

Partial Redundancy and Morphological Homeostasis: Reliable Development through Overlapping Mechanisms

Micah Brodsky¹

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

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? I show how these questions may have a common answer, a process of morphological homeostasis built from overlapping, partially redundant mechanisms.

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 (1942), 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 organisms within a single phylum start from dramatically different initial conditions yet converge to a common “phylogenetic” stage of development, before differentiating into their characteristic larval forms (Kirschner and Gerhart, 2005). 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 (Gierer et al., 1972).

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?

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”.

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 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. 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. 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. 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. Instead, over evolutionary time, any accidental copies that remain diverge and acquire new functions. 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 of Brodsky (2014c, 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. The remarkably robust results, evaluated informally, show how effective homeostasis through partial redundancy can be.

Model Background

On stylized, deterministic substrates like cellular automata, the value of feedback and redundancy is not so apparent, but deep challenges surface 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.

Recent years have seen the use of increasingly sophisticated physical models (e.g. Chen and Brodland (2008); Disset et al. (2014); Doursat et al. (2012)). The model employed here (Brodsky, 2014c, Ch. 2), unlike the more common mass-spring models, improves 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 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. Cell shapes, and hence vertex positions, are governed by surface tension and internal elasticity. The model is extended into 3d by the addition of flexural springs at each cell-cell junction, the spring constant determining stiffness.

Cells are regulated by simple software agents, the realization in terms of genes not a focus of this work. Cells can sense properties of their mechanical conformation such as elongation or curvature and can influence it through neighbor-neighbor tractions (an adjustment to edge tension) 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.

Decomposing the Problem

Natural biological structures are complicated, combining multiple subparts with differing characteristics. 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 pattern affects how actuation efforts interact, and updates to the global pattern

require 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 (Kirschner and Gerhart, 2005), 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 and local actuation are both 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 (Brodsky, 2014a), 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 (1993)). Morphogenetic fields with self-sustaining sources (e.g. as in Doursat et al. (2012)) might be usable, with some caveats due to geometry (Brodsky, 2014c, Ch. 4). However, the normal neighbors patterning mechanism of Brodsky (2014b,c, Ch. 4), where patterns are specified through a purely topological description (an adjacency graph) and maintained continuously – hence tolerating substantial distortion – fits almost perfectly.

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, I don’t consider changes in cell number here). 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 de-

grees 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.

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 1). First, we need a distributed, scale-invariant measure of curvature, built from local sensors.

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, measuring extensive quantities seems to require leader election or an equivalent broken symmetry (Brodsky, 2014c, 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 a local measure of curvature – for example, multiplying by average mean curvature – to produce a scale-invariant measure of global curvature.

By trial and error, I found an interesting variation that worked well. Rather than combining the ratio of area to perimeter with another global measure, I combined it with a purely local measure of curvature, producing a hybrid measure that is partly local, partly global. This mirrors the effects of actuation, also partly local, partly global. The measure I found most effective 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 cocktail is, interestingly, an example of sensor-level partial redundancy.

Actuating Curvature

Now that we have a sensor for curvature, we must build an actuator. How? 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. 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 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 size, it will spontaneously buckle cylindrically; the spherical conformation becomes unstable. 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 rearranging cells rather than simply distorting them, lest we demand that cells flatten into pancakes or stretch into spaghetti. A more useful actuator for large intrinsic curvatures is thus cell-cell traction, by which cells can intercalate with their neighbors.

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.

This sort of boundary-focused purse-string traction 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 ori-

entation from the gradient vector or the level curves (choosing depending on the sign), and this works. I used an alternative source, the principal axes of the Hessian (negative axis along the boundary, due to sources, and positive axis elsewhere), which seemed slightly more effective.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. Reorienting 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 coordination, produces a drastic improvement. The combination of purse string traction as above and an integral controller on extrinsic bending, both using the same curvature feedback sensor, yields a stable and robust algorithm for producing spherical caps of arbitrary desired curvature. Figure 1 shows this tandem actuation mechanism in action, illustrating 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

¹Note 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.

