the melanocytes to produce melanin. Bard has also suggested the possibility of having different diffusion rates for different parts of the domain, a suggestion also explored later in computer graphics [39].

Gierer and Meinhardt have proposed a number of RD models to explain both specific and generic, visual and structural patterns. In the book by Meinhardt [18] one can find a good overview of their early generic models. The specific problem of mammalian coat patterns is not directly addressed in any of their models, but indirectly, for example, through models that can generate stripes and therefore could explain striped animals such as zebras. Many of their models included more than two substances to account for more complex regulatory processes.

Most of the current work on RD is addressing theoretical issues rather than experimental ones. Only recently a simple real chemical system has been shown to produce patterns predicted by Turing 45 years ago [17, 26]. Whether or not such chemical systems can be reproduced on biological tissues is still open for discussion. Perhaps the main shortcoming related to the RD theory is the fact that so far experimental biologists have not found or isolated an actual morphogen, that is, a chemical substance responsible for "form" generation.

The *mechanochemical* approach explains pattern formation by mechanical forces acting on cells. The basis for these models was established by Odell [23] and extended by many researchers [25, 38]. The forces are usually considered to be chemically induced. Recently, a new mechanochemical model was introduced by Savic [32] to explain pattern formation in animal coatings. He suggests that coat patterns are an expression of a pre-pattern of polarized and unpolarized domains of epithelial cells. The process of cell polarization is local and regulated through a long range negative feedback mechanism due to elastic forces.

Finally, the *cellular automata* [34] approach has also been proposed to explain mammalian coat pattern formation. Young [41] introduced a cellular automata version of an RD system where the intercellular interaction is more localized than in Turing's original model. Cocho [4, 5] presented a pattern formation framework where the multiplication of cells is modeled assuming an initial small number of "clonal" cells; these advance in time to a more complex arrangement according to the automaton's rules. A clonal cell is a single cell which generates a visible element in the final pattern, such as a spot in a spotted pattern or a patch in the giraffe pattern.

Pattern Formation Models in Computer Graphics

An advantage of using biology-inspired models in computer graphics is their potential to deliver more realistic simulations which can usually be translated into more realistic looking results. The images generated can be used inside a biology context as a powerful argument either against or in favour of the validity of the model [30]. Within the context of mammalian coat patterns we review here the approaches by Turk [36], and Witkin and Kass [39].

The basic RD systems studied in biology can generate a set of interesting but visually limited patterns (simple stripes, simple spots, etc.). Turk [36] developed the idea of cascade RD processes proposed earlier by Bard [2] where an RD system is simulated having as a starting point another RD simulation. A typical example is the

pattern of large and small spots found on cheetahs. Variations on the way two or more RD processes interact can lead to many different patterns. Turk also introduced the idea of simulating the RD process on the surface of the object being textured, an important contribution which avoids many of the problems of texture mapping. However, his approach failed to use information about the geometry of the model to drive the pattern mechanism.

The main contribution of Witkin and Kass's work [39] was to extend the basic idea of RD by incorporating anisotropy into an RD system, a suggestion also made 10 years earlier by Bard [2]. In their work, anisotropy is introduced by assigning different diffusion rates in the RD system as a function of direction in a local frame of reference. In a classic RD model, the same diffusion rate is used for all directions. The control of different diffusion rates for different parts of the surface is achieved through diffusion maps defined by the user. In spite of their usefulness, the use of diffusion maps just transfers to the user the definition of the pattern, since the diffusion maps are often more responsible for the final result than the RD system itself is.

The Clonal Mosaic Model

The basic idea of the CM model, borrowed from an active area of research in developmental biology, is that groups of contiguous cells in an organ are clones, that is, descendents of common ancestors. Applying this to skin[1], it is natural to suppose that cells in differently colored areas derive from different progenitors. We do not know this for a fact, but there is observational evidence in support of such a model. The supposition we work from is that during the early development of the epidermis, some cells differentiate so that their descendants encourage expression of dark pigments, while others differentiate to encourage expression of lighter pigments.

The epidermis does not produce colored proteins; these are produced by specialized pigment cells (melanocytes) [29] which migrate into the hair follicles during embryonic development. However, we know a number of genes, expressed in the dermis or epidermis, which affect the expression of pigment. Chief among these is the well-known *agouti* gene which is responsible for the production of lighter colored bands on the hair of animals such as cats. Where the agouti gene is expressed prominently, the hair is (almost) completely yellow; where it is expressed minimally, the hair is darker, usually brown. Other genes control whether brown, black (or some other color) is the base, and also control the effectiveness of the agouti protein coded for by the gene.

[1]Clonal mosaicism has been demonstrated for most of the major organs in the body, however it has not yet been demonstrated for epidermis.

