

The leopard never changes its spots: realistic pigmentation pattern formation by coupling tissue growth with reaction-diffusion

MARCELO DE GOMENSORO MALHEIROS, FURG, Brazil

HENRIQUE FENSTERSEIFER, UFRGS, Brazil

MARCELO WALTER, UFRGS, Brazil

Previous research in pattern formation using reaction-diffusion mostly focused on static domains, either for computational simplicity or mathematical tractability. In this work, we have explored the expressiveness of combining simple mechanisms as a possible explanation for pigmentation pattern formation, where tissue growth plays a crucial role. Our motivation is not only to realistically reproduce natural patterns but also to get insights into the underlying biological processes. Therefore, we present a novel approach to generate realistic animal skin patterns. First, we describe the approximation of tissue growth by a series of discrete matrix expansion operations. Then, we combine it with an adaptation of Turing's non-linear reaction-diffusion model, which enforces upper and lower bounds to the concentrations of the involved chemical reagents. We also propose the addition of a single-reagent continuous autocatalytic reaction, called reinforcement, to provide a mechanism to maintain an already established pattern during growth. By careful adjustment of the parameters and the sequencing of operations, we closely match the appearance of a few real species. In particular, we reproduce in detail the distinctive features of the leopard skin, also providing a hypothesis for the simultaneous productions of the most common melanin types, eumelanin and pheomelanin.

CCS Concepts: • **Computing methodologies** → **Computer graphics; Modeling and simulation.**

Additional Key Words and Phrases: natural phenomena, texturing, pattern formation, reaction-diffusion, Turing model

ACM Reference Format:

Marcelo de Gomensoro Malheiros, Henrique Fensterseifer, and Marcelo Walter. 2020. The leopard never changes its spots: realistic pigmentation pattern formation by coupling tissue growth with reaction-diffusion. *ACM Trans. Graph.* 39, 4, Article 63 (July 2020), 13 pages. <https://doi.org/10.1145/3386569.3392478>

1 INTRODUCTION

Growth is an essential mechanism of life and a continuous process since the inception of all living beings. Although much research has been done on cellular biochemistry, the overall mechanisms responsible for growth have been only partially uncovered by now, and are already very complex [Wolpert et al. 2015].

Authors' addresses: Marcelo de Gomensoro Malheiros, FURG, Computer Science Center, Rio Grande, Brazil, mgmalheiros@gmail.com; Henrique Fensterseifer, UFRGS, Institute of Informatics, Porto Alegre, Brazil, rique.hf@gmail.com; Marcelo Walter, UFRGS, Institute of Informatics, Porto Alegre, Brazil, marcelo.walter@inf.ufrgs.br.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

0730-0301/2020/7-ART63 \$15.00

<https://doi.org/10.1145/3386569.3392478>

Reaction-diffusion is a well-known model conceived by Turing [1952], being later used in Computer Graphics by the pioneering works of Turk [1991], Witkin and Kass [1991] and Fowler et al. [1992] for simulating natural phenomena. Although interest in reaction-diffusion faded along the years for the graphics community, it continued to play a major role in theoretical models from Mathematical Biology [Maini et al. 2012; Meinhardt 2009; Murray 2003].

We are interested in the general problem of reproducing the appearance of animal skin patterns. Instead of focusing on procedural methods [Hu et al. 2019], exemplar-based techniques [Raad et al. 2018] or indirectly through machine learning [Zhou et al. 2018], we show that a reduced subset of simulated biological mechanisms is enough to generate accurate textures.

Our main contribution is showing that the adequate approximation of tissue growth combined with reaction-diffusion is the key to the emergence of complex yet realistic patterns. By carefully adjusting model parameters and setting up distinct growth phases, we can match skin pigmentation of mammals and other animals. In particular, we have produced *in silico* the characteristic leopard rosettes, depicted in a 3D rendering in Figure 1.



Fig. 1. Synthetic leopard coat: pigmentation generated by our technique and then rendered in 3D by assigning individual fur colors. The skin below the fur layer has a uniform pink color and it is barely visible.

2 RELATED WORK

The seminal work of Turing [1952] provided a plausible yet controversial hypothesis for morphogenesis, that is, the biological process that drives an organism to develop its shape and appearance. In Biology, a pattern refers to any visual or spatial organic structure, hence the development of life forms happens through many pattern formation processes. In this paper, we use the term more strictly, where **pattern** describes the variation of color over a surface.

Turing proposed an abstract chemical reaction, which could give rise to coherent patterns, even when starting from a random initial state. This mechanism was called **reaction-diffusion** (RD), describing the interaction of two **reagents** (chemical substances) and being modeled by a pair of differential equations.

Two papers introduced RD for Computer Graphics. Turk [1991] used a RD model adapted from Bard [1981], which in turn was derived from Turing's original non-linear equations. Turk achieved a variety of results by applying two phases to the simulations. First, the model was run with a set of parameters creating a base pattern. Then, a second simulation was performed with different parameters, while freezing part of the base pattern. Witkin and Kass [1991] provided parameterized geometric distortion and anisotropic patterns by varying the diffusion rates along the x and y dimensions.

The Gray-Scott model was presented by Pearson [1993], which produced complex pattern dynamics. Still, the fundamental visual **features** were the same as in other RD models, that is, spots and labyrinths. Later, Sanderson et al. [2006] showed that applying anisotropic diffusion and spatially modulating parameters produced controllable patterns. They generated images similar to fish skin pigmentation and were able to constrain patterns into desired shapes.

Many CG works already addressed the close relationship between form and texture, as in the traditional problem of texturing surfaces of arbitrary topology [Knöppel et al. 2015]. Few works pursued the link between dynamic growth and appearance. Walter et al. [2001] showed that cell division driven by growth produced the characteristic pattern of a growing giraffe and the big cats. Albeit distinct from a RD model, its basis was the discrete simulation of cell proliferation as the focus of the growth processes, running over an expanding 3D surface. Runions et al. [2005] described a model that incorporates growth to describe geometric leaf vein development.

Point-based methods are fairly common for applications like fluid or crowd simulations, but there are fewer works that address biologic patterns as the result of the interaction of discretized cells.

Fleischer et al. [1995] proposed a rule-based cellular system, focused on the modeling of surface details such as scales, feathers and thorns. Kider Jr et al. [2011] provided a reaction-diffusion model to generate growth patterns for areas of fungal and bacterial infection, changing both the visual and geometric properties of synthetic fruits. Volkening and Sandstede [2015] created a complex simulation to account for the formation of stripes in zebrafish, using direct cell signaling. Malheiros and Walter [2017] proposed a generic simulation system, implementing reaction-diffusion among a close arrangement of unorganized cells on a 2D domain. Recently, Gingras and Kry [2019] developed a new procedural 3D shape modeling approach, where a RD simulation affects the local expansion and bending of the surface of a growing thin shell.

3 OVERVIEW

Our focus is on using reaction-diffusion as a general technique for producing patterns.

For that, we have developed a particular RD model that provides more expressibility than prior works, by coupling it with an approximated tissue growth scheme.

We then systematically searched its parameter space, being able to closely match pigmentation patterns found in nature.

A piece of two-dimensional skin tissue can be approximated by a rectangular matrix, where each matrix element roughly describes a single biological cell. This representation is both compact in terms of memory usage and very convenient when locating nearest neighbors, as nearby cells are directly found through indexing access.

However, when the tissue grows, the cell division process has to be simulated, which implies creating new cells among the already existing ones. This is usually done through point-based methods, following the analogy to a dense yet unorganized set of cells as part of real living tissue, typically resembling Voronoi polygons.

