

Integrating Shape and Pattern in Mammalian Models

Category: research

Abstract

The giraffe and its patches, the leopard and its spots, the tiger and its stripes are spectacular examples of the integration of a pattern and a body shape. We propose in this paper a system which combines and integrates a pattern generation system which can effectively deliver a variety of patterns characteristic of mammalian coats, and a body growth and animation system that can use experimental growth data to produce individual bodies and their associated patterns totally automatically.

We use the example of the giraffe to illustrate how the models take us from a canonical embryo to a full adult giraffe in a continuous way, with results that are not only realistic looking, but also objectively validated. The flexibility of the system is demonstrated by examples of several big cat patterns, including an interpolation between patterns.

The method also allows a considerable amount of user control to generate *un-natural* patterns and/or shapes and to animate the resulting body with the pattern.

CR Categories:

Keywords: integration, animal models, animal patterns, clonal-mosaic, growth, animation

1 Introduction

The task of generating a computer graphics rendering of an object is traditionally broken down into two parts: the modeling of the shape or geometric attributes such as height, width, *etc.* and the modeling of the visual attributes, or how the object is going to look. A final *integration* step connects the two, that is, a visual attribute has to be defined for every point on the surface of the object.

The separation of these two modeling tasks makes the whole process highly flexible and powerful, and from a conceptual point of view easier to handle. While generally good for many classes of objects, this two-step process is nevertheless prone to problems when the geometry of the object is complex and therefore the mapping of visual characteristics to every point of the surface is non-trivial. Another potential drawback of this approach is that it implicitly assumes that there is no interplay between the two processes which define the shape and the visual attributes of an object. For some objects, however, the visual aspect is the result of a process where there is an interaction between the two and this interaction plays a significant role on the final result. The typical example are patterned animals such as giraffes and leopards, where the pattern visible on the fur of an adult animal is the result of a much earlier process which took place while the animal was in the womb. In this case to model the interplay between the embryo growing and the pattern formation process is as important as the modeling of the individual processes themselves.

Here we present a method for these cases whereby the visual attributes are defined directly on the surface of the object and, more importantly, where we take into account the dynamic change of shape undergone by the object because of growth or other reasons. We will show how a mammalian coat pattern can be generated by a biologically-plausible model simulated on the surface of a changing geometry, where the interplay between the two plays a decisive role in the final result.

2 Related Work

2.1 Mammalian Coat Patterns

An advantage of using biologically-plausible models in computer graphics is their potential to deliver more realistic simulations which can usually be translated into more realistic-looking results. In a biological context the images generated can be used as a powerful argument either against or in favour of the validity of the model [prus93].

The basic reaction-diffusion (RD) systems studied in biology [turi52, murr81a] can generate a set of interesting but visually limited patterns (simple stripes, simple spots, etc.). Turk [turk91] developed the suggestion made earlier by Bard [bard81] about cascade RD processes, where a 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. The web-like pattern similar to the reticulation pattern found on giraffes, for example, is explained by the simulation of two different RD systems (spots formation and stripe formation) together into one. 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. His method did not, however, use information about the geometry of the model to drive the pattern mechanism.

Witkin and Kass [witk91] extended the range of possible RD patterns. Their main contribution was to extend the basic idea of RD by incorporating anisotropy into a RD system, a suggestion also made 10 years earlier by Bard [bard81]. 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 transfer 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.

It is fair to say that while both papers show a giraffe-like pattern among their examples, neither patterns are very convincing (while the giraffe pattern is very striking, it is quite amazing how far from the true pattern even skilled artists allow themselves to go).

2.2 Integration of Pattern and Shape

Most of the work on the integration of the shape and the visual aspects of an object has been done in the context of texture mapping and focussed on how to deal with the problems intrinsic to the technique, such as texture placement and texture distortion. More than twenty years after texture mapping was first described by Catmull [catm74] the graphics community is still addressing such problems. This is clearly a strong motivation for researching alternative methods for pattern placement.

It is not of course that there has not been any progress made. In [litw94], for example, an interactive system for texture placement is proposed. The idea is to define indirectly the mapping via user manipulation of the texture while it is being placed on the surface of the object. Objects and camera are constrained to remain fixed during the whole texturing process, which can be a great limitation

is some cases. Also it is not very clear how the technique would work when the same texture is shared by many surfaces.