Most genes are *pleiotropic* — that is, they have multiple effects. We would expect the agouti gene to fit this pattern. If so we might look for other differences between the putative clones of "brown" cells and "light" cells. In fact there are such differences. For example, in the "dark" regions of a cat's coat, the skin thickness is noticeably greater. Also, the density of guard hairs (which carry most of the pigment) is greater than in light areas. We might hope that these diverse effects would be manifestations of some simpler effects of the agouti protein at the cellular level.

Our hypothesis is that one effect of the expression of agouti protein is to affect the growth rate of cells, such that "dark" cells will split faster than "light" cells. This is consistent with the anatomical evidence cited above. Also, other things (such as mobility) being equal, this should result in dark spots, or possibly dark blotches, on a light background. This is in fact what is seen in all members of the cat family.

This working hypothesis leads naturally to the idea that the shape of a pattern element will be the shape of a clone; the shape of a clone will be determined by the deformation induced by non-uniform stresses on the cells during development. The stresses on the epidermis induced by the expansion of the embryo are locally uniform, so that the explanation of non-uniform stresses must lie in non-uniform local expansion of the cell sheet, such as might be caused by non-uniform cell splitting (mitotic) rates.

The Implementation

Whether the clonal mosaic hypothesis is correct or not is obviously a biological problem. Our goal is to determine the characteristics of the CM model as a pattern generator, and to check if we can turn the model into a practical system to generate animal patterns for computer graphics purposes. If we can, on the way, contribute to the validation of CM from a biological point of view, so much the better.

The goal of the implementation described here is to produce a pattern expressed as a 2D image in a regular domain — a square with toroidal boundary conditions. The patterns produced by a given simulation can then be visually analyzed and used to evaluate the model in a feedback loop. Adjustments can be made regarding the parameters and/or specific strategies of implementation.

The current implementation is one possible algorithmic translation of the theoretical abstract model presented in the previous section and provides a powerful computational testbed. Our results from the implementation of the model show that it is possible to obtain fairly realistic looking patterns from various combinations of two parameters — mitotic rates and differential adhesion.

Cells and Groups of Cells

The number of biological cells necessary to represent a given pattern can be very large (on the order of millions at the time the pattern formation process takes place). It would be computationally prohibitive to implement a model which would represent each real biological cell. For this reason we defined a representation scheme where each cell in our implementation is actually a representation for a group of biological cells.

The assumption is that one cell in our system represents the behaviour of a group of biological cells. The issue is then to show that this assumption is plausible in both biological and mechanical terms. The only important biological trait that we have to assess is mitotic rates. Can a single system cell dividing represent many individual biological cells dividing? If the mitotic rates are context-insensitive then after many subdivisions, on average, we will have the same ratio of system cells to biological cells, that is, the assumption is valid.

In terms of mechanical behaviour, if many individual cells are all subject to the same force then we can replace the set of cells by one single cell subject to a force which can be think of as a resultant force. This trades off modeling of individual behaviour for computational efficiency. In other words we might be missing phenomena with scales smaller than the size of a system cell. We think the tradeoff is necessary. Throughout this description we will use the term *cells* to refer to a system cell.

General Description

The potential number of types of cells in the system is arbitrary. However, we restricted the system to 3 types of cells, since we can express all desired patterns with only 3 types. We call them *foreground* (F), *background* (B), and *intermediate* (M). The synthesis of a given pattern is done through two main procedures: *initialization* and *simulation*.

The initialization is responsible for distributing in the domain the initial set of background cells and determining the ones which differentiate into F or M. In a spotted pattern, for example, the foreground cells would correspond to the spots. Once the initialization is done, the simulation through time starts. The initial distribution of cells can also be determined at random, either all at once initially or progressively by the probability of B cells mutating into F or M cells.

The implementation assumes that the only forces acting on the cells result from cells maintaining their sizes under adhesion control. The mobility of cells is a response to these forces. Cell size is maintained by introducing a repulsive force between cells that depends on the distance between them and on pre-defined adhesion values. Cells are modeled as points for computing purposes. Points are usually the first choice to represent a biological structure as a cell [11]. Although points are a very simple primitive, they have proved adequate enough for our purposes. Equilibrium is reached by a relaxation scheme. The idea of using repulsion on a surface to achieve a uniform spatial distribution of points has been used before in biology [33] and computer graphics [36][40]. To turn cells into a tessellation of the surface, we compute their Voronoi polygons. The Voronoi polygon of a point in a given domain is the region of the domain which contains all the points closer to that particular point than any other [27]. The adequacy of Voronoi polygons to describe epithelial cells was studied by Honda in [16]. According to him, "...Voronoi polygons were shown to describe some cellular patterns (cultured monolayer cells, epithelial cells in tissue, etc.) with relatively small deviation values."