A major bottleneck of point-based techniques is their dependence on Nearest Neighbor Search (NNS) algorithms, as locating nearest cells must be performed for each cell, at each simulation step. Such algorithms are computationally intensive, even more for evolving domains, represented by a dynamic set of moving points. Though using specialized data structures like k -d trees, NNS still takes a significant portion of the simulation time [Weaver and Xiao 2016].

In our work, we take a different direction. We approximate the continuous development of a tissue by a series of discrete matrix expansion operations, which randomly insert new matrix elements akin to biological cell division. However, after one growth step, we still have the domain represented as a rectangular matrix, therefore keeping the benefits of having a fast and coherent data structure. In Section 4 we describe this scheme.

Although having developed this matrix expansion technique independently, we have later found the work of Binder et al. [2008], which also performs modification on a matrix. Differently, the authors employ cellular automaton rules, using only fixed-state cells and with no reaction-diffusion mechanisms taking place.

In Section 5, we describe the pigment formation process, employing reaction-diffusion over an expanding domain. The RD model employed in this paper is based on the equations from Bard [1981], subject to concentration limits and further modified to have more intuitive parameters to adjust pattern behavior and scale. In Section 6 we address the need for a mechanism to maintain the pattern during growth, developing a novel single-reagent autocatalytic model. Pattern maintenance is a topic first raised by Fowler et al. [1992], and for that, we generalize the discrete reinforcement mechanism discussed by Malheiros and Walter [2017].

Our results are discussed in Section 7. To achieve pattern diversity, instead of discretely altering the pattern or spatially controlling the parameters as in prior works, we opted to extensively explore the parameter space of our RD model and only employ domain expansion during the simulation.

In Section 8 we present our contributions, discuss current limitations and present future research opportunities.

4 TISSUE GROWTH

Growth is a fundamental process for any life form. Although biologic growth mechanisms are still not completely understood, they are known to be carefully regulated processes. We are particularly interested in approximating skin tissue growth.

4.1 Matrix expansion

Given a rectangular matrix with m rows and n columns, we can perform a **matrix expansion** by simply adding a new row or column. Figure 2 shows the result of expanding a matrix in the horizontal dimension, by adding a new column. Likewise, this operation may be performed in the other dimension, thus adding a whole new row of elements (here representing discrete biological cells).




Fig. 2. Column expansion. (a) Initial state, where one cell from each row is randomly chosen (marked with a strong outline). (b) After the selected cells are subdivided, with new cells inserted on their right side (all black cells), shifting the rightmost cells.

To approximate tissue growth, this expansion is performed in two passes: first creating a new row and shifting all cells accordingly in one dimension; then, a new column is created, shifting again part of the cells in the other dimension. After the division the parent and child cells are identical, so there is no practical difference in considering the insertion of the newborn cell on the right or left of its parent cell. Although this is a very rough approximation of the division of biologic cells, it is an adequate method for the domain expansion of a reaction-diffusion simulation, as discussed in Section 5. A timing comparison with simulations using point-based Nearest Neighbor Search algorithms is made in Appendix A.

4.2 Random cell division

The indices for cell division can be picked following any particular probability distribution. Typically, a uniform random distribution is used. In Figure 3 we show the outcome of applying different random generators for the location of new cells. Alternatively, we might have separate control by employing a distinct distribution for each dimension. We have statistically checked that shuffling row and column indices is equivalent to using a uniform distribution, so this latter approach is used in most experiments.

The implementation of these expansion operations is straightforward when applied to a rectangular matrix: values are simply copied into a larger memory area, adding a row or column offset after the indices of the divided cell. This is particularly useful because our simulation employs two matrices to represent the concentration of each reagent: the current and the next state, thus the expansion does not need to be done in-place.




Fig. 3. Distinct strategies when selecting row and column indices. The matrix starts with size 100 by 100, with a single dark color, and is expanded until it reaches 200 by 200. Lighter colors indicate subdivided cells added later. In the top row (from left to right): shuffling row and column indices, uniform distribution, clipped normal distribution and triangular distribution. In the bottom row, we show the associated index distribution histograms.

5 PIGMENT FORMATION

The process of pigment formation using reaction-diffusion is driven by a pair of parameterized differential equations. Such a system represents an idealized auto-catalytic chemical reaction involving two reagents, here named A and B.

Although there are many published reaction-diffusion models, we developed a variation of Turing's non-linear RD model because it has a simple implementation and provides efficient numerical computation. Furthermore, this model has been already addressed a few times in the literature, which provided in-depth analysis.

We made slight modifications to Turing's model to reduce the number of parameters, making the exploration of its parameter space more tractable. Then, we have provided an alternative set of artist-oriented parameters, which makes fine-tuning the generated results easier. Last, we also added configurable upper and lower bounds to the concentrations of the reagents, which greatly expands the model generational expressiveness.

5.1 Original model

Turing [1952] (page 65) proposed a particular non-linear reaction-diffusion model as the basis for pattern formation. Bard and Lauder [1974] were the first to computationally explore it, slightly modifying the original terms. Their model is shown as Equations 1 and 2, where a and b express the chemical concentrations of the two diffusible reagents A and B, respectively.

$$\frac{\partial a}{\partial t} = S_p(16 - ab) + D_a \nabla^2 a \quad (1)$$

$$\frac{\partial b}{\partial t} = S_p(ab - b - \beta) + D_b \nabla^2 b \quad (2)$$

Reaction-diffusion equations have a **reaction** part and a **diffusion** part. The reaction is interpreted as the system feedback, which either produces or consumes a given reagent based on the current concentrations of both reagents. In Equation 1, for example, the reaction part is given by $S_p(16 - ab)$, whereas the diffusion part is given by $D_a \nabla^2 a$. Diffusion accounts for the reagent exchange

among nearby locations within the two-dimensional domain (which here is discretized into a rectangular matrix). The diffusion rates are given by D_a and D_b , indicating how fast the diffusion happens for each reagent. Bard and Lauder [1974] also proposed the addition of an S_p parameter over Turing's original system to adjust the intensity of both reaction parts, which remained in ensuing literature.

5.2 Proposed model

We performed a reparameterization of the original model, to both reduce the number of parameters and to better control the generated patterns. Our model is shown in Equations 3 and 4.

$$\frac{\partial a}{\partial t} = 16 - ab + rs\nabla^2 a \quad (3)$$

$$\frac{\partial b}{\partial t} = ab - b - 12 + s\nabla^2 b \quad (4)$$

We have fixed the β parameter as 12. Originally, the parameter β would vary locally, having a slight fluctuation around the fixed value of 12. This was the source of randomness that made patterns appear, as the initial concentrations of both reagents were all set to 4. Choosing a different set of β values would thus result in a distinct pattern. We opted to use the simpler approach of making $\beta = 12$, where the variations are driven by controlled randomness in the initial concentrations of A or B. This approach of having fixed equation parameters and starting with varying initial conditions can be traced back to the work of Gierer and Meinhardt [1972], who proposed other well known RD systems.

We have also replaced D_a and D_b by more intuitive parameters, defining **ratio** $r = D_a/D_b$ and **scale** $s = D_b$. That is, we decouple the actual pattern development behavior (expressed by r) from the overall pattern size (indicated by s). These parameters can then be adjusted independently, providing finer control. As far as we could verify, prior RD research focused on obtaining distinct patterns by varying the reaction coefficients and keeping the diffusion rates fixed, as in the work of Miyazawa et al. [2010]. Differently, we assume that the constants 16 and -12 are immutable (being directly linked to an abstract chemical reaction) whereas diffusion rates might be biologically controlled by cellular mechanisms.