An alternative approach for texture placement is to paint a texture directly onto the surface of the object [hanr90]. The objects are modeled as a collection of many very small polygons from which a parameterization of the surface can be derived. The mapping function is indirectly established by the user “painting” on the surface. There is no actual distortion to be corrected since there is no *a priori* texture map to be distorted. The drawbacks are that *i*) the final result is still highly dependent on the artistic abilities of the user and therefore achieving a visually elaborate texture can be difficult and *ii*) as presented the approach does not handle a pre-existing texture, and therefore cannot be used to correct texture distortions by visual inspection.

A solution to minimize texture distortions when mapping pre-computed textures was presented in [mail93]. They define a metric for the distortion and try to minimize its value globally. This approach is limited since for very complex objects this minimization step might not be feasible. One alternative is to use a mapping function which is local and not necessarily continuous. In other words the object is split into *charts* and a collection of charts is called an *atlas*. The creation of charts takes into account surface curvature and the user interacts visually to achieve the best atlas for a particular object. The idea behind charts is to represent a non-developable surface as a set of developable surfaces. A developable surface is a surface that can be deformed to planar shape without changing length measurements in it [fari90]. In a similar fashion the work presented in [benn91] “cuts” a given 3D parametric surface into regions which can be flattened out without warping. The minimization of distortion is achieved through a compromise between cuts and distortions.

When the user has control over how the texture is being generated some more effective ways to avoid distortions are possible. The method in the paper already cited [witk91] uses models described parametrically as a collection of patches and synthesizes textures with reaction-diffusion systems. The problem of texture distortion caused by the mapping from the texture parameter space to the surface space was solved in an integrated manner. The texture synthesis incorporates a correction factor for the distortion, that is, the diffusion rates present in the reaction-diffusion system were controlled to account for the geometric distortions present on the surface. This correction, however, only works for surfaces which can be described by a single parametric function, usually not the case for complex surfaces. Seamless periodic textures were created using cyclic boundary conditions, i.e., points that shared a common boundary in different patches had the same boundary conditions. In the reaction-diffusion context this means that the chemicals involved have the same concentrations at the boundaries of the parametric patches.

There has been little work addressing integration as a task *per se*, outside the context of texture mapping. Two remarkable exceptions are the work by Greg Turk presented in [turk91] and by Deborah Fowler, Hans Meinhardt and Przemyslaw Prusinkiewicz presented in [fowl92]. Both papers present variations on the fundamental idea of computing the pattern on the surface of an object as a “growth-in-place” procedure.

Turk’s work, as mentioned in the previous section, also used reaction-diffusion textures. Instead of mapping the generated pattern in a polyhedral or parametric model, his approach simulates the reaction-diffusion system on the surface of model, without the intermediate mapping from texture space to object space. Basically, the surface of the model is divided into cells and the reaction-diffusion system is simulated directly on the mesh formed by these cells. The cells for the simulation are the regions of a planar variation of a Voronoi diagram [okab92] computed from a polyhedral description of the model. The approach does not have the usual problems of texture discontinuity and distortion since there is just one mesh over

which the reaction-diffusion system is simulated. Nevertheless, this solution does not make any use of the geometric information about the shape of the object being textured and therefore the texture on the surface can appear strangely regular and artificial. Moreover, to achieve a more natural-looking result Turk introduced adjustments so that the texture could grow faster in some areas and slower in others. This was achieved with the user defining “initiator-cells” and different diffusion rates for different parts of the surface.

Fowler *et al* [fowl92] similarly approached the modeling of seashells by discretizing the growing edge of a parametric model of a shell into polylines. Each segment of the polyline is treated as a cell for the one-dimensional reaction-diffusion simulation. The geometric and visual attributes of the shells are easier to integrate since both shape and texture can be unequivocally expressed as a function of time. The integration is therefore easier since time can be used as an integration factor. Their exceptional results suggest that the use of an integrated approach is in some cases not only useful but imperative.

Another closely related work is by Fleischer *et al* [fleis95] on cellular texture generation. In many ways their work is more general than ours, allowing for full programmable and context-sensitive behaviour of the cells creating the textures. As usual in modelling, however, too much power can be a bad thing as there is no obvious relationship between the equations governing the cells and the results. In other words it is possible to generate an extraordinary range of models, but hard to see what they are models for.

Our overall goal is to match for mammalian bodies and coat patterns the level of achievement reached by Fowler *et al*. There will be many critical differences in the methods used. First, for reasons elaborated in details elsewhere, we do not believe that mammalian coat patterns are well modelled by reaction-diffusion methods, and we will use a different model, called *Clonal Mosaic*, summarized below. Second the process obviously has to take place at least in two dimensions corresponding to the skin surface. And finally since in mammals the pattern is established in the fetal stage, and undergoes changes due to body growth both before and after birth, we will have to integrate pattern formation and growth in an effective way.