Attribute	Meaning	Type
Color	RGB	3 floats [0-1]
Division Rate	Mean time between divisions	float
Initial Prob.	Prob. to be of this type	float (0-1)
Mutation Prob.	Prob. to switch to other type	float (0-1)
Adhesion	Drag between types	float (0-1)

Table 1: Attributes of a cell

Initialization

A given user-specified number of cells is randomly placed on a 2D square domain. Typical initial numbers are between 500 and 1000. The initial position of these cells is given by a random uniform distribution function presented in [28]. Each cell is created with a given type which is related in the theoretical model to the level of expression of the agouti gene responsible for color. The type of a cell can be specified both by the user or randomly by the system. The type of a cell defines its behaviour in the system. The information attached to a given type is: color, division rate, probability for the cell to be of a particular type (only used when type is being determined by the system), probability for the cell to switch to another type (defined for every pair of types), and adhesion (also defined for every pair of types). The current implementation of the probability functions is context-independent, that is, not dependent upon the state of the neighbors. A summary of these attributes is given in Table 1.

The cells undergo a relaxation process in order to achieve a somewhat regular stable spatial configuration. In order to achieve this configuration, each cell moves as far away from all its neighbors as possible. Only cells within a given *repulsive radius* are considered neighbors. The repulsive radius is determined proportionally to the average "ideal" area for each cell. Intuitively, the idea is that each cell tends to occupy a constant area. For a given area A and m cells, the repulsive radius r is given

We want to compute, for a given cell C at position (x_C, y_C), the new position (x'_C, y'_C)

1. For each neighbouring cell P_i at position (x_i, y_i)

   (a) Compute dx_i, dy_i, and d_i

   $$dx_i = x_C - x_i$$
   $$dy_i = y_C - y_i$$
   $$d_i = \sqrt{dx_i^2 + dy_i^2}$$

   (b) Compute f_i

   $$f_i = 1.0 - \frac{d_i}{r}$$

   (c) Compute Dx_i and Dy_i

   $$Dx_i = \frac{dx_i}{d_i} f_i \left(1 - \alpha_{P_i C}\right) r$$
   $$Dy_i = \frac{dy_i}{d_i} f_i \left(1 - \alpha_{P_i C}\right) r$$

2. Compute new position for cell C according to

$$x'_C = x_C + w_x \sum_{i=1}^{n} Dx_i \qquad y'_C = y_C + w_y \sum_{i=1}^{n} Dy_i$$

where

- r is the repulsive radius
- f_i is a scalar which models the strength of repulsion. We need a function for f_i such that $f_i = 1$ when the distance between cells d_i is 0 and $f_i \to 0$ when $d_i \to r$
- w_x and w_y are user-defined weight factors specific for the repulsive displacements in the x and y directions (see section on Anisotropy)
- n is the number of neighbors which fall inside the area defined by the repulsive radius
- Dx_i and Dy_i are the individual displacement forces due to the neighboring cells P_i
- $\alpha_{P_i C}$ are user-defined adhesion values, specific for the kind of cells involved.

Figure 1: Pseudocode for computing the new position of a cell

as $r = w_r\sqrt{A/m}$, where w_r is a user defined scaling value. An adhesion parameter α controls the strength by which cells repel each other in the relaxation scheme. This strength is proportional to $(1 - \alpha)$ and $\alpha = 1$ means no repulsion at all. With this parameter we can, for example, force any two types of cells to remain loosely or strongly connected. In Figure 1 we present the pseudocode to compute the new position for a cell using the relaxation scheme. The individual displacement forces Dx_i and Dy_i are computed proportionally to a repulsive scalar force f_i (dependent upon the distance between cells) and to an adhesion factor $\alpha_{P_i C}$ between cells P_i and C. The adhesion factor is an expression of the fact that cells move at different rates depending on the cohesiveness between cells of various types.

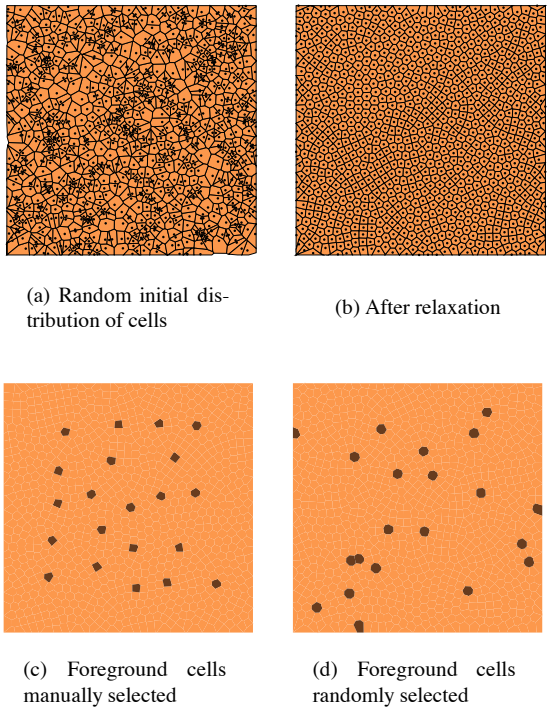

(a) Random initial distribution of cells

(b) After relaxation

(c) Foreground cells manually selected

(d) Foreground cells randomly selected

Figure 2: Initialization (1000 cells)