Finally, we have found the usage of S_p confusing, because it affects both the perceived scale of the pattern and the velocity of the simulation. As we can directly adjust the pattern scale through s , we have fixed $S_p = 1$, thus removing it from our model.

5.3 Model implementation

As the diffusion follows Fick's second law, we used a numerical implementation of the discrete Laplacian ∇ operator over the nearest matrix elements. We used as integration scheme the standard nine-point stencil, covering a Moore neighborhood of unitary size, representing isotropic diffusion¹.

For the numerical integration, we have evaluated several alternatives and settled on the simplest approach: first-order forward Euler integration [Sanderson et al. 2006]. We employed a fixed time step Δt during the simulation, which multiplies both the reaction and diffusion parts of the model equations. The default value was 0.01, being

¹That is, $[[1, 4, 1], [4, -20, 4], [1, 4, 1]]/6$.

decreased when numerical instabilities appeared. Higher-accuracy methods gave similar results, in spite of more computational cost.

All paper figures exhibit the concentration of one reagent mapped onto a linear color map². Unless noted, the B reagent is shown. The parameters and timings are listed in Appendix A.

A reference implementation is also publicly available³.

5.4 Upper and lower bounds for concentrations

Similar to other reaction-diffusion models, patterns generated by our model exhibit a self-adjusting behavior, being resilient to added noise and perturbations. The consistent size of features is traditionally called the pattern **wavelength**.

Malheiros and Walter [2017] proposed the addition of upper bounds for the concentrations of reagents A and B, noting that such constraints added more diversity to the resulting patterns. We employed the same mechanism, defining the upper limit parameters U_a and U_b , which are used to clamp the associated concentrations at the end of each simulation step. Although there is not yet either empirical or experimental evidence of that principle in a biological process, we observed that the usage of upper bounds breaks the wavelength maintenance of the RD model, yielding asymmetric Turing patterns that matched skin pigmentation for a few species.

Reaction-diffusion models assume that reagent concentrations are non-negative. Some models accomplish this intrinsically, like the Gray-Scott model [Pearson 1993] whereas others have to explicitly define a zero lower bound by clamping negative values, as originally done by Bard and Lauder [1974].

We have found that also setting this lower limit for values higher than zero introduces another kind of diversity in the generated patterns, so both lower bounds are parameterized by L_a and L_b , and enforced at the end of each simulation step. In fact, by imposing these lower bounds we can achieve a similar effect as changing the 16 and -12 constants, although in a more controlled way and still keeping the reaction part of our RD model unchanged.

Note that Equations 3 and 4 alone are not enough to generate a stable pattern, having the added requirement of B being non-negative. This was achieved as an extra step of the numerical integration calculation, which also enforced the lower and upper bounds for concentrations. This proved to be very efficient, providing the quick cycles of experimentation we needed.

However, we understand that a proper system of Partial Differential Equations (PDE) might be needed in other contexts, and thus briefly outline here an equivalent model. Both the lower L_a and upper U_a bounds for A can be explicitly enforced by adding rational terms to Equation 3. Provided the initial values for a are within the (L_a, U_a) interval, we may add a $\frac{1}{a^k - L_a^k}$ term for imposing its lower bound, and $\frac{1}{U_a^k - a^k} + \frac{1}{U_a^k}$ term for enforcing its upper bound. Likewise, similar terms would be added to Equation 4. We have found that $k = 2$ is enough to provide compatible results to our model, that is, given the same parameters a very similar pattern would be generated. In this case, the time step must be reduced by a factor of 20 to maintain numerical stability. Domain expansion, however, is still performed as a separate step (see Section 8.2).

²We use the *inferno* perceptually-uniform color map from the Matplotlib library.

³<https://github.com/mgmalheiros/reaction-diffusion>

5.5 Effect of parameters

A sample of patterns obtained by varying the ratio parameter r in a static domain is shown in Figure 4. Three situations are shown: standard behavior in the top row, with no imposed upper bounds, but still restricting B to non-negative values. In the middle row, the upper bound U_a is also set, thus demonstrating the break on the wavelength and appearance of large constant regions. In the bottom row, just the lower bound L_b is set, with positive values, thus altering the pattern features.

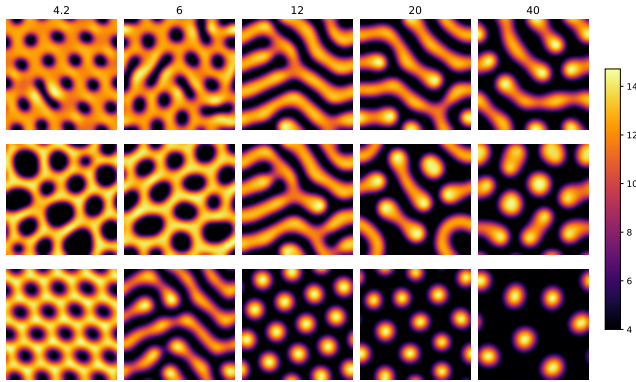


Fig. 4. Example patterns for a static domain, with matrix size 50 by 50. Representative ratio values r are set for each column. In the first row, the typical behavior is shown. In the middle row, only the upper bound U_a is set, while in the bottom row, only the lower bound L_b is set.




Fig. 5. Example patterns for a growing domain, with initial size 50 by 50 and final size 100 by 100. Representative ratio values r are set for each column. The first row shows normal behavior. In the middle row, only the upper bound U_a is set, whereas in the bottom row, only the lower bound L_b is set.

If we take the same patterns as in Figure 4 and then apply domain growth, from an initial matrix size of 50 by 50 to a final size of 100 by 100, we generate the patterns shown in Figure 5. A discussion about initial states for the simulations can be found in Section 7.1.

Suppose we have a rectangular domain where we have already performed many steps of a reaction-diffusion simulation. A reference pattern can be seen on the left column of Figure 6.




Fig. 6. Pattern development during domain growth. Top row with no reaction nor diffusion, middle row with standard reaction-diffusion and bottom row with saturated RD. Columns from left to right: initial pattern (size 100 by 100), after one matrix expansion, after ten expansions and after 100 expansions (size 200 by 200). Iteration counts are shown on top of each pattern.

Here we define a **growth rate**, establishing how many expansions per second should occur during the simulation. For the rows of Figure 6, we have uniform expansion happening every second, equivalent to every 100 simulation steps for $\Delta t = 0.01$.

In a situation with no reaction nor diffusion, after the application of a matrix expansion operation, cells divide simultaneously at random positions, thus duplicating the values of matrix elements. Repeating a few times this operation would result in a noisy image. The sequence is shown in the top row of Figure 6. Naturally, this is not desirable: although the overall pattern has increased in size to match the larger domain, the original pattern would be severely affected by noise, destroying the small feature details.

We may, however, continue to run the reaction-diffusion equations while we grow the domain. Following a biologic analogy, we could expect chemical diffusion to be a much more rapid process than the actual growth of living tissue, even during embryonic cell proliferation. This is shown in the middle row of Figure 6. Time-spaced expansions allow the pattern to gradually adapt to the new domain, effectively absorbing the noise introduced by the random cell divisions. We can see that the final pattern keeps its overall features and smoothness, despite fitting into a larger domain. The number of features has increased, each with the same wavelength.

If we enforce an upper bound to the concentration of A , it is possible to constrain the pattern development during growth, reducing the emergence of new features. One such example is shown in the bottom row of Figure 6, where the initial dark regions increase in size when we set $U_a = 5.75$, evolving into a reticulate of large polygons. As the RD process is still active, some regions may be fused into the larger ones, but the overall reticulate pattern is kept as the domain grows. This equilibrium is not static, however, so the pattern can still absorb the irregularities added by the expansion process.

