Current Biology, Vol. 12, 778–786, May 14, 2002, ©2002 Elsevier Science Ltd. All rights reserved. PII S0960-9822(02)00839-4

Robustness, Flexibility, and the Role of Lateral Inhibition in the Neurogenic Network

Eli Meir,[1,2,3] George von Dassow,[1,2] Edwin Munro,[2] and Garrett M. Odell[1,2]
[1]Friday Harbor Laboratories
University of Washington
620 University Road
Friday Harbor, Washington 98250
[2]Department of Zoology
University of Washington
24 Kincaid Hall
Seattle, Washington 98195

Summary

Background: Many gene networks used by developing organisms have been conserved over long periods of evolutionary time. Why is that? We showed previously that a model of the segment polarity network in *Drosophila* is robust to parameter variation and is likely to act as a semiautonomous patterning module. Is this true of other networks as well?
Results: We present a model of the core neurogenic network in *Drosophila*. Our model exhibits at least three related pattern-resolving behaviors that the real neurogenic network accomplishes during embryogenesis in *Drosophila*. Furthermore, we find that it exhibits these behaviors across a wide range of parameter values, with most of its parameters able to vary more than an order of magnitude while it still successfully forms our test patterns. With a single set of parameters, different initial conditions (prepatterns) can select between different behaviors in the network's repertoire. We introduce two new measures for quantifying network robustness that mimic recombination and allelic divergence and use these to reveal the shape of the domain in the parameter space in which the model functions. We show that lateral inhibition yields robustness to changes in prepatterns and suggest a reconciliation of two divergent sets of experimental results. Finally, we show that, for this model, robustness confers functional flexibility.
Conclusions: The neurogenic network is robust to changes in parameter values, which gives it the flexibility to make new patterns. Our model also offers a possible resolution of a debate on the role of lateral inhibition in cell fate specification.

Introduction

In this paper, we use a computer model to explore the properties of the neurogenic network, originally characterized in *Drosophila melanogaster*. This is but one example of the many networks of cross-regulatory genes at work in complex organisms. Other familiar examples include the networks of segment polarity genes, of cell cycle genes, of circadian clock genes, and so on. Each of these seems to have remained more or less intact through long periods of evolutionary time and across many phyla. Many of these networks have acquired novel functions in different organs and species. Why is that? Is there something special about each of these networks that causes natural selection to preserve them through deep time and coopt them to perform new functions? Or is it simply that evolution doesn't fix what first evolved if it isn't broken?

Only recently can we begin asking these network-level questions. A tremendous and ongoing experimental effort has generated an unprecedented catalog of the components and connections of many networks, while advances in computers give us the power to explore complex systems in ways unimaginable a decade ago. Recently, we found that a model of the segment polarity network in *Drosophila* embryos exhibits two very interesting properties, robustness to changes in kinetic parameters and robustness to changes in its initial conditions (i.e., due to prior expression patterns of upstream genes) [1]. This example is tantalizing, but it takes several such case studies to support general conclusions. We therefore investigate in this paper whether a second well-studied network, the neurogenic and proneural genes (hereafter simply referred to as "the neurogenic network"), has similar properties.

We use the "canonical" structure of this network as it operates during *Drosophila* development, especially in determining neuroblasts (NB) in embryos and sensory organ precursor (SOP) cells in imaginal disks (reviewed in [2]). Components of this network are used in a wide range of patterning processes, from vertebrate retinal development [3] to nematode anchor cell specification [4], with apparently similar interactions among the core genes. Neurogenesis is often cited as a classic case of lateral inhibition [5]. This hypothesis postulates that a small group of neighboring cells all start out competent to assume a particular fate. A stochastic fluctuation or external cue slightly favors one cell over the others. The favored cell then suppresses (laterally inhibits) its neighbors to keep them from also assuming that fate.