Once a stable configuration is achieved (that is, the maximum and minimum forces are relatively small) the initialization procedure is over and the system passes to the simulation phase, described in the next section. The exact timing for stopping the initialization step is not critical, since the cells continue to relax in the simulation step. Figure 2(a) shows the created cells before relaxation; in (b) the cells underwent the initial relaxation; in (c) the foreground cells were manually selected, and in (d) the foreground cells were randomly selected by the system.

Simulation

The simulation phase controls the evolution over time of the initial distribution of cells into the final pattern. We model the simulation through an event priority queue implemented as a heap [6]. The two possible events are *relaxation* and *division*. Typically the queue will have many evenly spaced relaxation events and some sparse division events. The rate of relaxation events is user controllable. For each time step, we have ρ relaxation events in the queue. The relationship between ρ and the division rate models the relationship between cell subdivision and cell

motion. A large value for ρ allows time for the relaxation forces to balance over the domain, that is, the cells are closer to equilibrium.

During a division event one cell splits into two, i.e., they undergo mitosis. We can think of these as *parent* and *child* cells. The child cell can be of a different type than its parent, based on a probability matrix given by the user. The child cell inherits all the attributes corresponding to its type. The position of the new cell is uniformly random within a circle of diameter arbitrarily chosen to be 1% of the repulsive radius centered at the position of the parent cell. The exact time for the division of a cell is given by a Poisson distribution with average equal to the rate of division for the cell. The Poisson distribution models small variations on the timing for mitosis, otherwise the cells would all split at the same time. There is no *a priori* end to the simulation. The simulation procedure keeps adding and handling events in the queue. The user can monitor the progress at any time by stopping the simulation and checking the pattern obtained up to that point in time.

Anisotropy

For some patterns we want to be able to set a preferred direction for the cells to move. This can be accomplished in three ways: *i)* when a given cell divides, the position of its daughter is not randomly uniform, but it moves in a preferred direction; *ii)* the repulsive forces acting on the cells have a preferred direction. The first solution can effectively produce anisotropy only if the cell rate of division is high with respect to the rate of relaxation. With the second solution we define two weighting factors w_x and w_y which control how strong the x and y components of a directional force are on a local frame of reference. The third way would be using an anisotropic adhesion factor, which would happen with cells of asymmetrical shapes.

Efficiency Considerations

The most computationally intensive task in the implementation of the CM model is the relaxation step, since we need to find all the neighbors for a given cell. The worst case cost of this procedure is $O(n^2)$ where n is the number of cells.

To avoid this cost we implemented a dynamic rectangular grid of *buckets* over the domain. The linear size of each bucket is the same as the repulsive radius. This scheme guarantees that we only have to check for neighbors within the 8 buckets around the bucket of a given cell plus the bucket that contains the cell itself. Each of these buckets has a pointer to a linked list of pointers to the cells it contains. Since the number of cells grows exponentially with time we need to adjust the grid structure as the number of cells grow. The adjustability of the grid is necessary because the overall domain size is maintained artificially constant. In a growing domain the bucket size would remain constant and the number of buckets would increase as the total area increases.

The grid information is updated (i.e., the number of buckets increase) every time the new number of cells is 50% greater than the previous one. This guarantees a relatively efficient computation scheme. To give a rough idea of timing, the worst case among all computed patterns (Figure 3(b)) took 173 seconds to compute on an Origin 2000 SGI (a 195Mhz processor) and the average time for all patterns was 84 seconds.

Results

In this section we present results of patterns generated with the CM model (the results are also available on-line at `www.cs.ubc.ca/spider/marcelow/cm.html`). In order to better assess the patterns visually, both computed and real patterns are presented. The real patterns were scanned from pictures of animals. The pattern we see on an adult animal is actually the result of two phases of the process, the first which took place some time during embryo development on a shape changing with time, and the second due to the growth of the body after birth. The patterns we produced so far prove that the model is capable of generating a planar 2D pattern which looks similar to a projection of a pattern which is actually defined on the curved surface of the animal's body.

