Partial Redundancy and Morphological Homeostasis: Reliable Development through Overlapping Mechanisms

Micah Brodsky[1]

[1]Massachusetts Institute of Technology, Cambridge, MA 02139
micahbro@csail.mit.edu

August 26, 2016

Abstract

How might organisms grow into their desired physical forms in spite of environmental and genetic variation? How do they maintain this form in spite of physical insults?

This paper presents a case study in simulated morphogenesis, using a physics-based model for embryonic epithelial tissue. The challenges of the underlying physics force the introduction of closed-loop controllers for both spatial patterning and geometric structure. Reliable development is achieved not through elaborate control procedures or exact solutions but through crude layering of independent, overlapping mechanisms. As a consequence, development and regeneration together become one process, "morphological homeostasis", which, owing to its internal feedbacks and partially redundant architecture, is remarkably robust to both knock-out damage and environmental variation. The incomplete nature of such redundancy furnishes an evolutionary rationale for its preservation, in spite of individual knock-out experiments that may suggest it has little purpose.

1   Introduction

The physical forms of multicellular organisms are amazingly robust, developing correctly in spite of substantial environmental and genetic variation. This phenomenon was dubbed the "canalization" of development by Waddington [28], reflecting the notion that there seems to exist some sort of restoring force pulling developing organisms back to their expected phenotype whenever perturbed. The most dramatic example may span entire phyla, as animals within a single phylum start from dramatically different initial conditions yet converge to a common "phylotypic" stage of development, before differentiating into their characteristic larval forms [21]. Similar convergence effects in spite of environmental perturbations can also be seen to varying degrees in the adult forms of animals, ranging from wound healing, to limb regeneration, to complete body reassembly after disaggregation, as in the hydra [18].

What sorts of principles and tools does nature employ to produce such astonishing robustness? Can we master them ourselves, whether for engineering robust systems or for a deeper understanding of natural phenomena?

1.1   Morphological Homeostasis

Waddington's hypothetical "restoring force" of development cannot be completely hypothetical. For the dynamics of a physical system, such as an organism, to converge to a common attractor, the dynamics must be sensitive to the present state of the system – there must be feedback. Though such sensitivity can be a natural consequence of inanimate dynamics, for example, the surface tension that draws a droplet into a sphere, the complexity of biological forms strongly suggests explicit feedback control – an idea explored in this paper. We might dub Waddington's phenomenon, as extended to the adult, "morphological homeostasis".

1.2   Redundancy and Partial Redundancy

Is feedback control as an organizing principle enough to explain the reliability of development? In engineered systems, high reliability is typically achieved through redundancy, not merely feedback. For maintaining homeostasis, however, redundancy brings hazards of its own.

Consider an example from engineering, the RAID-5 disk array. Such a system uses $n + 1$ hard drives to provide n drives' worth of space. Through clever use of parity bits, it can survive a single drive failure with no loss of data or availability and with negligible performance degradation. Unfortunately, a common consequence is that the first drive failure goes completely unnoticed, until the second drive fails some time later and the entire data set is lost.

Full redundancy has a fundamental flaw: because it is so successful in hiding early failures, necessary steps to restore the system to its original, healthy state are neglected [26]. Such systems are vulnerable to invisible deterioration: "rot".

Redundancy becomes an expendable resource, a finite buffer against damage that, once depleted, is gone for good.

Nature has long since explored the design trade-offs here. The excessive kidney capacity we are born with is a good example, likely an acceptable compromise given our limited lifespans [2]. On the other hand, for the integrity of our genes, such expendable redundancy is completely inadequate.

Animal genomes incorporate an entire hierarchy of redundancy measures. At the lowest level is a form of full redundancy reminiscent of RAID-1: the two-way mirroring within the double-stranded DNA polymer. There is, however, a crucial difference: in a cell, regular maintenance is tightly coupled into the system, not an afterthought to be handled by some outside process (i.e., a harried system administrator). Like in aviation, a cell that fails inspection is "grounded" – it enters senescence, ceasing to divide, or undergoes apoptosis, removing itself from the system [27]. Moreover, a cell doesn't need outside inspectors carefully following a maintenance protocol; it inspects and repairs itself. Such intrinsic self-maintenance is a significant improvement over the blind redundancy furnished by a hard drive array.

Atop this 2-way mirror set lie multiple layers of further redundancy, but none so rigid and symmetrical; not replication, but imperfect redundancy, where the replicas are only similar at best and sometimes very different. A striking example is the diploid structure of animal cells: two nearly-complete but non-identical copies of the program code are included and executed simultaneously. Identicality (homozygosity), indeed, is often downright harmful.

Genes are also duplicated throughout the genome, and rarely are the copies identical. Unless there is a selective advantage to increased RNA throughput through simultaneous transcription (as in the unusual case of ribosomal RNAs, for example), identical duplicates constitute expendable redundancy: so long as accidental damage to a gene is more likely than successful re-duplication, spare copies are likely to be lost through neutral drift [26, 17]. Instead, over evolutionary time, any accidental copies that remain diverge and acquire new functions [7]. Much of the original functionality remains, to the extent that a knock-out of either copy is often survivable, but not without some cost.

Most extreme is the case when redundancy is provided by completely unrelated genetic components through differing physical processes. Animal physiology is rife with highly divergent mechanisms converging on a common purpose. For example, blood loss at a wound is held in check simultaneously by the platelet clotting system, the thrombin/fibrinogen clotting system, and vasoconstriction. Why so many complex mechanisms? Why are these not pared down through neutral drift? In spite of their overlap, each independent mechanism seems to confer its own, unique selective advantage. That is, the mechanisms are not fully redundant, they are *partially* redundant. Damage to one is disadvantageous (e.g. as a hemophilia), but survivable.

Why should nature prefer partial redundancy? Why not do things one way and do it well? As a form of redundancy, of course, partial redundancy offers a buffer against damage and stress, bodily, genetic, or environmental. Unlike full redundancy, however, a component lost or weakened will cause detectable

degradation. The gaps in redundancy are visible, and precisely *because* they are visible, partial redundancy provides feedback – feedback that favors regeneration (either somatic or selective), or even learned avoidance of danger. Partial redundancy, much more than full redundancy, facilitates homeostasis.

This paper explores a detailed case study in partial redundancy, arising in the problem of morphological homeostasis: how an organism attains and maintains its physical form, in spite of external insults, environmental variation, and internal evolutionary changes. The physical substrate used is the deformable surface model developed in prior work [5, Ch. 2], a rich, "2.5"-dimensional physics that caricaturizes the mechanics of embryonic epithelial tissue. Taming this physics requires a fair amount of new mechanism for sensing and for actuation. In the course of developing this mechanism, the need for partial redundancy arises naturally. Three cases are explored: patterning, structural sensing, and mechanical actuation. In each case, several simple methods are available, none entirely satisfactory. However, in each case, a partially redundant combination of two such methods is trivial to construct and performs superbly. The final results are remarkably robust, showing how effective homeostasis through partial redundancy can be.