6 PATTERN ENLARGEMENT

We have initially expected that just coupling reaction, diffusion and cell division would lead to consistent patterns, but the dynamic nature of RD equations always disrupts an established shape. Even if we imposed upper bounds to one or both reagents, an initial pattern would still get affected by growth, where small features could disappear or new ones emerge in the middle of large constant areas (the so-called *tip splitting* phenomenon). Because the RD process is still operating, new local dynamics are continually introduced into an established pigmentation pattern as the tissue grows.

When examining patterns from large mammals with contrasting fur coatings, like zebras or cows, we can see that the adult exhibits large uniform areas of either black or white color. The borders between these colors are usually sharp, whereas the contour themselves are often curved shapes. On most animals, there are no visible irregularities on these color transitions. But the overall smooth shapes and relative size of patterns have been mostly kept since the newborn animal. So it seems that a more complex mechanism of pattern maintenance is still missing.

A regulation mechanism was described by Malheiros and Walter [2017], designed to smooth pattern borders but also keep the interior of constant regions stable, called **reinforcement**. That process was rule-based, working on a single reagent, and continually altered its concentrations towards either a lower or an upper value. We have designed a cubic polynomial function that can provide an equivalent, yet continuous, feedback as shown in Equation 5.

$$f(c) = \gamma(t - w - c)(t - c)(t + w - c) \quad (5)$$

The production and the consumption of the reagent C can be seen as a function of the current concentration c , where the three real roots are placed at $t - w$, t and $t + w$. The idea is to have negative feedback for concentrations of c in the interval $(t - w, t)$, and positive feedback for concentrations of c in the interval $(t, t + w)$. If we let $\gamma = \frac{3\sqrt{3}}{2w^2}$, we have $-w$ and w at the local minimum and maximum, respectively. This function is plotted in Figure 7.




Fig. 7. Plotting of reinforcement reaction term for $t = 1$ and $w = 1$. Roots are marked as red circles, while the local minimum and maximum are marked by downward and upward green triangles, respectively.

Setting γ as indicated is important because it provides invariance to both changes in the threshold value t and the interval width w . That is, if concentration values that define a pattern are offset or scaled, the respective change on t and w will produce the same visual result after the reinforcement mechanism is applied. Note also that the curve slope for $c < w - t$ and $c > w + t$ keeps the

concentration in the $[t - w, t + w]$ interval, without the explicit enforcement of upper and lower limits.

This function can be thought of as the reaction part of a single-reagent autocatalytic equation. The diffusion part could be represented by the spatial Laplacian operator ∇^2 multiplied by the D_c diffusion rate, as shown in Equation 6. This equation bears a strong analogy to RD systems, albeit with a simpler behavior.

$$\frac{\partial c}{\partial t} = \gamma(t - w - c)(t - c)(t + w - c) + D_c \nabla^2 c \quad (6)$$

Suppose that before this mechanism is started we already have an initial pattern defined by the concentrations of reagent C. During the reinforcement process, diffusion of C occurs normally. For each simulation step, Equation 6 is applied. The production and consumption intervals oppose the diffusion effect, reinforcing the regions with high and low concentrations, respectively. This keeps the border between those areas thin. The major point is: if tissue growth occurs, the overall shape of regions is maintained, despite some minor detail being lost. For example, in Figure 8 we apply the same growth as in Figure 6. If we also adjust the D_c diffusion rate, we can alter the range of the smoothing effect on borders.




Fig. 8. Growth with continuous reinforcement. From left to right: initial pattern (size 100 by 100), after one matrix expansion, after ten expansions and after 100 expansions (size 200 by 200). Iteration counts shown on top.

Note that the diffusion does not affect the large light or dark areas, but it is actively smoothing-out borders, even when new cells are created. This smoothing is countered by the reinforcement mechanism, which tries to enhance the borders, like an image processing high-pass filter.

7 RESULTS

In this section, we show our best results so far, focused on the faithful reproduction of the skin pigmentation patterns found on living animals from several species.

Following the usual approach for simulating reaction-diffusion, the domain is discretized into a regular square lattice, with m rows and n columns. For every discrete cell there are two associated reagent concentrations, called a and b . The implementation stores these values in two equally-sized $m \times n$ matrices of floating-point values. Note that we name all concentrations for the first reagent as A, and B for the second. For the leopard experiment, discussed in Section 7.4, we add a third reagent C, represented by a third matrix.

Initial concentrations are set for the matrices. The simulation then starts, performing many iteration steps. At the beginning of each iteration we numerically integrate Equations 3 and 4, as discussed in Section 5.3. After that we apply the concentration limits to the updated concentrations, as in Section 5.4. Finally, when mandated by an established growth rate, the expansion process is equally applied

to all matrices, so that the same randomly selected matrix elements are divided, following the operation outlined in Section 4.

Note that the experiments go through one or more phases, with model parameters kept the same during each phase. Moreover, the final state of a phase is the same initial state of the following phase. We use the term **prepattern** to describe any such initial state.

We have mapped the final concentration of a single reagent into an interpolated color map, designed to match the overall colors of living individuals illustrated in the photos. No further alteration was made. Iteration counts are shown on top of the images.

7.1 Initial state

It is important to carefully address the initial state of the experiments. Because it is crucial to the generated outcome, we defend that it must be based on a biologically plausible hypothesis. The major difficulty for specifying the initial state is that most of the available data is visual information about the skin of adult individuals, although many pattern formation studies, like the ones surveyed by Murray [2003], point to the definition of the visual appearance of mammals in the early stages of development on the womb.

Therefore, we employ only two types of initial states, setting up starting concentrations close to the levels exhibited after the patterns form and stabilize.

The first one is a constant concentration for A and a random distribution with a small variation for B, approximating an almost constant chemical state, as first suggested by Turing [1952]. For most patterns shown in this paper, the initial value for A was 4, whereas B values were uniformly distributed in the [4, 5] interval.

As a second type, we also explored the idea of a completely constant initial state, with A and B set to 4. For the emergence of patterns in this situation, randomness will be later introduced during the simulation by a localized addition of reagent B, as discussed in the squirrel and leopard experiments. The biologic analogy is straightforward: specialized organelles would start to produce specific substances, which then diffuse and react accordingly as they spread over the organism. This follows the hypothesis formulated by Bard [1981] that the *neural crest* could be the initiator for the pattern formation process in many mammals.

Further randomness is added by the cell division process, approximated by the matrix expansion operations. By changing the initial random seed we can produce variations of the resulting patterns, in the way individuals from the same species have distinct skin patterns yet the same overall visual features.

7.2 Exploratory approach

As our implementation was geared to interactive experimentation, we have systematically explored variations of the parameter set, focusing first on mapping the overall patterns possible, and then manually adjusting parameters to match a particular skin pattern.

Many 2D parameter maps were automatically generated to cover parameter sets, usually varying two of them in the horizontal and vertical axis, similar to Figures 5 and 6. For our model, the most important parameter is the ratio r , which defines the overall structure of the pattern. As needed, the lower bound L_b was altered from the default zero value, as in Section 5.4, to induce a pattern structure.

Then, given a base pattern, we mapped the effects of varying the upper bound U_a , responsible for breaking the feature wavelength. We also adjusted the growth rate, which induced increased pattern complexity. Finally, interesting patterns close to already known living species were selected and manually adjusted.