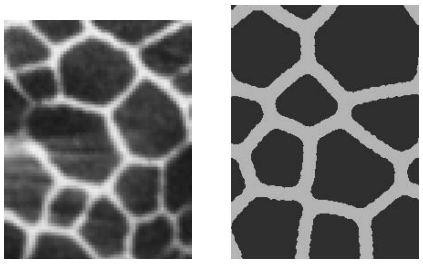

(a) *G. c. reticulata*

(b) Computed

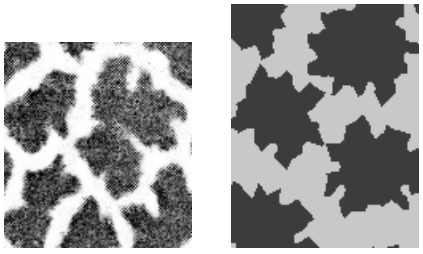

(c) *G. c. tippelskirchi*

(d) Computed

Figure 3: Giraffe patterns

The values for the parameters used to generate the patterns shown in the figures are given in Table 2.

Giraffe patterns

The main taxonomy for giraffes in use today classifies them into one species with 9 subspecies. The differences in giraffe markings have been used as a key feature to identify subspecies, even though this criterion has been replaced by more objective ones such as skull measurements. Visually speaking, the two most distinctive patterns are from *Giraffa camelopardalis reticulata* shown in Figure 3(a) and from *G. c. tippelskirchi*, shown in Figure 3(c). The first is described by Dagg [8] as "the large, smooth-edged liver-colored spots are placed closely together with only a fine network of light color dividing them". The second is also described by Dagg as "the spots are usually splintered, forming all shapes of sharply differentiated leaf or stellate designs, although some approach *reticulata* in design and color". We can easily go from *reticulata* to *tippelskirchi* patterns in our model by decreasing the adhesion between F cells and increasing adhesion between B cells as seen in Table 2.

Spotted patterns

Spotted patterns occur mainly in the cheetah and at the extremities (mainly legs, head, and tail) of other big cats such as the leopard and the jaguar. The cheetah presents usually two distinctive spot sizes whereas for the jaguar and leopard the spots are more regular in their size distribution. Figure 4 shows the real and two computed spotted patterns. In Figure 4(b) the initial probability of F cells was slightly smaller than in (c).

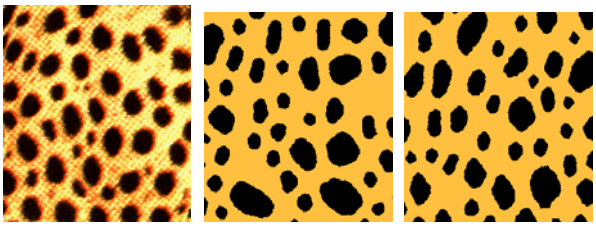

(a) Real (b) Computed (c) Computed

Figure 4: Spotted patterns

For the jaguar and leopard, the spots "break apart" and a third color appears inside the spot. This type of pattern is known as a rosette. We simulate this type of pattern by allowing the foreground type of cells to switch with a small probability to the intermediate M type. Figure 5 shows an example of this result. The extra parameters for this pattern, not mentioned in Table 2, are as following:

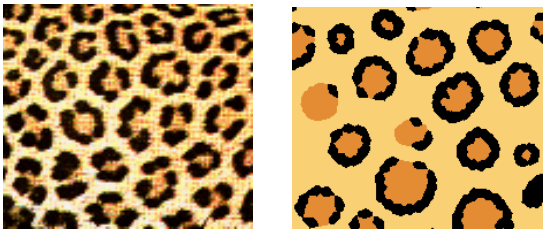

(a) Real (b) Computed

Figure 5: Rosettes

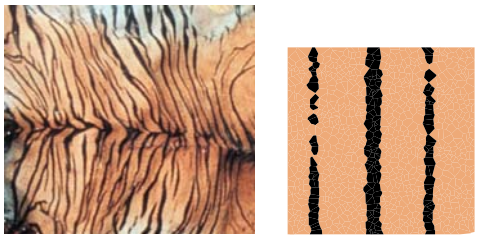

(a) Stretched Tiger Skin (b) Computed Pattern

Figure 6: Anisotropic patterns

mitosis M = 10, $\alpha_{FB} = 0.5, \alpha_{BF} = 0.5, \alpha_{FM} = 0.8$, $\alpha_{MF} = 0.5$, and $\alpha_{MM} = 0.8$.

Anisotropic patterns

Since the tiger is a close relative of all other yellow-black type of big cats, the mechanism for generating stripes in the tiger ought to be of the same type as the mechanism generating spots or rosettes in the other big cats. Therefore we have to consider mechanisms that allow a cellular-based system to eventually produce stripes. We believe that the CM model can easily provide such a mechanism. This possibility has been discussed earlier in biological research by Bard [1] who said on the problem of stripe patterning that "...the stripes might just appear or *spots might be generated on the dorsal line and be extended by an inductive wave moving ventrally*.".

There has been no further research detailing how exactly the wave process mentioned by Bard would work and actually a wave mechanism is not really necessary for the CM model to produce stripes. One clear point is that the growth tensions present on the embryo at the time the pattern is laid down have to play an important role on the final patterns. In order to assess the behaviour of the CM model with respect to anisotropic forces, we have done simulations where the forces are much stronger in one direction than in the other. One result is shown in Figure 6 where $w_y = 40w_x$. A full simulation of these effects de-

mands the simulation of the pattern formation sub-process on a surface which has the same topology as the embryo changing over time, goal of future work.