2   Model Background

Redundancy is often invoked as a buffer against mutational damage and as a means to smooth out rugged fitness landscapes [26, 17, 13]. However, in the absence of mutation and damage, the value of feedback and redundancy is not particularly apparent, especially under stylized, deterministic physics like cellular automata. On such predictable, local substrates, dead reckoning is highly effective, and spatially separated phenomena may be regarded as independent. However, these assumptions fall apart when a model's physics become sufficiently rich. With sophisticated mechanics, available strategies for development become more varied, but also, their effects less predictable, less modular. Feedback and redundancy prove valuable even with deliberately engineered developmental programs.

Recent years have seen the use of increasingly sophisticated physical models (e.g. [6, 9, 11, 20, 16]), exposing richer dynamics than earlier, idealized models of development (e.g. [25, 10, 8]). The model employed here [5, Ch. 2], unlike the more common mass-spring models, trades off generality for improved mechanical richness by specializing to epithelial (sheet-like) tissues. This work also focuses on development by embryomorphic mechanical transformations, not by cell proliferation. However, the concepts should be applicable to any rich, 3d physics where cells can sense and manipulate their mechanical environment.

The model used here is a "vertex model", a representation of a foam-like sheet of polygonal, tightly adhering cells in terms of the positions of their vertices [14]. Cell shapes, and hence vertex positions, are governed by surface tension and internal elasticity (Figure 1). A distinctive feature of vertex models is how they naturally accommodate cell rearrangement and plastic flow, through

$k_A(A_0 - A)$ σ_{ij} $k_B(\theta_0 - \theta)$

Figure 1: Principal mechanical forces governing the epithelial cell model. Left to right: bulk elasticity, surface tension, flexural elasticity.

the "T1" interchange process (Figure 2). The particular model used here is extended into 3d by the addition of flexural springs at each cell-cell junction, the spring constant determining stiffness. Dynamics are evolved via quasi-static energy minimization of a global mechanical energy function, consisting principally of the sum over all cells of the following terms:

$$E_{2d} = k_P \mathsf{P} + k_A (\mathsf{A} - A_0)^2 \tag{1}$$

$$E_B = \sum_{e \in edges} k_B \mathsf{L_e} (\Theta_\mathsf{e} - \theta_0)^2 \tag{2}$$

where E_{2d} is the energy of two-dimensional, in-plane distortion and E_B is the energy of bending in the third dimension. P is the cell perimeter, A is its area, and A_0 is a parameter specifying the equilibrium area. Θ_e is the bending angle at the junction, $\mathsf{L_e}$ is the length of the associated edge, and θ_0 is a parameter specifying the setpoint angle for the junction (by default shared for all a cell's edges). Differential adhesion and traction are implemented by modulating k_P on a per-edge basis depending on the cell types involved (and, in the case of traction, on which vertices the traction is applied to). Additional minor terms (contact forces and pathology-avoidance terms) and the values of all parameters can be found in Table 2.1 of the thesis [5, Ch. 2].

Cells are regulated by simple software agents, the realization in terms of genes not a focus of this work. Cell agents randomly receive the opportunity to execute whenever the mechanical energy minimization process converges sufficiently. Cells are allowed to read out properties of their mechanical conformation such as elongation or curvature and can influence it through neighbor-neighbor tractions or by modifying set-points such as flexural angle. Sufficient traction or external force will cause cells to intercalate, rearrange, and flow. How these effects can be profitably applied is a key focus of the paper.
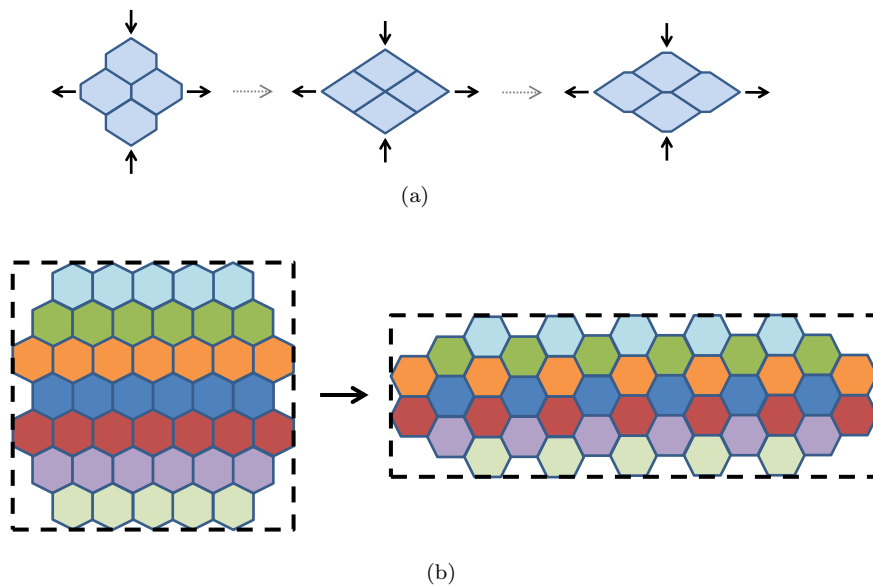
(a)

(b)

Figure 2: Cell rearrangement in vertex models. (a) Elementary T1 interchange, which occurs whenever forces reduce an edge to zero length, whether driven internally by traction or externally by applied stresses. (b) Plastic flow ("convergent extension") collectively resulting from many interchanges.

3   Decomposing the Problem

Natural biological structures are complicated, combining multiple subparts with differing characteristics. Borrowing inspiration from nature [21, 7], we can simplify the problem of engineering morphological homeostasis by breaking it into a cascade of two subproblems: patterning and actuation. Patterning – "what goes where" – consists of laying out a body plan for the structure. Actuation – "what happens here" – represents the processes of local mechanical transformation necessary to create the desired features, given a pre-existing global body plan. Of course, these problems are not independent – the global body pattern affects how actuation efforts cross-interact, and updates to the global pattern require corresponding updates to the local features. Similarly, local actuation alters the geometric properties of the substrate, modifying the patterning process, as well as rearranging already patterned cells. However, so long as the goals of patterning and actuation are compatible, I show that the combination of appropriately robust patterning and actuation algorithms can yield a robust and stable complete solution.

The presence of conserved compartment maps in animals, an invisible and highly conserved pattern of gene expression prior to detailed morphogenesis [21, 7], suggests that nature may use a similar decomposition strategy. Since perturbations in early, pre-morphogenesis development as well as local injuries to

the final form can heal, global patterning (e.g. as in [12]) and local actuation are both independently likely to involve feedback mechanisms.