3 Basic Tools and Data

3.1 How and When is the Pattern Established

The coloured pattern in mammalian coats is produced by the various colours of the hair constituting the *pelage*, as the collection of hair is called. The skin of mammals is made of two layers, a superficial layer called the *epidermis*, and an inner layer called the *dermis*. The epidermis is composed of flattened cells, and hair is produced by an *invagination* of the epidermis called a *follicle*. The hair is formed by division of cells in the *bulb* at the base of the follicle. Pigments giving the hair colour are incorporated into the hair by *melanocytes*, basal cells of the epidermis specialized in the production of *melanin*. Melanins are polymers synthesized from *tyrosine* (an amino acid) and have colour ranging from pale yellow to black, through buff, red-brown and brown. Basically the density of melanocytes is not what determines the colour of the hair, but the amount and the nature of the melanin they produce. Note also that melanin contributes to the colouration and protection of the skin as well in humans. The melanocytes have migrated during embryonic development from the *neural crest* to their final position in the epidermis as part of a complex of cells called an *epidermal melanocyte unit*.

Direct observation of mammals shows that their characteristic coat pattern is already established at birth, and after this is modified only due to differential growth of the body. For instance the spots on an adult giraffe can easily be recognized from the spots of the same individual at birth. Note that sometimes the spots can fade or




Figure 1: A 35-45 days giraffe embryo (457 days gestation) from [murr81b]

disappear due to a change in "colour map". For instance lions have spots at birth, which quickly fade and are not visible in the adult. So one can distinguish two phases in the creation of the pattern: the first phase takes place in the fetal stage, where growth and establishment of the pattern both take place, and the second phase, both before and after birth, where only growth affects the pattern. Of course the pattern is only visible after the pelage has grown, but one can assume that the distribution of the melanocytes responsible for the colour has already taken place.

One cannot go into details to justify the bounds that can be established for the starting time of the first phase and the second phase, but we will give the main results we will use and some pointers to the relevant literature. The start of the pattern formation is most likely as soon as the melanocytes have finished their migration from the neural crest. That means that it is about 35 days for the giraffe (out of 457 days for the total gestation time).

At this time the fetus already has a recognizable shape, as can be seen in Figure 1 which shows a giraffe embryo with an estimated age of 35 to 45 days.

The upper bound for the end of phase two is of course the end of gestation, but there are reasons to believe that it might be sooner. Murray [murr81b] from considerations based on a reaction-diffusion model, deduce that the total duration of the pattern formation phase is quite narrow, but that does not necessarily apply to other models, such as clonal mosaic. In our simulation we have generally assumed a time of about 200 days for the giraffe (about half of the gestation time). A very useful fact is that during that whole time one can safely assume a linear growth (that is the length of some part of the body is a linear Figure 2 shows the plot obtained for the length of the giraffe fetus from measurements found in [owen49] and [bedd06].

It is important to note that it actually is a distinct advantage to "go back" to the fetus, since we have only to use a *canonical* fetus, which will then be grown to the dimensions of a given individual by the process described in section 3.3.

3.2 The Clonal Mosaic Model

The model is explained and illustrated in a paper by Walter and Fournier [walt98] and we will only stress the main points here. 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 – they are clones. Hence the name *Clonal Mosaic* (CM). Different hair colors result from different types of underlying cells. The model takes into account important and re-




Figure 2: Fetal Length for Giraffes

cent biological experimental data such as the migration of and interactions among epithelial cells [gord78]. 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. Clonal mosaicism has been demonstrated for most of the major organs in the body. Even though it has not yet been demonstrated for epidermis, the same mechanism will be extended to its epithelial cells. The supposition 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.

It is known that there are genes which affect the expression of the pigments. A well known one is the *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. Most genes are *pleiotropic* – that is, they have multiple effects. Part of the model is the hypothesis that one effect of the expression of *agouti* protein is to increase the growth rate of cells. Another effect is a different adhesion between cells of the same types and cells of different types, influencing the tendency of cells to "stick" together or to mix more freely. This leads to the conclusion that the shape of a pattern element will be determined by the shape of a clone. In turn the shape of the 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 mitotic rates and adhesion factors.