Here, we first ask whether the neurogenic network, as the literature currently portrays it, can perform lateral inhibition. Simple models encapsulating the barest essential facts of the neurogenic network (the Dl/N interaction) succeed under certain conditions in differentiating cells with initially similar levels of Dl and N (e.g., [6]). The real pathway between Notch activation and Delta expression, however, has layers of feedback and modulation, including switch and homeostat-like subcircuits. These ought to (and do) make it more difficult to accomplish lateral inhibition. Thus, it is a nontrivial question to ask under what conditions, if any, the neurogenic network can accomplish lateral inhibition. And, if it can, how do we explain the inclusion of design features that seem to work against the apparent biological role of the circuit? This is where we begin our exploration.

Abstracting the Neurogenic Network and the Three Tasks It Performs

We focused our investigations of the neurogenic network on the genes that select neural cells in *Drosophila*

[3]Correspondence: meir@beakerware.org

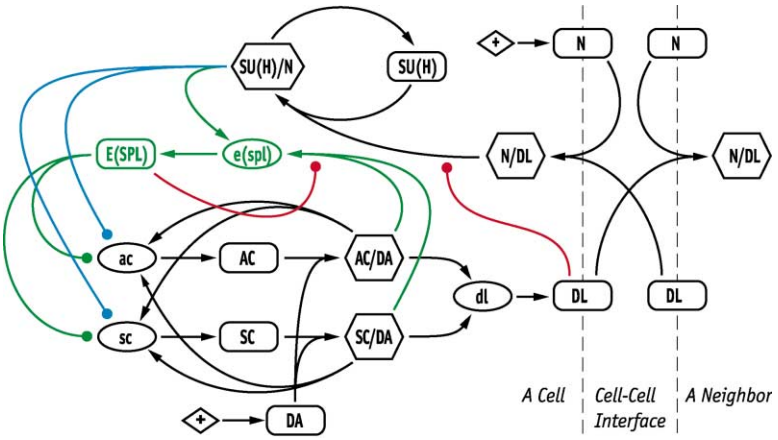

Figure 1. Models of the Neurogenic Network Used in Neural Specification in *D. melanogaster*

A graphical summary of the experimental literature on neural specification, which our model represents mathematically. Ovals represent mRNAs, rectangles represent proteins, and hexagons represent protein complexes. Diamonds containing plus signs indicate constitutive expression of N and DA where we have not diagrammed, but have modeled, the intermediate mRNA concentrations. Though not diagrammed, all gene products decay with parameter-specified decay rates, except SU(H), which we modeled as cycling between two states with no production or decay. Our augmented network model includes the black, green, and red connections. Our standard network model includes the black and green connections. Our reduced network model includes the black and blue connections such that SU(H)/N directly inhibits production of ac and sc with all E(SPL) products absent in that version. We adopt the following convention to distinguish between reality and model: as in the *Drosophila* literature, "*ac*" and "Ac" refer, respectively, to the *achaete* gene and its protein product; whereas, in referring to the model, "ac" and "AC" refer to the nodes in our model that represent *achaete* mRNA and protein.

embryos and imaginal disks. Figure 1 shows our summary of the core genes, their products, and their interactions. In crafting Figure 1, we approached the model-building process as a biochemist approaches in vitro reconstitution; by adding to the system piece by piece, we hope to figure out how each design feature contributes to the function of the essential core network. We rationalize our choice of this diagram in the Supplementary Material available with this article online, with a synopsis as follows (Below, "*ac*" and "Ac" refer to the real *achaete* gene and its protein product, whereas "ac" and "AC" refer to corresponding nodes in the model):