The first problem, body plan patterning, can be solved by a patterning mechanism that is robust to widely variable substrate geometries and produces meaningfully consistent patterns before and after deformation. The patterning mechanism must also self-correct in the face of perturbations, without requiring a clean slate restart; incremental corrections to pattern and geometry must eventually converge, after all. These requirements all but eliminate self-timed pre-patterning [3], which cannot respond to unexpected deviations, and likely disfavor fixed-wavelength Turing-type mechanisms, which have a preferred body size and may reconfigure under deformation (although note Meinhardt's success in a restricted case [23]). Morphogenetic fields with self-sustaining sources (e.g. as in [11]) might be usable, with some caveats due to geometry [5, Ch. 4]. However, the normal neighbors patterning mechanism developed in earlier work [4, 5, Ch. 4], where patterns are specified through a purely topological description (an adjacency graph) and maintained continuously – hence tolerating substantial distortion – fits almost perfectly.

3.1 Background: Normal Neighbors Patterning

How might a body plan pattern establish and maintain itself? A common theme seen in developmental biology is the Rule of Normal Neighbors [24]: a point in a patterned tissue knows what elements of the pattern belong adjacent to it, its "normal neighbors". If it finds its neighbors are wrong, it takes corrective actions, such as re-growing a more appropriate neighbor or changing its own fate to better fit its environment. This general rule captures many striking experimental results, such as the growth of additional, *inverted* segments in cockroach limbs when the distal portions of the limbs are excised and replaced with excessively long explants [15].

This idea, that the constraint of pattern continuity ultimately governs the dynamics of patterning, can be rendered into an algorithm through constraint propagation techniques [5, Ch. 4]. Soft constraint propagation via free energy minimization is an attractive strategy, requiring only local computation and minimal resources. We represent the topology of the desired pattern as an adjacency graph over discrete pattern states (e.g. Figure 3). Based on this graph, we construct an energy function using local interactions, for which the desired pattern is (usually) a minimum. Individual cells can then explore this potential by a process mathematically analogous to thermal (and simulated) annealing, seeking a minimum.

To augment the purely topological information conveyed by the adjacency graph, several more elements are necessary. An implicit self-edge is added to each pattern state, so that the pattern may be scaled up freely, independent of any lattice spacing (contrast with [29]). The self-edge is configured to be stronger than ordinary neighbor edges, giving rise to a virtual "surface tension" effect, favoring compact, geometrically well-behaved pattern regions. We also generally add quorum sensing morphogens, providing long range negative

7

Figure 3: Schematic illustration of normal neighbors patterning. Left: Pattern state adjacency graph. Right: Associated desired pattern.

feedback that stabilizes the relative sizes of regions.

Energy is minimized through Boltzmann-weighted mean field relaxation, yielding a purely analog algorithm. The result is reminiscent of a mean field version of the classic cellular Potts model [19], modified for the purpose of pattern constraint propagation rather than cell sorting mechanics. Formally, the steady state of the system is given by Equations 3 and 4 below, and the pattern is developed by relaxing the $\mathbf{p}(i)$ toward steady state. The net result is a well-behaved spatial patterning mechanism that demonstrates spontaneous symmetry breaking, approximate scale invariance, and self-repair [5, Ch. 4].

$$p_s(i) = e^{-E_s(i)/T} / \sum_{t \in states} (e^{-E_t(i)/T}) \tag{3}$$

$$E_s(i) = -k_h p_s(i) - \frac{k_q}{q_s + k_{soft}} - \frac{1}{|n_i|} \sum_{j \in n_i} (\hat{\mathbf{e}}_s^{\mathsf{T}} \mathbf{U} \mathbf{p}(j))) \tag{4}$$

For each cell i, $\mathbf{p}(i)$ represent levels of commitment to each pattern component, $\mathbf{E}(i)$ represent the associated energies, and $T = 0.2$ is the virtual temperature. A pattern fate s is considered selected when its $p_s(i)$ exceeds 0.5. n_i is the set of spatial neighbors of cell i, $\mathbf{U}$ is the adhesion energy matrix, equal to the abstract adjacency graph's adjacency matrix plus the surface tension constant $k_s = 0.5$ times the identity matrix, and $\hat{\mathbf{e}}_s^{\mathsf{T}} \mathbf{U}$ denotes the s-th row of $\mathbf{U}$. $k_q = 0.0125$ is the quorum sense gain constant, q_s is the quorum sense level for state s (scaled as a fraction of total population of cells), and $k_{soft} = 0.005$ is the softening constant. $k_h = 0.125$ is the hysteresis constant. Quorum sensing is computed through long-range diffusion, by relaxing

$$\nabla^2 q_s = -p_s^2 + \gamma_q q_s \tag{5}$$

where γ_q is a constant governing the decay range of the associated morphogens (which must be significantly larger than the substrate size to be effective).

8

Figure 4: Twist and twin defects in normal neighbors patterns (computed on a square lattice). Top left: Correct pattern for Figure 3. Bottom left: Stable local minimum with inconsistent handedness. Top right: Correct pattern for a related adjacency graph particularly prone to stable twinning; boundary conditions fixed at 5. Bottom right: Twinned configuration.

3.2   Partially Redundant Patterning

Is the body plan patterning problem solved? The prior work summarized above offers a method for generating and maintaining a wide variety of patterns on deforming substrates. However, there still remain some notable limitations. As a metastable symmetry breaking process, normal neighbors pattern formation is not particularly fast. It also may not break symmetry directly to the desired pattern but instead pass through several undesirable intermediates first. For certain classes of patterns, it can even remain stuck in local minima, such as twisted and twinned states (e.g. Figure 4).

What might be done about these drawbacks? Interestingly, they correspond quite closely to the key strengths of some of patterning mechanisms we originally dismissed. Gradient-based patterning [25, 10], for example, is fast and free of local minima but encounters difficulties under large deformations [5, Ch. 4]. Self-timed pre-patterning [3] has the additional advantage of converging directly to the correct solution with no intermediates, but it behaves erratically if the substrate's cellular topology rearranges beneath it.

What if we combined mechanisms? Pre-patterning could provide a rapid,

reliable initial layout, which would then be maintained in homeostasis by the normal neighbors algorithm. This would require specifying the pattern redundantly, in two different encodings, but this turns out not to be so onerous. The two specifications do not need to match up exactly, as the pre-pattern determines only the initial transients; the long-term behavior is still determined by normal neighbors. The pre-pattern need not even be particularly complete; a mere sketch will yield most of the benefits. The most important requirement is that it break all the relevant symmetries.

The combination, then, amounts to a partially redundant approach to patterning. If the pre-pattern and the normal neighbors specifications drift apart, the worst that happens is that the patterning process becomes slower and less reliable; it is still largely functional. We will not pursue this direction farther in this paper, as all the structures demonstrated here develop reliably under normal neighbors alone (with one exception, noted in Section 5). However, adding partially redundant pre-patterning can significantly improve development speed while eliminating transient developmental excursions.