²There 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.

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.

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, I found that no active controller is needed, just a routine to ensure 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 might use a smoothing geometric flow (e.g. 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 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 2a). Such virtual forces can often be relied on to produce a particular, final conformation in space.

Figure 2 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 2b); 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 associativity 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 3 shows several examples of structures grown this way.

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. It remains to be shown that the structures exhibit the robustness properties I have claimed, including self-repair, approximate scale invariance, and tolerance of unexpected parameter variations.

Geometric self-repair follows easily from the feedback control mechanism. One can even take geometric results of running under one program, switch to a different program, and watch the structure reform. The results are indistinguishable from structures produced starting with a sphere.

Approximate scale invariance can be demonstrated by running the same program on different size domains. Figure 4 demonstrates the program of Figure 3a 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




Figure 1: 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 ‘Complex Structures from Simple Pieces’). Three different target curvatures are illustrated, with ratios 1 : 3 : 5 respectively.




Figure 2: 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).

well (Figure 4b).³ Such twinning can be avoided with adjustments to the body plan patterning algorithm, trading off speed for better convergence (e.g. Figure 4c, where the quorum sense morphogens have been configured to persist over longer distances) or simply tuning for the larger size (e.g. by increasing the “temperature” parameter).

The most interesting case to explore is that of unexpected parameter variation. For this purpose, I vary the stiffness of the substrate. This also 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.

Table 1 summarizes the results of informal trials under several different values of bending stiffness constant k_B and

³In fact, the large size develops with extensive temporary twinning showing fully-actuated curvature, which only resolves through churn and domain wall migration. The highly curved lobe regions have a particular tendency to remain twinned, in spite of the body plan, probably due to the influence of their mutual mechanical repulsion.

with several different “knockout variants” of the curvature control algorithm. As claimed, only mechanisms that combine both extrinsic bending and traction are able to succeed in all cases (and at all with the middle stiffness). A lack of *directed* traction is a hindrance, but only inasmuch as it reduces the speed of convergence. Interestingly, there are cases where each of the other mechanisms still succeed. With high stiffness, one hardly notices the total loss of traction. With low stiffness, some patterns develop successfully even without bending (although their precise shapes are visibly altered).