The number of biological cells necessary to represent a given pattern can be very large (order of millions for the time the pattern formation process takes place). It would be computationally prohibitive to implement a model which would have each real biological cell represented in the system. For this reason we defined a group representation scheme. 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 individual behaviour modeling for computational efficiency. In other words we might be missing phenomena with scales smaller than the size of a system cell. We think that the tradeoff is necessary. Throughout this description we will use the term *cells* to refer to a system cell.

The patterns we will use are generated by a simulation of the clonal mosaic process. For all the patterns of interest we use a maximum three types of cells. We will call them *foreground* (abbreviated F), *middleground* (abbreviated M) and *background* (abbreviated B) cells. In the giraffe only F) and B) cells are needed.

The initialization distributes among background cells a number of foreground cells. The same result can be accomplished by having background cells switch to foreground cells in a controlled manner. Cells then divide at rates that are characteristic of their type (usually the foreground cells divide faster than the background cells). As new cells appear they try to maintain a constant area (in fact a volume, but we assume that the thickness of the layer is constant) by repulsing their immediate neighbours. The repulsive force in our model is a linear function of the distance to the neighbours, reaching zero at a distance equal to the cells *radius of influence*. For each cell and each relaxation step the resultant force results in a displacement which is in the direction of the force and proportional to the force intensity and to the adhesion characteristic of the pair of cells interacting. The new position of the cell is then computed from the displacement. This is done in parallel for all the cells. In order to model anisotropic displacement, the actual cell displacement is the weighted sum of the previously computed displacement and its projection in a given anisotropy vector. Table 1 sums up the main cell parameters and their meanings.


Attribute	Meaning
Color	RGB
Division Rate Mitosis	Mean time between divisions Controls the absolute and relative numbers of cells of a given type
Initial Probability	Probability to be of this type
Mutation Probability	Probability to switch to other type
Adhesion	Drag between types Controls the tendency of cells to stay together

Table 1: Attributes of a cell

For the sake of efficiency the relaxation step is not applied at every cell division event, but on a user controlled schedule.


This model and its simulation can produce an astonishing range of patterns closely resembling mammalian patterns, such as the spots of the giraffe, the stripes of the tiger (with anisotropy), the rosettes of the leopard and the jaguar (with three types of cells), etc.. A fundamental *motif* of these patterns is a Voronoi diagram structure, seen in its purest form in the reticulated giraffe (*Giraffa camelopardalis reticulata*) and shown in Figure 3. The simulation was computed in a square domain.

The similarity of the model and the real spots was verified with statistics on size distribution and closeness to an ideal Voronoi diagram (see [walt98]).



(a) Giraffe

(b) Real



(c) Computed

Figure 3: *Giraffa camelopardalis reticulata*




Figure 4: Giraffe body and 18 control cylinders

3.3 Controlling Shapes and Pattern Growing Parameters

With only a very small loss in generality, we assume that the patterns will be integrated with polygonal models of the animal bodies. Since the body has to be grown over a span of time with covers gestation up to adulthood, we need a flexible system to apply and control that growth, but at the same time we need to be able to deal with very sparse data. We will use the techniques described by Walter and Fournier in [walt97] and briefly summarize them here.

For each section of the body that will grow and generally be transformed independently, a cylindrical coordinate system is attached to it. The cylinder is positioned so that it encloses that part of the body it controls (this is not necessary, but is convenient for the user). These cylinders are predefined by the user, and are normally chosen so that they correspond to the available or obtainable growth data, as well as the need for animation. The cylinders are part of a classical modelling hierarchy. Figure 4 shows an example of the 18 cylinders initially defined for a giraffe's body

The main difference between a normal transformation hierarchy and the one used for growth is that since the growth data available is generally expressed in absolute terms, only translation and rotation are inherited from the ancestors in the growing hierarchy, but the scaling is local and in absolute terms. Each cylinder A has its own canonical coordinate system, and an associated matrix to transform to the world coordinate system:

$$M_{W \leftarrow A} = [T_{W \leftarrow A} R_{W \leftarrow A} G_{W \leftarrow A}]$$

where $T_{W \leftarrow A}$ is a translation matrix, $R_{W \leftarrow A}$ is a rotation matrix and $G_{W \leftarrow A}$ is a scaling matrix, the growth matrix. The growth matrix is given by:

$$G_A = \begin{bmatrix} L_A & 0 & 0 & 0 \\ 0 & R_A & 0 & 0 \\ 0 & 0 & R_A & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

where R_A and L_A are the world radius and length of the cylinder, which can be easily derived from such measurements as the length and girth of features of the real animal or its pictures (see for instance the features used for the giraffe fetus shown in Figure 1).