Assessing the patterns

In order to assess how close a given synthesized pattern is to a real one, we use a qualitative and a quantitative approach. In the qualitative, the generated patterns are visually compared with pictures of real animal patterns. Pictures provide an initial basis for comparison and have been widely used to validate much modeling of natural phenomena either in computer graphics (e.g. [9]) or in biology (e.g. [19]). In the quantitative approach, visually important features of a real pattern are measured and used as a metric for validating results.

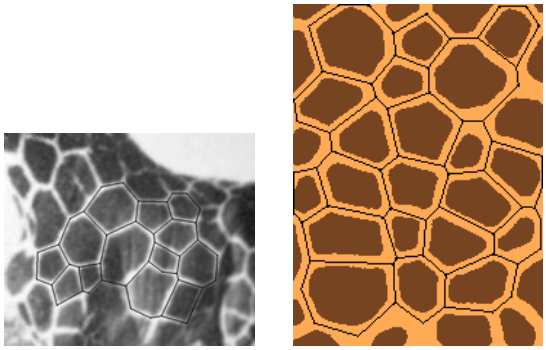

(a) Reticulated M = 1.43

(b) CM Computed M = 0.896

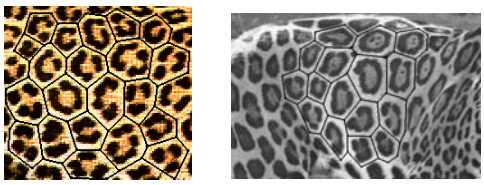

(c) Leopard M = 0.702

(d) Jaguar M = 1.95

Figure 7: Voronoi Measures for different patterns

The main reason the giraffe pattern is used as an example, even though in general our work is more focussed on the *Felidae* family, is that the reticulated pattern is a clear example of a simple geometric pattern, the Voronoi diagram.[2] The fact that the pattern is a Voronoi diagram can be established quantitatively. After scanning in the patterns, we drew by hand the outlines of the spots of the pattern, and applied a geometric construction [24] to each cell that can determine the center of the Voronoi cell. The estimated error on the position of that center is averaged across all cells with a valid center, and this is the number M we use to measure the closeness to a Voronoi diagram.

Figure 7 shows four patterns and the values of M for the *reticulata*, one of our generated patterns, the leopard, and the jaguar patterns. To give an idea of the meaning of the magnitude, a value of 2.41 is obtained if we randomly displace the vertices of the cells by 1% of the average perimeter of a Voronoi cell in an exact Voronoi diagram. Of course for an exact Voronoi diagram $M = 0$.

One can see from these numbers that the giraffe spots closely approximate a Voronoi diagram (the distortions due to the curvature of the body do affect that number). This pattern is quite basic, and it occurs in the big cats as well, although not as spectacularly as in the reticulated giraffe. The fit of the leopard pattern is very good, the fit for the jaguar is less so, but still convincing. These numbers are useful in guiding the choice of parameters, since we now know how close to a Voronoi diagram we have to be.

The CM model can easily explain why a Voronoi pattern is created. If the adhesion between cells of the same type is high, and the adhesion between cells of different types is relatively low, or even zero, then cells of the same type will stick together. If the foreground (spot) cells divide faster, they will crowd out the background cells and push them to lines between the spots. The process is similar to the so-called *prairie fire* model to produce a Voronoi diagram. Other quantitative measures for validation can be used. We reproduce in Table 3 statistical results about giraffe patterns presented in [7].

Species	Spot Area (%)	Number of Sides per Average Spot
tippelskirchi	59	12
reticulata	80	5
rothschildi	50	6

Table 3: Spot areas and spot shapes for giraffes (after[7])

The notion of *spot area* captures how much of the total giraffe body's area is covered with "polygonal spots". In the giraffe patterns produced with the CM model, we can compute an equivalent ratio of the number of foreground cells to the total number of cells and use this value to validate them. The numbers from the CM model are 55 for the *tippelskirchi* subspecies and 78 for the *reticulata* subspecies (Table 2). These numbers are close to the measured ones for the two subspecies, less than 7% variation, a small value considering that the numbers given by Dagg are actually for the whole animal's body.

The number of sides counted, while quite arbitrary for the *tippelskirchi*, is reliable for the *reticulata*, and correspond quite closely to the average number of sides for a

[2] We have to distinguish between Voronoi cells mentioned earlier and used to tessellate the 2D domain, and the Voronoi diagram created by the overall pattern, which are unrelated.