An account for the discovery of the frog pattern is made in the Supplementary Material.

7.3 Reaction-diffusion experiments

Starting with random initial concentrations, we could reproduce the intricacies of skin pigmentation from the reticulate whipray (*Himantura uarnak*), shown in Figure 9. To closely match that individual, we found that two stages were needed. A first stage is run on a static domain with $r = 30$, $s = 3$ and $L_b = 2$, starting from random values and resulting in equispaced dark spots. A second stage maintains all but one parameter, where the ratio is now changed to $r = 8$ and uniform growth occurs, at 2.5 expansions per second.




Fig. 9. Top row: static domain with size 100 by 100, using $r = 30$. Middle row: growing domain, now with $r = 8$, reaching final size of 200 by 200. Bottom: reticulate whipray, photo by Brian Gratwicke (Flickr, CC BY 2.0).

If we modify the previous experiment by setting $s = 6$ and $L_b = 3$, initializing again from random values, it takes longer to generate equispaced dark spots, shown in the top row of Figure 10. Then we can proceed with new parameters $r = 8$ and $s = 3$ during growth with the same rate, where each spot develops into the polygonal shapes, similar to a honeycomb whipray (*Himantura undulata*).

Here we observe the hypothesis from Watanabe and Kondo [2012], where closely related species with distinct patterns might have the same biochemical mechanism but with distinct parameters. Several works already noted the ubiquity of equally spaced spots as a prepattern [Headon and Painter 2009; Murray 1988; Turk 1991]. Recently,




Fig. 10. Top row: static domain with size 100 by 100, using parameters $r = 30$ and $s = 6$. Middle row: growing domain, now with $r = 8$ and $s = 3$ (final size is 120 by 120). Bottom: honeycomb whipray (photo by the authors).

Cooper et al. [2018] described a linear RD system generating many discrete shark features, producing the same spotted prepattern.

Another experiment focuses on the formation of stripes on elongated domains, a phenomenon first noted by Murray [1981]. We first simulated reaction-diffusion in a static domain of size 50 by 10, with toroidal boundary conditions, $r = 4.5$ and $s = 4$. Starting from random concentrations, horizontal stripes were formed, as shown in the top left image of Figure 11. After the stripes stabilized, the domain was grown with a rate of 0.5 expansions per second. The dark stripes retained their size, while the light strips widened. Given enough “spare space” was available inside the light stripes, small dark spots started to emerge, which then grew into irregular dark blobs. The final pattern closely matches the pigmentation of a yellow-banded poison dart frog (*Dendrobates leucomelas*). Doan et al. [2012] showed that early stages for this particular species do present a longitudinal striped pattern with no small spots.

A current hypothesis for the origin of oriented patterns is the cell migration originated on the neural crest. In Figure 12 we show an experiment that simulates this phenomenon, starting from a domain with constant concentration 4 of A and B. We then applied uniform growth with a rate of 1.25 expansions per second. We added a small uniformly distributed random amount at each iteration to B in the top row, with amplitude 0.01, as an approximation of localized production of B along the neural crest. In the initial images from the simulation, we can see that this localized production induces stripes oriented along the neural crest. After a while, and because of the continued expansion, the darker regions are wide enough to generate smaller spots. The resulting pattern is similar to the thirteen-lined ground squirrel (*Ictidomys tridecemlineatus*).




Fig. 11. Top: growing domain with only changes in parameter $s = 2$, the initial size is 50 by 10 and the final size is 125 by 85. The A reagent is shown. Bottom: yellow-banded poison dart frog (*Dendrobates leucomelas*), photo by Adrian Pingstone (Wikimedia Commons, public domain).

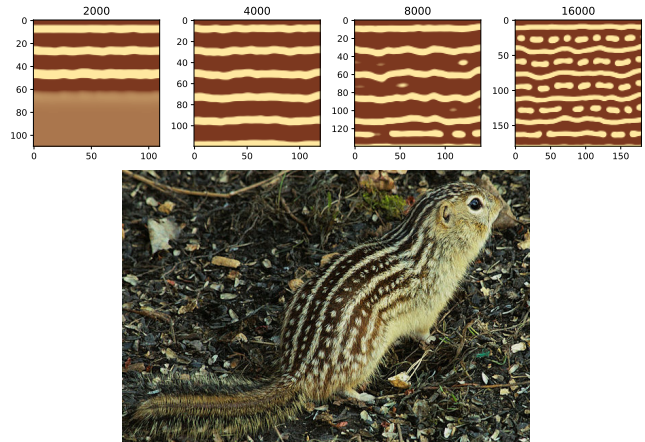


Fig. 12. Top row: simulation on a growing domain with initial constant values and localized random production. The initial size is 100 by 100, reaching a final size of 180 by 180. Bottom: thirteen-lined ground squirrel, photo by Mnmazur (Wikimedia Commons, public domain).

7.4 Combined reaction-diffusion and reinforcement

Our most significant results came from searching for an explanation for the formation of the leopard rosettes. We strived to create a plausible hypothesis of this pattern formation process, traced from the secretion of some reagent over the embryo skin, during the early days of embryogenesis, to the actual pattern of an adult leopard. Several other works also addressed the leopard patterns [Koch and Meinhardt 1994; Kondo and Shirota 2009; Liu et al. 2006; Malheiros and Walter 2017; Murray 1988; Turk 1991; Walter et al. 1998], but we believe to have developed the most complete hypothesis for its formation, thus providing the more accurate reproduction to date.

Even though only reaction-diffusion and tissue growth are not enough for us to reproduce the full rosettes, when we add the reinforcement mechanism, we can explain several phases of its formation, namely definition, development and maintenance.

The **definition phase** started with constant 4 concentrations for reagents A and B. We then added random concentrations of B with amplitude 0.01 only along the central row of the domain, shown on the left side of Figure 13. The analogy is that the domain is part of the skin of a leopard embryo, and the horizontal line in the center coincides with the neural crest, locally producing the diffusible reagent B. In this phase, we supposed a static domain, which could be equivalent to a situation where RD is much faster than actual growth. There were many possible parameter choices here, from the amount of B produced to the r ratio and the lower and upper bounds for both reagents. Albeit the diffusion is isotropic, most combinations generated parallel lines to the neural crest, producing a *generative front* that defined equispaced stripes. We have found a particular set of values that produces exactly the desired behavior. We had $r = 40$, $s = 5$ and $L_b = 2$. These values were kept unchanged through the three phases.

We have further split this simulation phase into two figures, to better describe the reaction dynamics. For the initial 5,000 iterations, we had the formation of stripes, induced by the localized production of B, as is shown in Figure 13. The B reagent is shown, with higher concentrations appearing as lighter colors.




Fig. 13. The formation of stripes in a static domain with size 130 by 130. Random production of B is done over the central row.

After stripes are initially defined, the pattern seemed to almost stabilize. However, minor local variations along the stripes, together with the continued random production at the center caused them to oscillate. This waving movement then increased and the stripes (at different moments) started to break apart and form round equispaced spots. This is shown in Figure 14. We kept the simulation running until iteration 220,000, where spots have evenly spread. Note that the particular arrangement of spots followed the direction of the early stripes. Particularly, given the continued effect of local production, the last stripe to dissolve into spots was the central one. Due to RD, the pattern was nearly stable, so continued simulation would only slightly adjust the spots, staying this way indefinitely. The definition phase has ended.