4 Controlling Shape

The core of this case study, then, is devoted to the problem of "what happens here": how to produce and maintain simple geometric features in spite of perturbations. We have at our disposal several mechanical actuation mechanisms, including cell shape change, apico-basal constriction, and neighbor traction forces (for simplicity, changes in cell number are not considered). Producing geometric features using these mechanisms is not too hard, given a known initial state. However, given perturbations, the initial state is *not* known. Instead, we must find techniques that respond appropriately to the system's pre-existing state.

Sensitivity to the state of the system – feedback – requires either that the intrinsic physics of the system be sensitive to system state (e.g. mechanical restoring forces) or that explicit feedback sensors be deployed by the control algorithm. Geometric structure involves numerous degrees of freedom, many of which are uninteresting (e.g. the relative arrangement of equivalent cells) or undesirable (e.g. high-frequency Fourier components). It can be valuable to leave such degrees of freedom to autonomous energy-minimization dynamics, for example, viscous relaxation, avoiding the control algorithm having to treat them explicitly. On the other hand, certain degrees of freedom represent key control targets. For these, we require sensors.

4.1 Sensing Curvature

For our first attempt at controlling geometry, consider spherical curvature – to produce spherical caps of varying radii, and hence varying subtended angle (e.g. Figure 8). First, we need a distributed, scale-invariant measure of curvature, built from local sensors. We assume cells can sense local properties of their

shape, such as bend angles, and that they can probe collective properties by, for example, the diffusion of morphogens.

Classical local measures of spherical curvature, such as Gaussian curvature and mean curvature, are not scale-invariant but instead provide curvature radii; they indicate how tightly curved the surface is locally but not how much curvature the surface encompasses as a whole. Gaussian curvature can be integrated over area to produce a dimensionless invariant related to the subtended angle (by the Gauss-Bonnet theorem), but this is an extensive quantity. In general, determining extensive quantities through local measurements seems to require leader election or an equivalent broken symmetry [5, Ch. 5]. It would be preferable to avoid this complication.

Another approach is to consider global properties based on length and area. For example, on a spherical cap, the ratio of area to the square of some linear dimension (e.g. perimeter) uniquely identifies the angle subtended. Without a leader, area and perimeter may not be directly measurable. However, the ratio of area to perimeter is easily measured (the 2D analogue of surface area to volume ratio, inverted), providing a second non-scale-invariant measure of curvature. This can be combined with the average of some local measure of curvature – for example, multiplying by average mean curvature – to produce a scale-invariant measure of global curvature that can be computed collectively by the cells.

A variation that seems to work particularly well is to combine the area-to-perimeter ratio with a purely local measure of curvature (not an average), producing a hybrid measure that is partly local, partly global. This seems to suit the mechanical effects of curvature actuation, also partly local, partly global. The measure chosen here is the product of the area-perimeter ratio and the extrinsic radius of curvature along the axis parallel to the region boundary – that is, the local circumferential curvature. Such a combination is an example of sensor-level partial redundancy, combining multiple independent mechanisms to greatly broaden the range of applicability.

4.2   Actuating Curvature

The previous section showed how to build a sensor for spherical curvature. This section explores how to build an actuator. As noted before, surfaces have numerous degrees of freedom; all of them need to be stable, and some of them need to reach particular control targets. In almost any representation, they are cross-coupled, due to the constraints of surface geometry and the complicated dynamics of deformation and flow.

For example, one might instruct each cell to bend itself in accordance with the sign of the error reported by the curvature sensor. Such "extrinsic" curvatures can be driven by, e.g., apical/basal constriction (see Figure 5). This approach, however, suffers from two serious flaws: it is geometrically inconsistent, and it does nothing to keep undesirable degrees of freedom under control. It is inconsistent for the same reason one cannot flatten an orange peel without tearing it: extrinsic curvatures require, in general, non-Euclidean geometries

Figure 5: "Extrinsic curvature" vs. "intrinsic curvature". Left: Extrinsic curvature reflects the tendency for neighboring patches of surface to be rotated in space, i.e. due to a wedge-shaped cross section – reflected in cell-to-cell bend angles. Middle: Intrinsic curvature reflects non-Euclidian distance relationships within the surface, i.e. a circular patch's circumference less than or greater than $2\pi r$. Extrinsic curvature is a property of the shape of cells, but intrinsic curvature is principally a property of how they are arranged. A curved surface ordinarily involves both intrinsic and extrinsic curvature. Right: Cells may intercalate among one another in a coordinated fashion as per Figure 2b to adjust intrinsic curvature. Cells intercalating so as to decrease a region's circumference is known as "purse string contraction".

within the surface. Distances between points within the surface must change in order to accommodate the extrinsic curvature. If a surface is deformed extrinsically, non-Euclidean "intrinsic curvature" will necessarily be generated by elastic deformation and plastic intercalation, at the cost of high stresses, which fight against the bending forces and often lead to buckling instabilities, oscillations, and worse.

For example, a small circular disc subject to uniform extrinsic bending will yield a spherical cap, but beyond a certain critical curvature, it will spontaneously buckle cylindrically; the spherical conformation becomes unstable (see Figure 6). Ideally, plastic deformation would set in before buckling, and the equilibrium intrinsic curvature would relax toward a spherical configuration. This is difficult to achieve, however, requiring substrates that are plastically soft yet flexurally quite stiff, and the high stresses involved remain a liability.

The complementary strategy, actuating on intrinsic curvature, is similarly geometrically inconsistent but has some notable properties. Unlike extrinsic curvature, which cells can directly manipulate, the relationship between what a cell can do locally and the resulting effects on intrinsic curvature is quite nontrivial (given by the Brioschi formula). Small changes to curvature can be produced by each cell changing its size and shape – adjusting its aspect ratio, for example. The effect on curvature is then a function of the differences in changes expressed among nearby cells. However, large changes must be achieved by plastically re-arranging cells rather than simply distorting them, lest we demand that cells flatten into pancakes or stretch into spaghetti. A more useful actuator for large

Figure 6: Uniform disc subject to extrinsic curvature such as due to apical/basal constriction, showing spontaneous cylindrical buckling beyond a certain critical curvature. Curvature set-points of (a) 100mrad/cell, (b) 200mrad/cell, (c) 300mrad/cell, (d) 400mrad/cell. Substrate stiffness $k_B = 0.8$.

intrinsic curvatures is thus cell-cell traction, by which cells can intercalate with their neighbors (as illustrated in Figure 2).

How should cells exert traction forces in order to produce a given curvature? This is complicated. For the case of axisymmetric curvature, however, as in a spherical cap, the "purse string" strategy is an option: if curvature is too small, cells near the edge should pull on their circumferential neighbors, so as to shrink the mouth of the region. If curvature is too large, cells should pull on their radial neighbors, so as to enlarge it (see Figure 5).