Parameters	ρ	wr	time	wx	wy	mitosis F	mitosis B	α FF	α BB	number of cells	spot area
Giraffe (fig. 3(b))	18	2.6	78	0.066	0.066	10	120	0.9	0.2	B=965 F=3385	78
Giraffe (fig. 3(d))	18	0.6	70	0.066	0.066	10	150	0.2	0.9	B=979 F=1197	55
Cheetah (fig. 4(b))	18	2	15	0.033	0.396	8	60	0.8	0.5	B=1512 F=991	-
Cheetah (fig. 4(c))	18	2	15	0.033	0.396	8	60	0.8	0.2	B=1420 F=1177	-
Rosette (fig. 5(b))†	18	2	60	0.066	0.066	12	30	0.8	0.5	-	-
Tiger (fig. 6(b))	18	2.4	70	0.033	1.32	10	120	0.5	0.5	-	-

Table 2: Table of Parameters for the Computed Patterns.
† For the rosette pattern, the F cell had a 70% probability of switching to an M type of cell.

Voronoi polygon in a Voronoi diagram, which is near 6 [24].

Conclusions

This paper introduced the Clonal Mosaic model for generating mammalian coat patterns and described an implementation of the same. We focussed our study on the patterns from the giraffe and members of the *Felidae* family (e.g., cheetah, tiger). The model proposes that these patterns are an expression of an underlying spatial arrangement of epithelium cells. Different types of cells are responsible for the different hair colors seen in these animals, and the patterns arise as the result of variations in division rates, cell adhesion, and anisotropy in the motion of cells.

The results so far have confirmed the potential of the CM model to deliver an array of patterns visually similar to real ones. In general, fairly realistic looking patterns were obtained from combinations of 2 parameters, mitotic rates of cells and different levels of adhesion between cells. For the giraffe patterns we determined that the basic pattern is very close to a simple Voronoi diagram, and the CM model can account for this easily, both conceptually and with the produced patterns. Another measure is the percentage of the surface area that is covered with spots; we showed that the giraffe patterns produced by the CM model are within 7% of the real patterns with respect to spot area.

The next important step is the implementation of the system on geometric models of animals, coupling the CM model with the growth of the animal's body. We want to generate the patterns directly on the surface of the model, without a texture mapping step. This has to be done while taking into account the growth of the body [37] both at the fetal stage and after birth, which will affect rates of cell division and mobility. This should allow us to obtain realistic patterns fully integrated with the body shape. We are also pursuing the validation of the model from the point of view of the patterns produced, especially the quantitative statistical analysis.

The model we described in this paper has deliberately been limited to context-free rules of behaviour for the cells. We wanted to explore first the range of patterns possible with this simple model (this parallels the evolution in power of L-systems [31] for plant simulation). There are legitimate reasons to extend the model to context-sensitive rules: in order to simulate any reaction-diffusion system, context, in the form of the concentration of the morphogens, is necessary. In real biological systems the behaviour of the cells is clearly affected by context, in the form of signaling chemicals sent across cells.

From the biological point of view, we hope that our implementation will encourage further studies on the mitotic rates, transition probabilities, and expression of the agouti locus in the dermis that can confirm the biological validity of the model. Finally, it should be noted that as it stands the implementation does not help the user much in selecting the parameters for a desired type of pattern. This is an issue that has to be addressed seriously when we feel that the system is mature enough to be used by others.

Acknowledgments

We would like to thank Michael McAllister for the Voronoi code. The first author gratefully acknowledges the financial support of CNPq – Brazil and FUNDEPE (Unisinos). We gratefully acknowledge the support of the Canadian National Science and Engineering Research Council through Research Grants.

References

[1] J. B. L. Bard. A unity underlying the different zebra striping patterns. *Journal of Zoology*, 183:527–539, December 1977.

[2] J. B. L. Bard. A model for generating aspects of zebra and other mammalian coat patterns. *Journal of Theoretical Biology*, 93(2):363–385, November 1981.

[3] E. E. Catmull. *A Subdivision Algorithm for Computer Display of Curved Surfaces*. Ph.D. thesis, University of Utah, December 1974.

[4] G. Cocho, R. Pérez-Pascual, and J. L. Rius. Discrete systems, cell-cell interactions and color pattern of animals - conflicting dynamics and pattern formation. *Journal of Theoretical Biology*, 125(4):419–435, April 1987.

[5] G. Cocho, R. Pérez-Pascual, and J. L. Rius. Discrete systems, cell-cell interactions and color pattern of animals - clonal theory and cellular automata. *Journal of Theoretical Biology*, 125(4):437–447, April 1987.

[6] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, second edition, 1990.

[7] A. I. Dagg. External features of giraffe. *Mammalia*, 32:657–669, 1968.

[8] A. I. Dagg and J. B. Foster. *The Giraffe: its Biology, Behavior, and Ecology*. Van Nostrand Reinhold Co., 1976.