The operation to transform a point P_B in feature B (associated with a cylinder B) to a point P_A in the coordinate system of its parent feature A (associated with cylinder A) is:

$$P_A = [T_{A \leftarrow B} G_A^{-1} R_{A \leftarrow B} G_B] P_B$$

where:

$$T_{A \leftarrow B} = [M_{A \leftarrow W} T_{W \leftarrow B} T_{A \leftarrow W} M_{W \leftarrow A}]$$

One then applies these transformations to convert to the coordinate system of the ancestor of A , until the root is reached, at which point we have the point in world coordinate system. The growth process then consists in applying these transformations to each needed point for each time at which we have measures for the feature. The data can be interpolated between time using any suitable interpolating formula.

It is important to note that this has two practical consequences: any measurement from a real animal once applied to the polygonal model has the effect of changing the proportion of the original model (which could be an idealized or otherwise fictitious animal) to the proportional of the real measured animal. That means that for a single giraffe model we can create individual giraffe bodies with their own measurements.

The second point is that the same hierarchy can and will be used for animation, since we can independently apply rotation and translation to the various elements of the body.

4 Patterns, Shape and Growth

4.1 Simulation of the CM model on an arbitrary surface

Our models of animal bodies are polygonal meshes, and for convenience we restrict the polygons to be triangles (this does not limit generality, since any polygon can be triangulated in pre-processing). For the simulation we need to be able to initialize the cells, to subdivide the cells and to relax the cells. Initialization requires the ability to control the distance between cells. The relaxation steps require to compute distances as well, to obtain the forces and apply the resulting displacements. We also need to be able to compute a Voronoi diagram on the polygonal mesh, because the Voronoi cells together with the colour inherited from their cells form the tessellation of the surface defining the pattern. Note that as a result the pattern can be output as uniformly coloured convex polygons (the cells of the Voronoi diagram). Alternatively, to save space, if there is a very large number of cells per triangle, the original triangles of the mesh can be output with their colour pattern as a local texture.

The distance between two points on a polyhedral surface is measured along a sequence of straight line segments which intersect the polygons edges in such a way that they all aligned if the two polygons sharing an edge are rotated to have a common plane. We use this fact to compute distances by considering each triangle in turn and the triangles which share an edge or a vertex. We rotate the three triangles sharing an edge until they share a plane, and map the other triangles which share a vertex so that they are in the common plane and fill the angles between the first three triangles.

The distances between points in the first triangle can then be computed as two-dimensional distances in the common plane. We make the assumption that all the cells acting upon cells within a triangle are either within that triangle or within a triangle that shares an edge or a vertex with it. That is equivalent to stating that the shortest height of a triangle is longer than the effective radius of repulsion. We can guarantee that by pre-processing. The other similar assumption is that all the Voronoi neighbours of a cell contributing to the diagram within a triangle are either within the same triangle or within a triangle that shares an edge or a vertex with it. This can be guaranteed with high probability by making sure that there is a large (about 10) number of cells per triangle.




Figure 5: Voronoi diagram of cells generated directly on the polygonal mesh.



Figure 6: Giraffe pattern generated directly on the polygonal mesh.

The initialization involves computing randomly distributed points on the surface. We mostly use a uniform Poisson distribution, and a method described by Turk [turk91] and also used by Van Gelder and Wilhelm [geld97] to pick a triangle with a probability proportional to its area and a point within it uniformly distributed by area.

To relax the cells on the surface, for each triangle we transform the adjacent triangles and their cells as described above, and we compute the forces and the displacements in the common plane. Then with the displacement magnitude and direction known, we displace the cell in that direction. If the magnitude is such that the cell stays within its triangle, we are done. If not, we compute the intersection with the first edge, subtract the distance to that intersection from the displacement, and repeat the procedure in the adjacent triangle with the intersection as the starting point and the new displacement, and the direction the current direction rotated to put it in the plane of the new triangle.

To compute the Voronoi diagram of the cells after division and relaxation, we do it for each triangle and its neighbours, again using the assumption that all the Voronoi neighbours will be there, and we draw and store only the part of the Voronoi diagram which is inside the triangle. There is some redundancy involved, since the elements of the Voronoi diagram between cells on different triangles will be computed twice (or more for triangles sharing only vertices), but this is more than offset by the fact that triangles and their neighbours are in effect used as *buckets* to limit the search for neighbours. Figure 5 shows a close up of the Voronoi diagram on the giraffe body (as a triangular mesh).



Figure 7: Giraffe at birth and 11 months.