This sort of boundary-focused purse-string traction (demonstrated in Figure 7) can be orchestrated, for example, by having the boundary emit a decaying gradient proportional in strength to the locally reported curvature error. The shape of the gradient then informs cells which direction and how hard to pull on their neighbors. The simplest approach might be to derive the orientation from the gradient vector or the level curves (choosing depending on the sign), and this works. Here, we use an alternative source, the principal axes of the Hessian (negative axis along the boundary, due to sources, and positive axis elsewhere), which appeared slightly more effective.

Formally, the cells attempt to compute by relaxation a gradient V,

$$\nabla^2 V = \gamma V \qquad (6)$$

$$V\big|_{d\Omega} = g_V (C_0 - C) \qquad (7)$$

where γ is a decay coefficient, C is the output of the curvature sensor, C_0 is the target curvature, and g_V is a gain parameter. Through traction, they apply a stress proportional to

$$\sigma_{ij} = sat(d^2V/dx_i dx_j - \delta_{ij}\nabla^2 V/2) \qquad (8)$$

where sat is an arctan-based saturation function. In practice, this is approximated by applying a traction proportional to the saturated magnitude of σ to each of the opposite vertices most closely aligned to the positive principal axis of σ (and with the aforementioned sign hack applied to boundary cells), with magnitude scaled to reach at most twice the ordinary surface tension of an

Figure 7: Spherical shell with lobe running curvature controller, showing purse-string style traction actuation (red arrows) and extrinsic bending set-points (blue shading); comparable run to Figure 8c. Black base region exerts random tractions and no extrinsic bending.

edge.[1]

The effects of such purse-string traction are several. The application of traction forces leads to net stresses and bending moments in the surface, tending to open up or close the mouth of the region, as intended. In response, cells intercalate as expected, circumferentially or radially, leading to changes in intrinsic curvature. However, so long as curvature error persists, the rearrangement is incessant. Re-orienting after each rearrangement, cells continue to grapple on one another, rearranging repeatedly. This continuing churn nullifies the yield strength of the cellular lattice and leads to viscous-like relaxation, which is both an asset and a liability. Churn relaxation is helpful because, as alluded to previously, it provides a natural mechanism for uninteresting and undesired degrees of freedom to relax and stabilize, without explicit control. It is problematic because the desired target degrees of freedom relax as well, making it difficult to sustain more than small deformations.[2]

The complementary problems exhibited by extrinsic bending and purse-string traction suggest that their combination might be more successful than either in isolation. Indeed, merely running them simultaneously, with no co-ordination, produces a drastic improvement. The combination of purse string

---

[1]Note that such actuation profiles are not scale-invariant, due to the fixed characteristic length scale of the gradient's decay. However, because the feedback sensors are scale-invariant, the resulting control algorithm is still quite flexible across a range of scales.

[2]There is also a subtle mathematical limitation to purse-string traction and other intrinsic actuation methods: they become singular when the surface is flat. Starting from a flat conformation, purse-string traction is weak and has no way to influence which way the surface will buckle. The sign of its influence depends on the sign of the existing extrinsic curvature.

traction as above and a trivial integral controller on extrinsic bending,

$$d\theta_0/dt = min(max(g_\theta(C_0 - C), -\dot{\theta}_{max}), \dot{\theta}_{max}) \tag{9}$$

(plus sensible limit stops), both using the same curvature feedback sensor, yields a stable and robust algorithm for producing spherical caps of arbitrary desired curvature. Figure 7 shows this tandem actuation mechanism in action, and Figure 8 illustrates the results for several different target values of curvature.

At first glance, one might expect that the two actuation mechanisms ought to be tightly correlated, so that consistent intrinsic and extrinsic curvatures would be produced. However, the precise combination turns out to be quite forgiving. As the integral controller governing extrinsic bending ratchets up, intrinsic churn relaxation begins to lead towards rather than away from the desired equilibrium. At the same time, as cells rearrange, both autonomously and deliberately, the stresses generated by inconsistent curvatures are relaxed. Indeed, even without any coherent direction at all to the traction forces – a traction random walk – the combination of traction and extrinsic bending is sufficient. Convergence is slower and stresses are higher, but it works. In general, the relative calibration of intrinsic and extrinsic control affects the time to convergence and the stress profile, but the ultimate equilibrium is robust.

4.3 Complex Structures from Simple Pieces

Now that we have the beginnings of an understanding of geometric control for simple, homogeneous regions, how might we proceed to more complicated structures? Rather than developing a slew of more complicated sensors, actuators, and controllers, each with multiple degrees of freedom, it would be simpler if we could instead assemble multiple elementary features along a body plan pattern, each feature region running some simple control law. With actuation controllers like our example above, however, simply cutting and pasting regions together does not work well. Controllers must behave compatibly along shared boundaries, or they will fight each other. Even if curvatures can be carefully selected to match up, evolvability is impaired, because further revisions will require consistent modifications in multiple places simultaneously.

Instead of directly coupling tightly controlled components to each other, a better strategy might be to connect them through special combiner regions (or "combinators", to borrow a term from computer science) – a special type of actuation controller that furnishes sort of weakly controlled glue to couple otherwise incompatible boundaries together. Instead of tightly specifying all properties of the structure, one could specify only certain key regions and features, relying on combiner regions to interpolate between them for the remainder. Such combiner regions would insulate individual components from the geometrical and mechanical side effects of other components, allowing their controllers to operate quasi-independently.

Through the principle of relaxation, simple combiners are constructed easily. For small structures, no active controller is needed, just a routine to ensure

(a)   (b)

(c)

Figure 8: Lobes with controlled curvature – spherical surfaces divided into two regions (via normal neighbors), where green pursues a target curvature using purse-string traction and extrinsic bending, while blue relaxes passively (see Section 4.3). Three different target curvatures are illustrated, with ratios $1 : 3 : 5$ respectively.

cells are reset their default properties. The churn injected from the jostling of neighboring regions' actuators is enough to cause mechanical relaxation, producing smooth connector regions with minimal curvature. For larger structures, it's necessary to add a controller that deliberately relaxes the surface through cell-cell traction; a simple random walk of traction will often suffice. A more aggressive approach, chosen here, is to use a smoothing geometric flow (in this case, exerting traction along the major axis of the Hessian of Gaussian curvature).