Delta (Dl) is a ligand for the receptor Notch (N). When Dl activates N, a cleaved-off cytoplasmic piece of N binds to the transcription factor Suppressor of Hairless (Su(H)), and that heterodimer activates *Enhancer of split* (*E(spl)*) complex genes. The proneural genes *achaete* (*ac*) and *scute* (*sc*) encode transcription factors that actually specify neural fate. Both Ac and Sc are autoactivating and cross-activating: they promote their own, and each others', transcription. Thus, the proneural genes constitute a bistable switch at the heart of the neurogenic network. They also activate transcription of *E(spl)* and *Dl*. E(spl) in turn represses transcription of *ac* and *sc*. Thus, the loop works as follows: something activates *ac* and/or *sc* in the neural-competent cluster. They upregulate *Dl*, whose product activates N in neighboring cells, which, through Su(H), activates *E(spl)*. E(spl) represses *ac* and *sc* in those neighboring cells. To achieve a neural fate, a cell must upregulate *ac* and *sc* enough that their autoactivation overwhelms E(spl)-mediated repression due to neighboring cells signaling through N.

We constructed three different models of the network in Figure 1, which we call "augmented", "standard", and "reduced". The standard network includes all components and interactions shown in Figure 1, except for *cis*-negative regulation of N activity by Dl and E(spl) autorepression (Figure 1 without red or blue connections). Experimental evidence for each of the latter interactions exists (see the Supplementary Material), but the literature has not given them much attention. Neither did we initially, but our results below regarding the augmented network (which adds the red connections) indicate that these may indeed be important. Our reduced network eliminates intracellular negative feedback from AC and/or SC to suppress ac and sc transcription (blue connections replacing red and green connections and their E(spl) hub). Such a simplified network could have functioned in a precursor to the *Drosophila* network since the similar process of anchor cell specification in the worm *Caenorhabditis elegans* appears to take place without E(spl)-like genes or function (X. Karp and I. Greenwald, personal communication) (though the *C. elegans* version is likely derived).

We explored three spatial patterning functions with each network (Figure 2). The first was lateral inhibition, leading to a separation in states between two neighboring competent cells, as in a related network in *C. elegans* [4]. We started two cells with different concentrations of ac and sc products and tested whether, within 300 min, the initially higher cell achieved a high concentration while the other turned off AC (see the legend of Figure 2). The second test asked the central cell in a group of seven to "win", inspired by SOP-forming clusters in an epithelial sheet [7]. The third test asked a double line of cells to refine to a single line, very roughly mimicking the pattern seen at the margin of wing imaginal disk [8]. We made and ran all models using our Ingeneue software [1, 9], a general-purpose program for modeling genetic networks. These tasks are intended to be stereotypes that abstract many biological phenomena rather than accurate mimics of any one case, and we designed them to insure that the network must use lateral inhibition in order to succeed. All files needed to recreate our results (as well as the Supplementary Material accompanying this paper) are available at www.ingeneue.org.

All Three Work!

Our first question was whether the three versions of the network in Figure 1 can perform lateral inhibition. A network model consists not only of its topology (Figure 1), but also of dozens of parameters specifying the strength, rate, and functional form of each connection.