Figure 8: Giraffe at 22 months and 45 months (adult).

Figure 6 shows the result of growing a pattern directly on a giraffe body. Note that the spots are distributed over the whole body without visible discontinuities and are very similar to the pattern on the real giraffe or the two-dimensional simulation. The big difference is that the size of the spots is roughly the same over the whole body, contrary to what can be observed on the real giraffes.

4.2 Growth without pattern generation

The main point of integrating the growth and pattern processes is to create more accurate models of patterned animals. For our examples we applied transformations to data from the giraffe embryo (Figure 1), a newborn giraffe and an adult giraffe.

The polygonal model for the canonical giraffe body was obtained from the Viewpoint database (for more information about Viewpoint see www.viewpoint.com). For the purpose of illustrations we used a relatively simple models, which contains 308 vertices and 600 triangles. The real measurements were determined from photographs of real giraffes and averaged from few individuals.

Figures 7 and 8 show four steps in the growth of the body after the pattern is already established. The giraffe body is grown from the newborn proportions to the adult (about 5 year old). Figure 9 shows the first and last body (at birth and adult) in a close-up. It is important to note that if the pattern is close to a Voronoi diagram, and if some parts of the body grow anisotropically, such as the neck, whose length grows by a factor of 4 while its diameter grows by a factor of 3, then the final pattern cannot be a true Voronoi diagram. This can be verified on Figure 9 (even though it is not obvious without measurement, and more importantly this has been verified experimentally (these results will soon be submitted to a mathematical biology journal).




Figure 9: Giraffe at birth and adult with same pattern.




Figure 10: Fetus and cells at 35 days and 90 days.

4.3 Pattern generation with growth

The whole point of integrating the pattern formation and the growth is that the two are tightly linked in the fetal stage. Growth modifies distances (obviously) but also changes division rate. Following our basic assumption that cells areas (and volumes) remain essentially constant, in the absence of specific external drive, the division rate is entirely determined by the local area stretching, which is in turn determined by the determinant of the Jacobian of the local transformation. When controlled by cylinders, the local increase in area is the increase in radius multiplied by the increase in length.

Figures 10 and 11 show four phases in the development of the giraffe pattern on the fetus at 35 days (start of pattern development), 90 days, 150 days and 300 days. It should be compared to the pattern of the giraffe at birth.

Figure 12 shows the same fetus at 35 days, but put in a position similar to the one in Figure 1, for easier comparison (note that we




Figure 11: Fetus and cells at 150 days and 300 days.



Figure 12: Fetus and cells at 35 days in fetal position.



Figure 13: Fetus, new born and adult at same scale.

do not suggest that the pattern is actually that visible on the embryo, just that it is there).

5 Control of Parameters

To give a better idea of the real difference in size between the bodies, Figure 13 shows the fetus, the new born and the adult at the same scale. It actually changes by a linear factor of 100, from 3cm to 3m.

From our simulations it is clear that differences in growth rates alone do not explain the differences in spot sizes, for instance between the main body and the leg. Other factors, such as the anisotropy direction for the tiger and the specific markings on the face, show clearly that there are phenomena external to the CM system that affect the parameters. To deal with this, and to also be able to introduce arbitrary effects (such as writing our names with the spots of the cheetah), we introduced a way to control externally the parameters of the CM model. Since the cylinders are already in place to control the shape and growth, they are also convenient to provide support for textures which can control any of the parameters used in the CM system. Each cylinder can point to ordinary image files, with a range to map the texel value to parameters. It is of course convenient that the cylinder have a natural parameterization, being developable surfaces. For each point of the polygons, its cylindrical parameters are determined from the intersection of a line between the point and the point on the cylinder axis with the same X (coordinate along the cylinder axis) and the cylinder. This is mapped to the (0,1) range and used to index into the texture file. This is very similar to the technique of *two-pass texture mapping* described by Bier and Sloan [bier86].

Figure 14 shows the giraffe with the presence and the size of spots controlled by textures on its cylinders.

Figure 15 shows a striped tiger (actually elongated spots), where the orientation of stripes was controlled by textures on the cylinders.

Since the parameters for different patterns are part of the same set, it is very easy to interpolate between different patterns by interpolating between parameters. Of course there is no guarantee that the pattern will look to us as a legitimate intermediary between the two extreme. As a simple example, consider the pink panther and the pink tiger shown in Figure 16. The interpolated result is shown




Figure 14: Giraffe pattern controlled by textures.



Figure 15: Tiger anisotropy and colour controlled by textures.




Figure 16: Pink panther and pink tiger.