By definition, a weakly controlled relaxation combiner does little to dictate the relative positions of the regions it connects, beyond the topological constraints imposed by the body plan. Where, then, do the regions end up? The body plan patterning mechanism may initially lay out the connected regions in some predictable fashion, but they effectively "float" upon the combiner, and in the long run, they move to occupy positions that minimize mechanical energy. Typically, this process is dominated by the bending energy. Regions can be modeled, in a sense, as interacting by virtual forces, dependent on their curvatures. Regions of the same sign of curvature typically repel, while those of opposite sign attract. If the global conformation leads to the formation of a

(a)

(b)

Figure 9: Simple compound structures and their associated normal neighbors body plans: (a) 3-lobe structure where red, green, and cyan regions control curvature while blue combiner region relaxes geometry. (b) 4-lobe structure where the lobes are split across two combiners (yellow and blue).

bend in the combiner region, subsidiary regions may interact with this curvature as well. For example, when several regions of positive curvature float within a spherical combiner, they frequently align themselves along a circumferential ring, like spokes of a wheel (e.g. Figure 9a). Such virtual forces can often be relied on to produce a particular, final conformation in space.

Figure 9 shows a few examples of this approach, where independently controlled lobes are arranged by virtue of their interaction forces within a relaxation combiner. The number of lobes, their sizes and curvatures, and the divisions of the combiner can all be independently specified. However, there is no direct control available over the relative positions of the lobes. Even breaking the lobes into groups under different combiners does not meaningfully affect their positions (see Figure 9b); pure relaxation combiners are, to a good approximation, fully associative.

A more sophisticated combiner might manipulate the layout of its subsidiary regions by adding deliberate tractions and bending moments so as to customize the virtual interaction forces. More simply, however, we can break the asso-

17

(a)

(b)

(c)
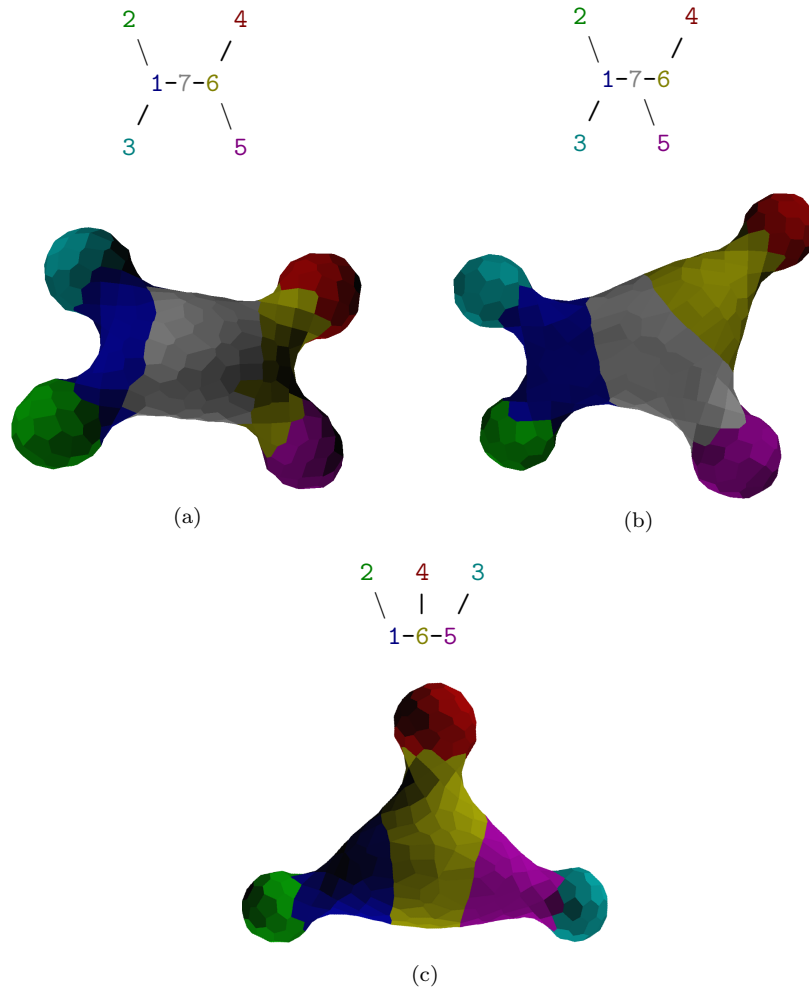
Figure 10: Compound structures using relaxation combiners with associativity broken by surface tension. Leaf nodes control curvature; non-leaf nodes are combiners. Combiner cells are configured with adhesive self-affinity and mutual disaffinity such that internal edge tension is reduced and mutual edge tension increased by (a) 40% and (b), (c) 80%. (The stronger surface tension in the latter two helps produce more distinct conical features.) Pattern regions are of unequal size in part due to deliberate adjustments to quorum feedback (k_q) – halved in leaf nodes, and, in (a), doubled in region 7.

ciativity of the combiners with additional passive forces and use the resulting non-associative combiners to produce more complex shapes. An easy way to do this is with differential adhesion, such that different combiners have mutual disaffinity and hence are shaped by the surface tension forces along their boundaries. Figure 10 shows several examples of structures grown this way.

5   Evaluation

In spite of our meager toolbox consisting of one control law and two closely related combiners, the variety of structures we can declaratively produce is beginning to get interesting. In this section we consider how well these structures exhibit the robustness properties I have claimed, including self-repair, approximate scale invariance, and tolerance of unexpected parameter variations. The first two properties are evaluated informally, while the third, and the associated role of partial redundancy, are investigated quantitatively.

Geometric self-repair follows naturally from the feedback control mechanism. Initial geometrical conditions in the form of a sphere, a cigar, and a mature, multi-lobed structure have all been evaluated. The final structures that result are indistinguishable, differing only in time to convergence.[3]

Approximate scale invariance can be demonstrated by running the same program on different size domains. Figure 11 demonstrates the program of Figure 10a running on different sized substrates. Using the same set of parameters as before, originally tuned for the middle size (400 cells), the small size (192 cells) works perfectly. The large size (900 cells) has a tendency to twin lobes but otherwise converges well (Figure 11b). Quantitatively, such persistent twinning occurs in about 30% of runs of the simple 3-armed structure of Figure 9a initiated on a 900-cell cigar, and more frequently on complex patterns such as Figure 10a.[4] Twinning originates in the body plan patterning algorithm and can be avoided by fine-tuning it: trading off speed for better convergence (e.g. Figure 11c, where the quorum sense morphogens have been configured to persist over longer distances) or adjusting the "temperature" parameter for the larger body size. Alternatively, twinning may be eliminated entirely (and convergence time drastically reduced) with no fine tuning at all by adding a partially redundant pre-pattern, as sketched out in Section 3.2. The resulting structure is indistinguishable from Figure 11c.

---

[3]Highly elongated initial conditions such as a thin cigar have, however, been observed to increase the occurrence of infrequent patterning defects such as transient twinning and pattern elements wrapping around the cylindrical circumference.

[4]In fact, the large version of such a complex structure develops with extensive temporary twinning showing fully-actuated curvature, which only resolves through churn and domain wall migration. The twinning that persists appears reminiscent of transient twinning rather than stable twinning. The highly curved lobe regions have a particular tendency to remain thus twinned, perhaps due to the influence of their mutual mechanical repulsion.