This first phase probably occurs early in the embryo life. The formed prepatter is similar to many others possible to generate from RD systems, in which the initial state is composed of many round spots. But, dissimilar to other works, we have selected this particular set of parameters first for maintaining the dorsal stripe, and second for generating spots aligned along with it, thus forming




Fig. 14. Continued simulation for a static domain. The stripes have oscillated and broken into round spots, which spread apart. The domain size is still 130 by 130. Random production of B continued over the central row.

approximate rows of round spots. Many other parameter combinations for r and L_b would generate unorganized spots, but we sought a much more structured prepatter for the leopard.

In the **development phase** we hypothesized the production of a third reagent C to match the exact concentration of B at each cell, which could take some time. Until now, we have supposed that there was a slow but steady development of the embryo, slow enough to not affect the appearance of the prepatter. Then, in this second phase, an intense growth happened, where random cells quickly subdivided. Both A and B continued to be subject to the *same parameters* as before, but with cell subdivisions happening 13.33 times per second, while reaction-diffusion is still active. This phase was very short, lasting only 12,000 iterations. In the end, the domain has grown to size 370 by 370. And, more importantly, the spots have split into the leopard rosettes, as shown in Figure 15.




Fig. 15. Spots evolve into the leopard rosettes. Concentrations for A and B are kept and growth is performed with the same parameters as before. The domain ends with size 370 by 370. There is no local production of B.

We have balanced the duration of this phase and the growth rate to match the final appearance of the rosettes to the grown adult leopard shown in Figure 19. Moreover, we have decreased the probability of division of the middle rows by half, which made nearby spots to vertically develop less. This has led to elongated rosettes that are similar to the reference leopard depicted. Therefore, a reduced and localized growth rate on the skin near the spine is a possible explanation for this local pattern variation.

At the same time, the concentration of reagent C was subject only to the reinforcement mechanism, with diffusion rate $D_c = 1$. We adjusted the specific parameters for the reinforcement to keep the overall size of spots, setting $t = 5$ and $w = 5$, although at such small a scale the former round spots were changed into irregular shapes. In general, reinforcement is more effective in maintaining larger regions. Still, the overall positions of spots and their relative size were maintained during the same 12,000 iterations, as in Figure 16. Note that the outcome was a mostly bicolor domain, where the majority of cells have either a concentration of 0 or 10 for C.




Fig. 16. The spots expanded into large irregular areas, made by reinforcement on concentrations of C, which started identical to B at iteration 220,000. Growth is the same as in Figure 15. The domain ends with size 370 by 370.

After the intense surge of cell division, we considered the development phase finished. We might suppose that each cell had acquired the concentrations of B and C that will determine the amount of melanin produced by it and their offspring cells during the animal life. This phase could probably occur when the animal was still in the womb of its mother, but having its pattern already assigned to skin cells. Such a pattern would be visible much later on the development of the fetus, when the fur grows.

In the **maintenance phase** the reinforcement process has taken care of B and C reagents (we believe A is not necessary anymore), amplifying and smoothing-out the pattern as the domain was further enlarged. We hypothesize that, as a fetus grows, a reinforcement-like mechanism could function after its birth and until the animal death. In this phase we performed a slow growth rate of 1.33 expansions per second, which run for 10,000 iterations, with the same parameters as before ($D_c = 1$, $t = 5$ and $w = 5$). Both outcomes are shown in Figure 17. After that, the maintenance phase has ended.




Fig. 17. Final states for B (on the left) and C (in the center). Each concentration was reinforced separately for 10,000 iterations. On the right, the final states are combined into a normalized melanin texture, mapped into gray values. All images have size 390 by 390.

We assume that B controls the amount of eumelanin (black pigment), whereas C accounts for the production of pheomelanin (reddish pigment). If we also suppose that eumelanin takes precedence over pheomelanin, higher concentrations of B would be visible over lower concentrations of C. We thus map the final concentrations of B and C into the normalized pigmentation parameter for the fur shader from Chiang et al. [2016]. In this range, 0 stands for white fur, passing from blonde to reddish hair, until almost black at 1. The resulting mapping is shown in Figure 17. We map the low concentrations from B and C to the value 0.3, the high concentrations from B to 1.0 and the high concentrations from C to 0.5.

We have used the grayscale image as the control texture for the pigmentation parameter of 100,000 fur strands over a simple plane, rendered in Blender 2.81 using its native implementation of the

fur shader from Chiang et al. [2016]. Each strand is assigned a single color. The skin below the fur layer has a constant pink color. Two point light sources were used and a typical setup of global illumination through path-tracing was employed (using the Cycles rendering engine). The result is shown in Figure 18, and a close-up of the same model is displayed in Figure 1. For comparison, we show the actual leopard (*Panthera pardus*) and a close-up used for reference in Figure 19. It is interesting to note that the generated pattern is simple, consisting of only two distinct “flat” areas besides a constant background. Yet, it is enough to produce a realistic result, as most of the visual complexity derives from fur orientation and self-shading, which we captured in the 3D rendering.



Fig. 18. Full 3D rendering: final melanin parameters were mapped into 100,000 fur strands placed into a plane, and rendered by a specialized fur shader. The skin below the fur layer has a constant pink color.



Fig. 19. Leopard: photo by Derek Keats (Flickr, CC BY 2.0) on the left, close-up detail on the right.

Now we briefly examine the transition from spots to rosettes. In the right column of Figure 20, we illustrate several results for the same three-phase process described. However, the length of the development phase is modified, lasting 10,000 iterations for the top image, and gradually reduced until 5,000 iterations for the bottom one. Only a part of the domain is shown for clarity. We can see that as the tissue grows less, the spots only partially develop, assuming shapes visually similar to the continuous variation already seen along the leopard body. For comparison, on the left of Figure 20, we depict the rear left leg of the same leopard individual shown earlier, taken from the same photo series.



Fig. 20. Variation on a leopard's leg. Left: rear left leg detail from the leopard, photo by Derek Keats (Flickr, CC BY 2.0). Right: resulting patterns by the development phase lasting only 10,000 (top), 9,000, 8,000, 7,000, 6,000 and 5,000 (bottom) iterations.

In Figure 21 we present a visual comparison of leopard patterns synthesized in prior works. Most use reaction-diffusion, but the models employed and the combination of techniques vary greatly.

8 CONCLUSIONS

In this paper, we have shown the role of tissue growth in the establishment of realistic skin pigmentation patterns. By providing a simple matrix expansion scheme, we could run an adapted reaction-diffusion model over a growing domain. By exploring a large portion of the parameter space, we have found interesting emerging patterns, which have been adjusted to match existing species, namely whipsnakes, poison dart frogs, squirrels and leopards.

We have emphasized that it is enough to anchor the pattern formation process in simple techniques to reproduce a few realistic patterns. Also, taking special care for the choice of prepattern is crucial to the generated outcome. We have also employed a single-reagent continuous model for reinforcement, similar to the typical reaction-diffusion equations. This model enabled the maintenance of the overall appearance of a pattern as its domain grows. We




Fig. 21. Visual comparison of generated leopard spots from the literature: top left image from Turk [1991] (detail of figure 2), top center from Koch and Meinhardt [1994] (figure 9b), top right from Walter et al. [1998] (figure 5b), bottom left from Liu et al. [2006] (figure 4b), bottom center from Kondo and Shirota [2009] (figure 4d) and bottom right from Malheiros and Walter [2017] (detail of figure 15).

thus have speculated that reinforcement could be an approximation for the still unknown regulatory mechanism that maintains tissue coherence as cells proliferate during healthy biological development.