[9] D. R. Fowler, H. Meinhardt, and P. Prusinkiewicz. Modeling seashells. *Computer Graphics (SIGGRAPH '92 Proceedings)*, 26(2):379–387, July 1992.

[10] H. M. Fox and G. Vevers. *The Nature of Animal Colours*. Sidgwick and Jackson Ltd., 1960.

[11] R. Gordon. *Computational Embryology of the Vertebrate Nervous System*, pages 23–70. Elsevier Biomedical Press, 1983.

[12] R. Gordon and A. G. Jacobson. The shaping of tissues in embryos. *Scientific American*, pages 106–113, june 1978.

[13] P. Hanrahan and P. Haeberli. Direct WYSIWYG painting and texturing on 3D shapes. *Computer Graphics (SIGGRAPH '90 Proceedings)*, 24(4):215–223, August 1990.

[14] P. S. Heckbert. Survey of texture mapping. *IEEE Computer Graphics and Applications*, 6(11):56–67, November 1986.

[15] I. Heran. *Animal Coloration*. Hamlyn, 1976.

[16] H. Honda. Description of cellular patterns by dirichlet domains: The two-dimensional case. *Journal of Theoretical Biology*, 72:523–543, 1978.

[17] I. Lengyel and I. R. Epstein. Modeling the Turing structures in the chlorite-iodite-malonic acid-starch reaction system. *Science*, 251:650–652, february 1991.

[18] H. Meinhardt. *Models of Biological Pattern Formation*. Academic Press, 1982.

[19] H. Meinhardt and M. Klinger. A model for pattern formation on the shells of molluscs. *Journal of Theoretical Biology*, 126:63–89, 1987.

[20] J. D. Murray. On pattern formation mechanisms for lepidopteran wing paterns and mammalian coat markings. *Philosophical Transactions of the Royal Society of London B*, 295(1078):473–496, October 1981.

[21] J. D. Murray. A pre-pattern formation mechanism for animal coat markings. *Journal of Theoretical Biology*, 88:161–199, 1981.

[22] J. D. Murray. *Mathematical Biology*. Springer Verlag, 1989.

[23] G. M. Odell, G. Oster, P. Alberch, and B. Burnside. The mechanical basis of morphogenesis: I. epitelial folding and invagination. *Developmental Biology*, 85:446–462, 1981.

[24] A. Okabe, B. N. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley & Sons, 1992.

[25] G. F. Oster. Mechanochemistry and morphogenesis. In A. Oplatka and M. Balaban, editors, *Biological Structures and Coupled Flows*, pages 417–443. Academic Press, 1983.

[26] Q. Ouyang and H. L. Swinney. Transition from a uniform state to hexagonal and striped Turing patterns. *Nature*, 352:610–612, August 1991.

[27] F. Preparata and M. Shamos. *Computational Geometry - An Introduction*. Springer-Verlag, 1985.

[28] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C : the Art of Scientific Computing*. Cambridge University Press, 1992.

[29] G. Prota. *Melanins and Melanogenesis*. Academic Press, 1992.

[30] P. Prusinkiewicz. Modeling and visualization of biological structures. In *Proceedings of Graphics Interface '93*, pages 128–137, May 1993.

[31] P. Prusinkiewicz, A. Lindenmayer, and J. Hanan. Developmental models of herbaceous plants for computer imagery purposes. *Computer Graphics (SIGGRAPH '88 Proceedings)*, 22(4):141–150, August 1988.

[32] D. Savic. Model of pattern formation in animal coatings. *Journal of Theorethical Biology*, 172:299–303, 1995.

[33] M. Tanemura and M. Hasegawa. Geometrical models of territory. *Journal of Theoretical Biology*, 82:477–496, 1980.

[34] T. Toffoli and N. Margolus. *Cellular Automata Machines: a New Environment for Modeling*. MIT Press, 1987.

[35] A. M. Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B*, 237:37–72, 1952.

[36] G. Turk. Generating textures on arbitrary surfaces using reaction-diffusion. *Computer Graphics (SIGGRAPH '91 Proceedings)*, 25(4):289–298, July 1991.

[37] M. Walter and A. Fournier. Growing and animating polygonal models of animals. In *Eurographics '97*, pages C–151–C–158. Springer-Verlag, September 1997.

[38] M. Weliky and G. Oster. The mechanical basis of cell rearrangement. *Development*, 109:373–386, 1990.

[39] A. Witkin and M. Kass. Reaction-diffusion textures. *Computer Graphics (SIGGRAPH '91 Proceedings)*, 25(4):299–308, July 1991.

[40] A. P. Witkin and P. S. Heckbert. Using particles to sample and control implicit surfaces. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 269–278. ACM SIGGRAPH, ACM Press, July 1994.

[41] D. A. Young. A local activator-inhibitor model of vertebrate skin patterns. *Mathematical Biosciences*, 72(1):51–58, November 1984.