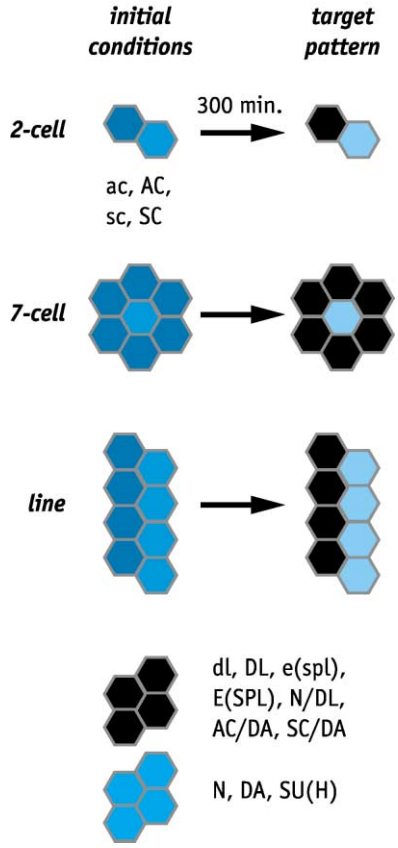

Figure 2. 2-Cell, 7-Cell, and Line Patterns Used to Test Our Models

Each test starts some cells with twice the concentrations of the mRNA and protein of both proneural genes as its neighbors (brighter colors show higher concentrations). We surround the field of competent cells with other cells (not drawn) containing no initial ac/sc products; hence, no ac/sc ever. All other model concentrations start either off (black) or at a uniform high level (blue) across all cells. We run the model for 300 min and measure whether the initially higher cell(s) achieve(s) a high concentration of AC (greater than 20% of the highest steady-state concentration possible) and whether the initially lower cell(s) turn off (below 2% of maximum steady state). This could happen without lateral inhibition if there was a threshold of initial ac and sc concentrations above which the ac/sc switch would turn on and below which it would turn off in isolated cells. To ensure that the separation is due to lateral inhibition rather than mere thresholding, we *always* run a second test with the same 2-fold initial difference between cells, but in which the lower-concentration cells in the second test have the same initial concentrations of *ac* and *sc* products as the higher cell in the first test (the one exception was in testing the effect of noise in initial conditions in which we did not use a second test). We call a model with its parameter values a "solution" only if both pairs achieve the correct pattern at 300 min and, additionally, if both pairs maintain that state for 1000 min. We conducted identical tests using much smaller initial concentration differences (down to 1%) and found qualitatively similar results, albeit with commensurately lower solution frequencies.

The reduced, standard, and augmented versions of the network have 53, 63, and 69 such parameters, respectively. Asking whether the network can do lateral inhibition amounts to asking whether there exist values for these parameters that enable the network to pass any of the three tests in Figure 2. We have practically no idea what appropriate values for these parameters would be, so we instantiated networks by sampling values randomly from within broad but biologically realistic ranges [1, 9]. Any set of randomly chosen parameter values that allows the model to pass some functional test we call a "solution".

All three models exhibited lateral inhibition in the 2-cell test, but the standard model had a much lower frequency of solutions than the augmented model, while the reduced model had the highest solution frequency of the three (Table 1). By this naive measure, the reduced model seems the most robust of the three. However if p is the probability of picking a "good" value for each parameter independently, we expect that the solution frequency for models with equivalent robustness per parameter would scale as p^n, where n is the number of parameters. p is highest for the augmented model (Table 1). That measure makes the augmented network most robust. All three networks could also mediate lateral inhibition between a central neural cell and six surrounding cells (7-cell pattern, Figure 2), with similar solution frequencies as for the 2-cell pattern (Table 1). As with the segment polarity network [9], cooperativity (Hill coefficients > 1) in transcriptional activation and repression is essential for the network to function (Table 1). In many neurogenic tissues, when the winning cell is removed early in the patterning process, a cell originally destined for a nonneural fate replaces it. We took parameter sets from our 2-cell test in the augmented network and removed the presumptive winner at a point at which concentrations of proneural genes in the loser were beginning to drop toward zero. We found parameters in which the presumptive loser could recover at that point and assume a neural fate. We also found the same result using a 3-cell test (the largest number in which all cells can still touch in our hexagonal grid).

There are two differences between our standard and augmented models, E(SPL) autoinhibition and DL inhibition of N activation (red lines in Figure 1). Both are supported by experiments, and both contribute to improving the solution frequency (Table 1), but E(SPL) autoinhibition contributes more than *cis*-inhibition of N by DL. We hypothesized that the reduced model has a higher solution frequency than the other two because it lacks activation of e(spl) by AC and SC (which implements a homeostat; see the Discussion). Indeed, eliminating AC and SC activation of e(spl) in the augmented model raised the frequency of solutions to almost the same as that of the reduced model (Table 1).