Figure 11: Program of Figure 10a running on different domain sizes. (a) Small, 192-cell domain. (b) Large, 900-cell domain, showing typical twinning that fails to resolve. (c) Large, 900-cell domain that avoids twinning via 10x reduction in decay rate of pattern region quorum sense morphogens.

5.1   Quantitative Evaluation

The most interesting case to explore is that of unexpected parameter variation. For this purpose, substrate stiffness is varied here. This additionally affords the opportunity to explore the relative roles of the two actuation mechanisms in tandem. When substrates are stiffer, one should expect the extrinsic actuation to be more powerful, while on softer substrates, intrinsic actuation should be stronger. Several variants of the control algorithm are evaluated, with modifications to reduce or eliminate functionality of one or the other actuation mechanism.

In the absence of an evolutionary fitness function, however, quantitative evaluation of the robustness of morphogenesis is a challenge. Structures may be complicated, and some variation in structure is entirely allowable, whether due to run to run nondeterminism or changing environmental parameters. Comparison against a single golden output would introduce spurious biases. Even measuring the time to steady-state convergence is fraught, because the algorithms presented here never completely stabilize but instead wander slowly around the envelope of acceptable structures.

To avoid these difficulties, I evaluate specially-designed structures on the development of symmetries emergent within their body plans. Trials here are conducted using the 3-armed structure of Figure 9a, starting from spherical

initial conditions, which under any successful developmental trajectory should discard its initial shape and generate a strong 3-fold prismatic symmetry. This symmetry is evaluated using low-frequency spherical harmonic spectral components, in the form of a hand-constructed score,

$$D3hScore = \frac{|Y_3^3|}{|Y_0^0|/2 + 2|Y_2^2| + 2\sum_{m\in 0..2}|Y_m^3|} \tag{10}$$

where the Z-axis is aligned with the minor axis of covariance. Trials are run until a pre-specified threshold score is reached or a failure or timeout is encountered, and experimental configurations are evaluated on the basis of success rate and mean time to completion.

Table 1 summarizes the results of trials under several different values of bending stiffness constant k_B and with several different "knockout variants" of the curvature control algorithm. As claimed, only mechanisms that combine both extrinsic bending and traction able to succeed reliably in all cases (and at all with the stiffer substrates). In all cases, the full, combined algorithm is both fastest and most reliable. Interestingly, there are cases where each of the other mechanisms are still viable. A lack of *directed* traction is a hindrance, but in two of the three cases, only inasmuch as it reduces the speed of convergence. With high stiffness, the total loss of traction can be tolerated. With low stiffness, some patterns develop reliably even without bending (although their precise shapes are visibly altered). Table 2 illustrates some example structures.

The tandem actuation mechanism thus exhibits partial redundancy: for many situations, multiple overlapping mechanisms are available, such that reduced function or complete failure of one pathway is quite survivable. However, due to the physical constraints of the problem, employing the full complement of mechanisms is often still helpful and sometimes absolutely necessary. The resulting combination mechanism is quite robust but irregularly so, giving confusing and seemingly contradictory results to knock-out experiments: Is the bending pathway necessary for curvature development? Is the traction pathway necessary for curvature development? Is the gradient field that directs traction necessary for curvature development? Differing conditions may produce differing "answers" to these questions. The situation is surprisingly reminiscent of the difficulties encountered in knockout experiments on real, live organisms [22, 26].

6 Discussion: Partial Redundancy

The canonical benefit of partial redundancy is resistance to rot, hidden degradation. The more "verbose" a system – more details in its specification, more parts in its realization – the more vulnerable it is to random damage. Yet, the above examples, each combining two partially redundant mechanisms, fared quite well under both knock-out damage and mechanical disruption. By making degradation visible and detrimental yet survivable, partial redundancy facilitates homeostasis – both at the genetic level, through cross-over and selection,

	$k_B = 0.8$	$k_B = 2.4$	$k_B = 8$
Tandem actuation	1714	1945	1900
	0%[a]	*0%*	*0%*
Tandem w/half-strength traction	2344	4581	2952
	0%	*0%*	*0%*
Bending + random traction	3781	2781	3298
	0%	*3%*	*47%*[b]
Bending only	5597	4194	3360
	0%[bd]	*25%*[b]	*20%*[b]
Traction only	3089	—	—
	0%[cd]	*100%*[c]	*100%*[c]

[a]Failure rates approaching 100% due to lobes pinching off are observed here under certain altered parameters. Evidence suggests this is due to excessively strong actuation collapsing the base of the lobe rather than allowing sufficient time for the main combiner body to slowly relax. Pinch-off can be prevented by putting tighter limit stops on actuation of *either* bending angle or traction strength, or by reducing the shear stiffness (i.e. internal surface tension) of the combiner body.

[b]Failure declared when simulator fails to make progress after 10000 mechanical iterations. This is generally due to geometric pathologies, often tight, hyperbolic creases, which suggest a crumpling or tearing of the surface.

[c]Unsuccessful cases never produce definitive lobes; only slight curvatures form.

[d]More complex patterns have a significant failure rate.

Table 1: Development time and failure rates for the 3-arm structure across varying substrate bending stiffnesses k_B under default algorithm and several "knockout" variants. Mean agent iteration count for success (reaching D3h score of 0.8) and percent failure rate in a minimum of 18 runs per category.

and at the somatic level, through regeneration – making complex and verbose systems sustainable.

The benefits run deeper. The case study here showed how to use partial redundancy as a weapon to attack a messy, hard-to-characterize system. No mechanism alone furnished an exact solution, but each was able to cover for the other's bugs and limitations. In case of actuation, even when both mechanisms were not essential, performance was better with both, if for no other reason than that the total force could be increased by combining multiple modes of actuation.

Partial redundancy also facilitates modularity: redundant mechanisms may be repurposed to new functions, and stresses placed on the system by changes are more easily absorbed. Companion techniques, such as passive homeostasis by relaxation (as in combiners), further help to neutralize cross-interactions between components. Partial redundancy thus fosters exploration.

7 Conclusions and Future Work

This paper explored development and regeneration as single framework, morphological homeostasis via explicit feedback control. Focusing on mechanical

	$k_B = 0.8$	$k_B = 2.4$	$k_B = 8$
Tandem actuation			
Tandem w/half-strength traction			
Bending + random traction			*
Bending only			
Traction only		*	*

Table 2: Ensemble of 3-arm structures generated with the same random number seed but differing actuation algorithms and stiffness parameters. Asterisk (*) marks instances considered failures.

remodeling rather than cell proliferation, several techniques were proposed. Ultimately, no one technique was best; instead, partially redundant combinations were fastest and most robust. Complex structures were then produced by introducing "combiners", using passive relaxation to decouple key features. A unifying theme, with applications to both biomimetic design and developmental theory, was partial redundancy and the feedback it entails. Still, many questions remain.