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 cur-

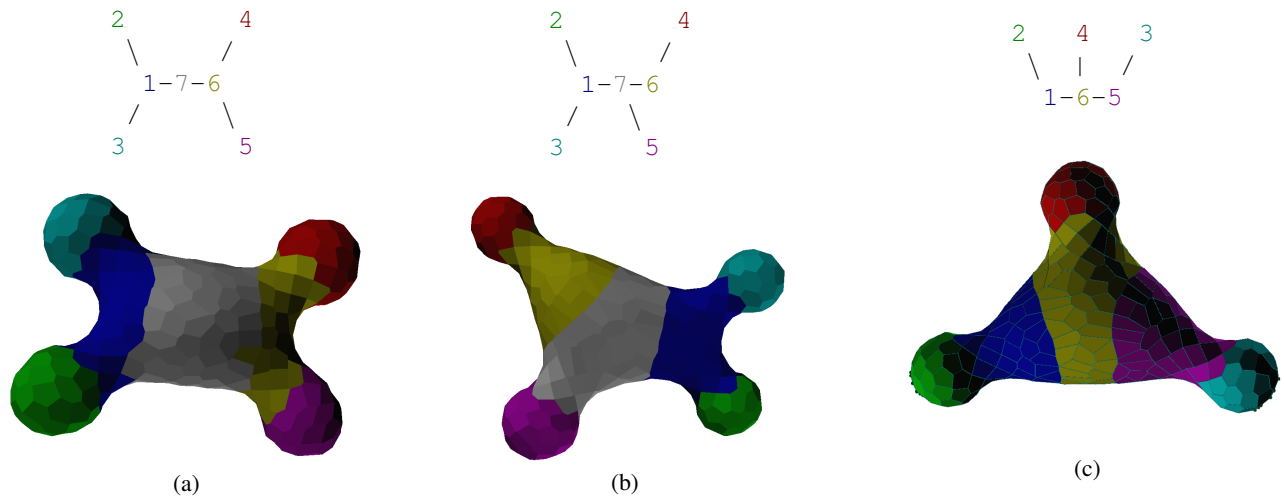


Figure 3: Compound structures using relaxation combiners with associativity broken by surface tension. Leaf nodes control curvature; non-leaf nodes are combiners. Combiner cells have adhesive self-affinity and mutual disaffinity such that internal edge tension is reduced and mutual edge tension increased 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.




Figure 4: Program of Figure 3a 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.

vature 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 (Lazebnik, 2002).

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 example, 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 home-

ostasis – both at the genetic level, through cross-over and selection, and at the somatic level, through regeneration – making complex and verbose systems sustainable.

The benefits run deeper. The example here showed how to use partial redundancy as a weapon to attack a messy, hard-to-characterize system. Neither mechanism alone furnished an exact solution, but each was able to cover for the other’s bugs and limitations. 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

	$k_B = 0.8$	$k_B = 2.4$	$k_B = 8$
Tandem actuation	Default: Fail ^a Reduced limit stops: Ok	Ok	Ok
Bending only	Marginal (slow; complex patterns fail) ^b	Marginal (slow; high failure rate) ^b	Ok
Traction only	Marginal (complex patterns are slow or unsuccessful) ^c	Unsuccessful ^c	Unsuccessful ^c
Bending + random traction	Slow	Slow	Ok

^aFails by a lobe pinching off. I hypothesize 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. The latter gives somewhat more consistent shapes.

^bFails through the development of tight, hyperbolic creases.

^cUnsuccessful cases never produce definitive lobes; only slight curvatures form.

Table 1: Summary of results with varying substrate bending stiffness k_B for default algorithm and several “knockout” variants.

cross-interactions between components. Partial redundancy thus fosters exploration.

Conclusions and Future Work

This paper explored development and regeneration as single framework, morphological homeostasis via explicit feedback control. Focusing on mechanical 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 used 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 (Beal, 2010).

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.

References

- Beal, J. (2010). Functional blueprints: An approach to modularity in grown systems. In *International Conference on Swarm Intelligence*.
- Brodsky, M. Z. (2014a). Self-timed patterning. In *7th International Workshop on Spatial Computing (SCW 2014)*.
- Brodsky, M. Z. (2014b). Spatial patterning with the rule of normal neighbors (ext. abstract). In Sayama, H. et al., editors, *Artificial Life 14*, pages 817–818. MIT Press.
- Brodsky, M. Z. (2014c). *Synthetic Morphogenesis: Space, time, and deformation*. PhD thesis, MIT.
- Chen, X. and Brodland, G. W. (2008). Multi-scale finite element modeling allows the mechanics of amphibian neurulation to be elucidated. *Physical Biology*, 5(1):015003.
- Disset, J., Cussat-Blanc, S., and Duthen, Y. (2014). Self-organization of symbiotic multicellular structures. In Sayama, H. et al., editors, *Artificial Life 14*, pages 541–548. MIT Press.
- Doursat, R. et al. (2012). Embryomorphic engineering: Emergent innovation through evolutionary development. In *Morphogenetic engineering: toward programmable complex systems*, pages 275–311. Springer.
- Gierer, A. et al. (1972). Regeneration of hydra from reaggregated cells. *Nature/New Biology*, 239(88):98–101.
- Kirschner, M. W. and Gerhart, J. C. (2005). *The Plausibility of Life: Resolving Darwin’s Dilemma*. Yale University Press.
- Lazebnik, Y. (2002). Can a biologist fix a radio? Or, what I learned while studying apoptosis. *Cancer Cell*, 2(3):179 – 182.
- Meinhardt, H. (1993). 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.
- Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150:563–565.