Figure 17: Interpolation between pink panther and pink tiger.

in Figure 17. It is clear in this case that the pattern does not yet exhibit much anisotropy, but the number of spots has visibly increased.

In this case we kept the same polygonal mesh for both bodies (the tiger), even though it is possible to morph the bodies as we morph the pattern. Table 2 shows the parameters used for these and the basic giraffe pattern. There ρ is the number of relaxation step per simulated day, wr the radius within which cells can cause a displacement, $time$ is the total simulated time (days), wd is the strength at which a cell repulses its neighbours, $mitosis$ are in average days between divisions, α is the adhesion (all cross-adhesions were set to 0.0), and the anisotropy is the length of the anisotropic vector (to be compared to wd). In the tiger images, the displacement is 80% in the anisotropy direction.

6 Conclusions

We presented in this paper a system which combines and integrates a pattern generation system which can effectively deliver a variety of patterns characteristic of mammalian coats, and a body growth and animation system that can use experimental growth data to produce individual bodies and their associated patterns totally automatically.

We used mainly the example of the giraffe, because the reticulated giraffe presents a pattern that can be objectively validated, and also because of the tremendous change in overall size and proportion between the fetus and the adult. We showed as well that the system can "do" most of the big cats, even though more detailed work is still needed in this respect.

The method also allows a considerable amount of user control through cylinders which control the growth, the animation and the parameters through textures.

There are many avenues left for extensions of this work. Immediately we want to use the texture control to generate realistic face markings (again the CM model won't be directly responsible for it).

Our current more long term goal is to extend the morphing capabilities of the system by using *generalized cylinders* instead of ordinary cylinders to better control the growth and the shape. This will also allow to define the shape parameters from real animals and pictures thereof in a more powerful way. Of course in mammalian coats the pattern is expressed as fur, and we plan to use our system to assign colour to simulated fur, as the final look very much depends on the *filtering* role of hair. The fur can be distributed on

Parameters	ρ	wr	$time$	wd	mitosis F	mitosis B	α FF	α BB	anisotropy
Giraffe (fig. 6)	18	2.6	50	0.066	10	120	0.9	0.2	0.0
Pink panther (fig. 16)	18	2.6	10	0.066	10	120	0.9	0.2	0.0
Pink tiger (fig. 16)	18	1.1	80	0.01	10	50	0.7	0.2	0.5
Interpolation (fig. 17)	18	1.85	45	0.038	10	85	0.8	0.2	0.25

Table 2: Parameters for some of the computed patterns.

our polygonal models in a manner similar to the one described in work by Van Gelder and Wilhelms [geld97]. In addition we have a method to control the behaviour of the entire pelage while allowing individual motion to single hairs.

And to end with the obvious, a more thorough exploration of the parametric space (maybe *a la* Design Galleries [mark97] for the clonal mosaic model will certainly be rewarding.

7 Acknowledgment

To be added later, naming all the little people who helped us on our way to the (relative) top.