The pinch-off pathology, briefly mentioned in Table 1, represents a larger problem only crudely addressed: substrates have limits, beyond which they fail. Actuation must be careful not to exceed these limits, or it will destroy its own substrate. The solution suggested here, enforcing fixed bounds on actuator outputs, is crude both because it is hand-tuned and because it may unnecessarily limit outputs (and hence speed and control authority) even where there is no imminent danger of damage. A more elegant mechanism might be for the substrate to recognize its own limits and express "pain" when over-exerted, causing actuation to back off [1].

A significant limitation with the approach in this paper is that all patterning happens simultaneously in a single stage, which is both biologically unrealistic and limits the amount of complexity that can be implemented without getting stuck in local minima. Hierarchical and cascaded patterning would alleviate this limitation, but how can such sequential mechanisms be reconciled with

regeneration? The answer is not clear; perhaps backtracking is involved.

The strategy of partial redundancy is not limited to physical or biological systems. For example, multiple versions of a software library might, like chromosomes, run in tandem. Different, loosely-coupled mechanisms might cooperate to ensure system homeostasis and sustainable resource usage. Virtual "pain" mechanisms might restrain over-taxing activities. The possibilities for bio-mimetic software systems are wide open.

Acknowledgements

References

[1] Jacob Beal. Functional blueprints: An approach to modularity in grown systems. In *International Conference on Swarm Intelligence*, 2010.

[2] Ariela Benigni, Marina Morigi, and Giuseppe Remuzzi. Kidney regeneration. *The Lancet*, 375(9722):1310 – 1317, 2010.

[3] Micah Z. Brodsky. Self-timed patterning. In *7th International Workshop on Spatial Computing (SCW 2014)*, May 2014.

[4] Micah Z. Brodsky. Spatial patterning with the rule of normal neighbors (ext. abstract). In Hiroki Sayama et al., editors, *Artificial Life 14*, pages 817–818. MIT Press, 2014.

[5] Micah Z. Brodsky. *Synthetic Morphogenesis: Space, time, and deformation*. PhD thesis, MIT, 2014. http://dspace.mit.edu/handle/1721.1/92963.

[6] Xiaoguang Chen and G. Wayne Brodland. Multi-scale finite element modeling allows the mechanics of amphibian neurulation to be elucidated. *Physical Biology*, 5(1):015003, 2008.

[7] Enrico Coen. *The Art of Genes: How Organisms Make Themselves*. Oxford University Press, 2000.

[8] Frank Dellaert and Randall D. Beer. A developmental model for the evolution of complete autonomous agents. In *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 393–401. MIT Press, 1996.

[9] Jean Disset, Sylvain Cussat-Blanc, and Yves Duthen. Self-organization of symbiotic multicellular structures. In Hiroki Sayama et al., editors, *Artificial Life 14*, pages 541–548. MIT Press, 2014.

[10] Rene Doursat. The growing canvas of biological development: Multiscale pattern generation on an expanding lattice of gene regulatory networks. *InterJournal: Complex Systems*, 1809, 2006.

[11] René Doursat et al. Embryomorphic engineering: Emergent innovation through evolutionary development. In *Morphogenetic engineering: toward programmable complex systems*, pages 275–311. Springer, 2012.

[12] Avigdor Eldar, Ruslan Dorfman, Daniel Weiss, Hilary Ashe, Ben-Zion Shilo, and Naama Barkai. Robustness of the BMP morphogen gradient in Drosophila embryonic patterning. *Nature*, 419(6904):304–308, 2002.

[13] Diego Federici and Tom Ziemke. Why are evolved developing organisms also fault-tolerant. In *Proceedings of the 9th International Conference on Simulation of Adaptive Behavior (SAB '06)*, pages 449–460, 2006.

[14] Alexander G. Fletcher, Miriam Osterfield, Ruth E. Baker, and Stanislav Y. Shvartsman. Vertex models of epithelial morphogenesis. *Biophysical Journal*, 106(11):2291 – 2304, 2014.

[15] V. French. Pattern regulation and regeneration. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 295(1078):601–617, 1981.

[16] Chikara Furusawa and Kunihiko Kaneko. Robust development as a consequence of generated positional information. *Journal of Theoretical Biology*, 224(4):413 – 435, 2003.

[17] C. Gershenson, S. A. Kauffman, and I. Shmulevich. The role of redundancy in the robustness of random Boolean networks. In L. M. Rocha, L. S. Yaeger, et al., editors, *Artificial Life X*, pages 35–42. MIT Press, 2006.

[18] Alfred Gierer et al. Regeneration of hydra from reaggregated cells. *Nature/New Biology*, 239(88):98–101, 1972.

[19] François Graner and James A. Glazier. Simulation of biological cell sorting using a two-dimensional extended Potts model. *Phys. Rev. Lett.*, 69:2013–2016, September 1992.

[20] Michal Joachimczak, Taras Kowaliw, René Doursat, and Borys Wróbel. Controlling development and chemotaxis of soft-bodied multicellular animats with the same gene regulatory network. In Pietro Liò et al., editors, *12th European Conference on the Synthesis and Simulation of Living Systems (ECAL 2013)*, pages 454–461. MIT Press, 2013.

[21] Marc W. Kirschner and John C. Gerhart. *The Plausibility of Life: Resolving Darwin's Dilemma*. Yale University Press, 2005.

[22] Yuri Lazebnik. Can a biologist fix a radio? Or, what I learned while studying apoptosis. *Cancer Cell*, 2(3):179 – 182, 2002.

[23] Hans Meinhardt. A model for pattern-formation of hypostome, tentacles, and foot in hydra: how to form structures close to each other, how to form them at a distance. *Developmental Biology*, 157:321–333, 1993.

[24] Jay E. Mittenthal. The rule of normal neighbors: A hypothesis for morphogenetic pattern regulation. *Developmental Biology*, 88(1):15 – 26, 1981.

[25] Radhika Nagpal. *Programmable Self-Assembly: Constructing Global Shape using Biologically-inspired Local Interactions and Origami Mathematics*. PhD thesis, MIT, 2001.

[26] Martin A. Nowak, Maarten C. Boerlijst, Jonathan Cooke, and John Maynard Smith. Evolution of genetic redundancy. *Nature*, 388(6638):167–171, 1997.

[27] Tina Rich, Rachel L. Allen, and Andrew H. Wyllie. Defying death after DNA damage. *Nature*, 407(6805):777–783, 2000.

[28] C. H. Waddington. Canalization of development and the inheritance of acquired characters. *Nature*, 150:563–565, November 1942.

[29] Daniel Yamins. *A Theory of Local-to-Global Algorithms for One-Dimensional Spatial Multi-Agent Systems*. PhD thesis, Harvard, December 2007.