8.1 Contributions

Here we summarize the highlights from this paper:

- We show that tissue growth is successfully approximated by a set of matrix expansions, where random cells are subdivided in parallel. This is both simpler and more efficient to implement than a point-based cell simulation, yet being sufficient for the emergence of intricate patterns.
- We propose a reaction-diffusion model derived from Turing's original non-linear model. The focus is both to provide more intuitive parameters for adjusting the resulting pattern (like the ratio and scale parameters), and to expand the model expressiveness (by setting the lower and upper bounds for concentrations).
- We present a new continuous reinforcement model, providing an autocatalytic chemical mechanism, which can be expressed as a single equation akin to traditional RD systems. This enables the overall pattern maintenance during growth.
- We discuss the importance of the careful definition of the initial state and show the effect of providing biologically plausible prepatterns. In particular, we draw attention for both sparse spots and localized production, as typical states that induce the development of more complex skin patterns.
- Finally, we have generated a few unprecedented 2D patterns matching real species. For those, we were not just concerned about the visual result, but also interested in establishing a hypothetical set of phases that could model the actual development of the pigmentation patterns. For that, we tried to minimize parameter changes between stages and avoided using any artificial control of the simulation outcome.

8.2 Limitations and further research

As our major research goal, we strived to use only biologically plausible pattern formation mechanisms. Therefore, we did not focus on actual explanations for biochemical mechanisms or cell-to-cell communication. We also did not address the stability analysis of the mathematical models used. We believe there is much room to be explored in those directions.

It is still an open research point whether a PDE system that also takes into account continuous domain growth would achieve the same results. We have made a brief comparison between our matrix expansion scheme and the equivalent image interpolation into a slightly higher resolution. That is, instead of expanding an $m \times n$ matrix once, we would interpolate its values into a new $(m+1) \times (n+1)$ matrix. For example, using the same parameters from the frog experiment (Figure 11), we have tested nearest, bilinear and bicubic interpolation. In those three cases, the distinct frog spots did not appear, only forming straight horizontal stripes, as in Figure 22. So we believe that the noise introduced by our matrix expansion scheme is a crucial part of the process, approximating the actual cell division and introducing small local irregularities to the pattern formation. Perhaps a continuous growth term would be too smooth to produce some patterns; this is an important research question.




Fig. 22. Growing domain by using bicubic image interpolation, following the same prepattern and parameters as in Figure 11. Stripes split but there are no spots formed due to lack of “division noise”.

Another perceived limitation is the length of the exploratory work to discover new patterns. However, we believe our approach could be extended to match other biologic pigmentation patterns, needing a technique for pattern similarity comparison and visual characterization, providing a utility function to automate pattern matching. Then, it would theoretically be possible to infer the correct parameters and prepattern to reproduce a given input. But our technique certainly is not able to match arbitrary textures, by not being image nor exemplar-based.

Our process can provide pattern diversity by changing both the random seed and slightly altering the parameters from an already defined pattern, albeit this is manually done by now. In the long run, our approach could be streamlined to provide finer detail and variation to artistic production lines.

In Appendix A, we make a brief comparison of the computational effort against a point cloud arrangement. We opted to use matrices based on their performance, as our setup aimed interactive exploration. But our approach could be directly applied to a triangulated 3D mesh, provided the adapted Turing model runs over its surface and the idea of local growth is executed by adding random new vertices, approximating cell division.

We believe that although focused on coat patterns, our results are essential for Computer Graphics because they restate that complex pigmentation patterns seem mostly induced by growth and randomness. Besides, 3D form development induces localized surface growth. Therefore, a proper animation of an embryo with our technique applied onto its surface would yield local patterns resulting from the composed effects of tissue development.

As a first step towards this goal, we succeed in simulating in 2D the overall arrangement of spots of a leopard along its spinal column. We also matched the variation along its leg, explaining that the differences of spot sizes and shapes might be simply due to distinct growth during animal development.

Naturally, exploring this process over a 3D model is a promising next step, where many scientific works have already laid the basis for further research in that front.

Finally, another large venue for exploration is to try to reproduce many other biological patterns seen in nature, including those from the plant kingdom.

REFERENCES

- Jonathan Bard and Ian Lauder. 1974. How well does Turing’s theory of morphogenesis work? *Journal of Theoretical Biology* 45, 2 (1974), 501–531.
- Jonathan B.L. Bard. 1981. A model for generating aspects of zebra and other mammalian coat patterns. *Journal of Theoretical Biology* 93, 2 (1981), 363 – 385.
- Benjamin J Binder, Kerry A Landman, Matthew J Simpson, Michael Mariani, and Donald F Newgreen. 2008. Modeling proliferative tissue growth: a general approach and an avian case study. *Physical Review E* 78, 3 (2008), 031912.
- Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. 2016. A practical and controllable hair and fur model for production path tracing. *Computer Graphics Forum* 35, 2 (2016), 275–283.
- Rory L Cooper, Alexandre P Thiery, Alexander G Fletcher, Daniel J Delbarre, Liam J Rasch, and Gareth J Fraser. 2018. An ancient Turing-like patterning mechanism regulates skin denticle development in sharks. *Science advances* 4, 11 (2018), eaau5484.
- Tiffany M Doan, Anthony M Nowacki, and Perry A Roberts. 2012. *Dendrobates leucomelas*. In *Catalogue of American Amphibians and Reptiles (CAAR)*. Society for the Study of Amphibians and Reptiles, Austin, TX, USA.
- Kurt W Fleischer, David H Laidlaw, Bena L Currin, and Alan H Barr. 1995. Cellular texture generation. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. Association for Computing Machinery, New York, NY, USA, 239–248.
- Deborah R. Fowler, Hans Meinhardt, and Przemyslaw Prusinkiewicz. 1992. Modeling seashells. *Comp. Graphics* 26, 2 (1992), 379–387.
- Alfred Gierer and Hans Meinhardt. 1972. A theory of biological pattern formation. *Kybernetik* 12, 1 (1972), 30–39.
- Charles Gingras and Paul G Kry. 2019. Procedural Modelling with Reaction Diffusion and Growth of Thin Shells. In *Proceedings of Graphics Interface 2019*. Canadian Information Processing Society, Mississauga, Canada, 1–7.
- Denis J Headon and Kevin J Painter. 2009. Stippling the skin: Generation of anatomical periodicity by reaction-diffusion mechanisms. *Mathematical Modelling of Natural Phenomena* 4, 4 (2009), 83–102.
- Yiwei Hu, Julie Dorsey, and Holly Rushmeier. 2019. A Novel Framework for Inverse Procedural Texture Modeling. *ACM Trans. Graph.* 38, 6, Article Article 186 (Nov. 2019), 14 pages.
- Joseph T Kider Jr, Samantha Raja, and Norman I Badler. 2011. Fruit senescence and decay simulation. *Computer Graphics Forum* 30, 2 (2011), 257–266.
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe patterns on surfaces. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–11.
- AJ Koch and Hans Meinhardt. 1994. Biological pattern formation: from basic mechanisms to complex structures. *Reviews of modern physics* 66, 4 (1994), 1481–1507.
- Shigeru Kondo and Hideaki Shirota. 2009. Theoretical analysis of mechanisms that generate the pigmentation pattern of animals. *Seminars in cell & developmental biology* 20 (2009), 82–89.
- RT Liu, SS Liaw, and PK Maini. 2006. Two-stage Turing model for generating pigment patterns on the leopard and the jaguar. *Physical review E* 74, 1 (2006), 011914.
- Philip K Maini, Thomas E Woolley, Ruth E Baker, Eamonn A Gaffney, and S Seirin Lee. 2012. Turing’s model for biological pattern formation and the robustness problem. *Interface focus* 2, 4 (2012), 487–496.
- Marcelo Malheiros and Marcelo Walter. 2017. Pattern formation through minimalist biologically inspired cellular simulation. In *Proceedings of Graphics Interface 2017 (GI*