References

- [bard81] J. B. L. Bard. “A Model for Generating Aspects of Zebra and Other Mammalian Coat Patterns”. *Journal of Theoretical Biology*, Vol. 93, No. 2, pp. 363–385, November 1981.
- [bedd06] E. Beppard. “Description of the External Characters of an unborn Fetus of a Giraffe”. *Journal of Zoology*, pp. 626–631, 1906.
- [benn91] C. Bennis et al. “Piecewise surface flattening for non-distorted texture mapping”. *Computer Graphics (SIGGRAPH '91 Proceedings)*, Vol. 25, pp. 237–246, July 1991.
- [bier86] E. A. Bier and K. R. Sloan, Jr. “Two Part Texture Mappings”. *IEEE Computer Graphics and Applications*, Vol. 6, No. 9, pp. 40–53, September 1986.
- [catm74] E. E. Catmull. *A Subdivision Algorithm for Computer Display of Curved Surfaces*. Ph.D. Thesis, University of Utah, December 1974.
- [fari90] G. Farin. *Curves and Surfaces for Computer Aided Geometric Design*. Academic Press, 1990.
- [fle95] K. Fleischer, D. Laidlaw, B. Currin, and A. Barr. “Cellular Texture Generation”. *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pp. 239–248, August 1995.
- [fowl92] D. R. Fowler, H. Meinhardt, and P. Prusinkiewicz. “Modeling seashells”. *Computer Graphics (SIGGRAPH '92 Proceedings)*, Vol. 26, No. 2, pp. 379–387, July 1992.
- [geld97] Allen Van Gelder and Jane Wilhelms. “An interactive fur modeling technique”. *Graphics Interface '97*, pp. 181–188, May 1997.
- [gord78] R. Gordon and A. G. Jacobson. “The Shaping of Tissues in Embryos”. *Scientific American*, pp. 106–113, June 1978.
- [hanr90] P. Hanrahan and P. Haerberli. “Direct WYSIWYG Painting and Texturing on 3D Shapes”. *Computer Graphics (SIGGRAPH '90 Proceedings)*, Vol. 24, No. 4, pp. 215–223, August 1990.
- [litw94] P. Litwinowicz and G. Miller. “Efficient Techniques for Interactive Texture Placement”. *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pp. 119–122, July 1994.
- [mail93] J. Maillot, H. Yahia, and A. Verroust. “Interactive Texture Mapping”. *Computer Graphics (SIGGRAPH '93 Proceedings)*, Vol. 27, pp. 27–34, August 1993.
- [mark97] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. “Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation”. *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pp. 389–400, August 1997.
- [murr81a] J. D. Murray. “On Pattern Formation Mechanisms for Lepidopteran wing patterns and mammalian coat markings”. *Philosophical Transactions of the Royal Society of London B*, Vol. 295, No. 1078, pp. 473–496, October 1981.
- [murr81b] J. D. Murray. “A Pre-pattern Formation Mechanism for Animal Coat Markings”. *Journal of Theoretical Biology*, Vol. 88, pp. 161–199, 1981.
- [okab92] A. Okabe, B. N. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley & Sons, 1992.
- [owen49] R. Owen. “Notes on the Birth of the Giraffe at the Zoological Society Gardens”. *Transactions of the Zoological Society of London*, Vol. 3, pp. 21–28, 1849.
- [prus93] P. Prusinkiewicz. “Modeling and visualization of biological structures”. *Proceedings of Graphics Interface '93*, pp. 128–137, May 1993.
- [turi52] A. M. Turing. “The Chemical Basis of Morphogenesis”. *Philosophical Transactions of the Royal Society of London B*, Vol. 237, pp. 37–72, 1952.
- [turk91] G. Turk. “Generating textures on arbitrary surfaces using reaction-diffusion”. *Computer Graphics (SIGGRAPH '91 Proceedings)*, Vol. 25, No. 4, pp. 289–298, July 1991.
- [walt97] M. Walter and A. Fournier. “Growing and Animating Polygonal Models of Animals”. *Computer Graphics Forum (Eurographics '97)*, Vol. 16, No. 3, pp. C–151–C–158, September 1997.

- [walt98] M. Walter, A. Fournier, and M. Reimers. “Clonal Mosaic Model for the Synthesis of Mammalian Coat Patterns”. *Graphics Interface*, pp. 82–91, June 1998.
- [witk91] A. Witkin and M. Kass. “Reaction-diffusion textures”. *Computer Graphics (SIGGRAPH '91 Proceedings)*, Vol. 25, No. 4, pp. 299–308, July 1991.

Integrating Shape and Pattern in Mammalian Models

Marcelo Walter
Department of Computer Science
University of British Columbia
2366 Main Mall, Vancouver, BC, V6T 1Z4
Tel.: (604) 822-2218
Fax.: (604) 822-5485
Alain Fournier
Daniel Meneveaux
{marcelow/fournier/danielm}@cs.ubc.ca

Category: research

Format:

Contact: Alain Fournier
Dept. of Computer Science
The University of British Columbia
2366 Main Mall, Vancouver, B.C.
CANADA V6T 1Z4
phone: (604) 822-6770
fax: (604) 822-5485
email: fournier@cs.ubc.ca

Estimated # of pages:

Keywords: integration, animal models, animal patterns, clonal-mosaic, growth, animation

The giraffe and its patches, the leopard and its spots, the tiger and its stripes are spectacular examples of the integration of a pattern and a body shape. We propose in this paper a system which combines and integrates a pattern generation system which can effectively deliver a variety of patterns characteristic of mammalian coats, and a body growth and animation system that can use experimental growth data to produce individual bodies and their associated patterns totally automatically.

We use the example of the giraffe to illustrate how the models take us from a canonical embryo to a full adult giraffe in a continuous way, with results that are not only realistic looking, but also objectively validated. The flexibility of the system is demonstrated by examples of several big cat patterns, including an interpolation between patterns.

The method also allows a considerable amount of user control to generate *un-natural* patterns and/or shapes and to animate the resulting body with the pattern.