- 2017). Canadian Human-Computer Communications Society, Mississauga, Canada, 148 – 155.
- Hans Meinhardt. 2009. *The algorithmic beauty of sea shells*. Springer Science & Business Media, Berlin, Germany.
- Seita Miyazawa, Michitoshi Okamoto, and Shigeru Kondo. 2010. Blending of animal colour patterns by hybridization. *Nature communications* 1 (2010), 66.
- J. D Murray. 1981. A pre-pattern formation mechanism for animal coat markings. *Journal of Theoretical Biology* 88, 1 (1981), 161–199.
- James D Murray. 1988. How the leopard gets its spots. *Scientific American* 258, 3 (1988), 80–87.
- James D Murray. 2003. *Mathematical Biology. II Spatial Models and Biomedical Applications*. Springer-Verlag, Berlin, Germany.
- John E Pearson. 1993. Complex patterns in a simple system. *Science* 261, 5118 (1993), 189–192.
- Lara Raad, Axel Davy, Agnès Desolneux, and Jean-Michel Morel. 2018. A survey of exemplar-based texture synthesis. *Annals of Mathematical Sciences and Applications* 3, 1 (2018), 89–148.
- Adam Runions, Martin Fuhrer, Brendan Lane, Pavol Federl, Anne-Gaëlle Rolland-Lagan, and Przemyslaw Prusinkiewicz. 2005. Modeling and visualization of leaf venation patterns. In *ACM SIGGRAPH 2005 Papers*. Association for Computing Machinery, New York, NY, USA, 702–711.
- Allen R Sanderson, Robert M Kirby, Chris R Johnson, and Lingfa Yang. 2006. Advanced reaction-diffusion models for texture synthesis. *Journal of Graphics, GPU, and Game Tools* 11, 3 (2006), 47–71.
- A. M. Turing. 1952. The Chemical Basis of Morphogenesis. *Phil. Trans. Roy. Soc. London B*, 237 (1952), 37–72.
- Greg Turk. 1991. Generating textures on arbitrary surfaces using reaction-diffusion. *ACM SIGGRAPH Computer Graphics* 25, 4 (1991), 289–298.
- Alexandria Volkening and Björn Sandstede. 2015. Modelling stripe formation in zebrafish: an agent-based approach. *Journal of the Royal Society Interface* 12, 112 (2015), 20150812.
- Marcelo Walter, Alain Fournier, and Daniel Menevaux. 2001. Integrating Shape and Pattern in Mammalian Models. In *SIGGRAPH 2001, Computer Graphics Proceedings (Annual Conference Series)*. Association for Computing Machinery, New York, NY, USA, 317–326.
- Marcelo Walter, Alain Fournier, and Mark Reimers. 1998. Clonal mosaic model for the synthesis of mammalian coat patterns. In *Graphics Interface*, Vol. 98. Canadian Human-Computer Communications Society, Mississauga, Canada, 82–91.
- Masakatsu Watanabe and Shigeru Kondo. 2012. Changing clothes easily: connexin41.8 regulates skin pattern variation. *Pigment cell & melanoma research* 25, 3 (2012), 326–330.
- T Weaver and Zhidong Xiao. 2016. Fluid Simulation by the Smoothed Particle Hydrodynamics Method: A Survey. In *Proceedings of the 11th Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*. Science and Technology Publications, Setúbal, Portugal, 215–225.
- Andrew Witkin and Michael Kass. 1991. Reaction-diffusion textures. *ACM Siggraph Computer Graphics* 25, 4 (1991), 299–308.
- Lewis Wolpert, Cheryll Tickle, and Alfonso Martinez Arias. 2015. *Principles of development*. Oxford University Press, New York, USA.
- Yang Zhou, Zhen Zhu, Xiang Bai, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2018. Non-stationary Texture Synthesis by Adversarial Expansion. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 37, 4 (2018), 49:1–49:13.

A TIMINGS AND PARAMETERS

Our implementation is written in Python with the NumPy and SciPy libraries, being run in a Linux desktop powered by an Intel Core i3-9100F CPU running at 3.6 GHz.

We have compared the performance between our matrix expansion approach and the two point-based NNS techniques available in the publicly available source code from Malheiros and Walter [2017], namely spatial sorting and k -d tree. We have set up a 50 by 50 initial domain with reaction-diffusion that grows during 12,000 iterations until reaching about 18,000 cells in both our matrix expansion (taking on average 6.06 s) and spatial sorting (taking 44.02 s). Our technique is about seven times faster on average. In a second test, we set up a 100 by 100 initial matrix, again growing during 12,000 iterations until reaching about 19,000 cells. The comparison was made between our expansion scheme (on average 8.57 s) and

the k -d tree (taking 153.65 s), which gives a result nearly eighteen times faster.

In Tables 1 and 2 we present the timings and parameters for most figures in our work.

Table 1. Timing and information for the figures.

Figure	Time (s)	Final size	Growth rate (/s)	Iterations	Boundary condition
4	1.72	50 × 50	-	≤ 33,333	toroidal
5	5.46	100 × 100	0.5	≤ 33,333	toroidal
6 top	0.94	200 × 200	5	5,000	toroidal
6 mid.	6.14	200 × 200	5	5,000	toroidal
6 bot.	6.14	200 × 200	5	5,000	toroidal
8	1.41	200 × 200	5	5,000	toroidal
9	14.45	200 × 200	2.5	14,000	toroidal
10	13.68	120 × 120	2.5	24,000	toroidal
11,22	5.87	125 × 85	0.5	19,000	no flux
12	15.60	180 × 180	1.25	16,000	no flux
13,14	208.55	130 × 130	-	220,000	no flux
15	45.52	370 × 370	13.33	12,000	no flux
16	15.17	370 × 370	13.33	12,000	no flux
17	54.23	390 × 390	1.33	10,000	no flux

Table 2. Parameters for the figures.

Figure	Parameters
4,5	$r = [4.2, 6, 12, 20, 40]$, $s = 3$, $U_a = [6.2, 6.0, 5.2, 4.4, 4.0]$ $L_b = [3, 3, 3, 3.5, 4]$, $\Delta t = [0.01, 0.01, 0.008, 0.006, 0.003]$
6 mid.	$r = 8$, $s = 10$, $\Delta t = 0.004$
6 bot.	$r = 8$, $s = 10$, $U_a = 5.7$, $\Delta t = 0.004$
8	$D_c = 1.5$, $t = 5.8$, $w = 4$, $\Delta t = 0.004$
9	$r = 30$, $s = 3$, $L_b = 2$, $\Delta t = 0.004$ $r = 8$, $s = 3$, $L_b = 2$, $\Delta t = 0.004$
10	$r = 30$, $s = 6$, $L_b = 3$, $\Delta t = 0.002$ $r = 8$, $s = 3$, $L_b = 3$, $\Delta t = 0.002$
11,22	$r = 4.5$, $s = 4$, $U_a = 6.8$, $U_b = 6.2$, $\Delta t = 0.01$ $r = 4.5$, $s = 2$, $U_a = 6.8$, $U_b = 6.2$, $\Delta t = 0.01$
12	$r = 20$, $s = 4$, $U_a = 7.2$, $U_b = 7.2$, $\Delta t = 0.004$
13,14 15,17	$r = 40$, $s = 5$, $L_b = 2$, $\Delta t = 0.0015$
16	$t = 5$, $w = 5$, $D_c = 1$, $\Delta t = 0.0015$