To see whether the network is robust to noise in initial conditions, we took solutions from each model and asked whether they could pass the 7-cell test from initial conditions as in Figure 2, but with Gaussian noise added to the initial concentrations of ac, AC, sc, and SC in each of the focal cells. For all three networks, we found solutions that were robust to this noise (Table 2). Similarly, both the augmented and reduced networks can select a single winner with very little initial bias in the prepattern by amplifying small amounts of noise (Table 2). Solutions for the reduced network were more likely to tolerate noise than solutions for the others. Since all networks can pass our tests, the results in this and the previous paragraph provide no evidence for which network is more likely to be an accurate representation

Table 1. Robustness of Neurogenic Network Models to Parameter Variation

Network/Task; (Number of Free Parameters)	Success Rate for Random Parameters, f; (f∧1/n)	Recombination Success
Augmented/2-cell; (69)	1/3,800; (0.89)	5%
Augmented/7-cell; (69)	1/8,000; (0.88)	5%
Standard/2-cell; (63)	1/113,000; (0.83)	1%
Standard/7-cell; (63)	1/153,000; (0.83)	not done
Reduced/2-cell; (53)	1/570; (0.89)	9%
Reduced/7-cell; (53)	1/2,900; (0.86)	4%
Augmented/2-cell; restricted range (69)	1/4; (0.98)	58%
Augmented/7-cell; restricted range (69)	1/34; (0.95)	45%
Variations on Augmented Network		
Augmented without DL *cis*-inhibition of N (67)	1/7,200; (0.88)	
Augmented without E(SPL) autoinhibition (67)	1/43,000; (0.86)	
Augmented without AC, SC activation of e(spl) (65)	1/600; (0.91)	
Augmented with no cooperativity in transcriptional activation/repression (54)	0/210,000	
Augmented with constant DL input in all cells (71)	1/157,000 (0.84)	

The second column shows the proportion of randomly selected parameter sets that successfully formed the test pattern. Success rates are based on finding hundreds (reduced and augmented networks) or dozens (standard network) of successful parameter sets. The n^{th} root of this success rate is in parentheses, where n is the number of free parameters in the model. Note that small differences in the n^{th} root lead to large differences in solution frequency because of the high number of parameters. The third column shows the percentage of successful offspring from 10,000 recombinations of 100–200 randomly found successful parameter sets (standard 2-cell used only 16 parental sets).

of what exists in the fly. But they do point to the relative importance of the different links in conferring robustness on the network and raise a puzzle about the role of *E(spl)* to which we return below.

To summarize, realistic dynamical models of the neurogenic network can, in fact, carry out classic lateral inhibition over an extraordinarily large portion of its parameter space, even when, unlike earlier models (e.g., [6]), we include the intermediate players between N activation and Dl expression. As we showed previously for the segment polarity network, these results predict that the lateral inhibition function of the core neurogenic network is very robust to changes in parameters; this network is a lateral inhibition module, whereas the segment polarity network is a boundary-stabilizing module.

Table 2. Robustness of Neurogenic Network Models to Variation in Initial Conditions, Prepattern

Network	Number of Solutions Found Robust to Noise
Augmented with prepattern	5/40
Standard with prepattern	1/10
Reduced with prepattern	12/40
Augmented without prepattern	1/40
Reduced without prepattern	4/40

The first three rows show the number of solutions to each network that formed the correct final pattern from at least 75% of noisy 7-cell prepatterns. We tested 100 prepatterns against 40 solutions per model (only 10 for the standard model). Each prepattern was as described in Figure 2, but the initial ac and sc product concentrations in the focal cells were randomly perturbed by a value drawn from a Gaussian distribution with a standard deviation equal to 40% of the initial difference between high and low cells. The last two rows show results for the same 40 solutions per model, but starting from initial concentrations drawn from a Gaussian distribution with the same mean for all cells and a 20% standard deviation (we used only 50 different prepatterns here). In both cases, we ensured that the middle cell had a higher initial concentration of each product than all surrounding cells (thus, there is still effectively a weak prepattern in the second case), but the difference could be arbitrarily small.

Parameters Are Not Restricted in Their Values

Figure 3 (also see the Supplementary Material) shows histograms of the values of selected parameters from all solutions we found while randomly searching parameter space using the 2-cell test on the augmented neurogenic network. Although some parameters tend to cluster in

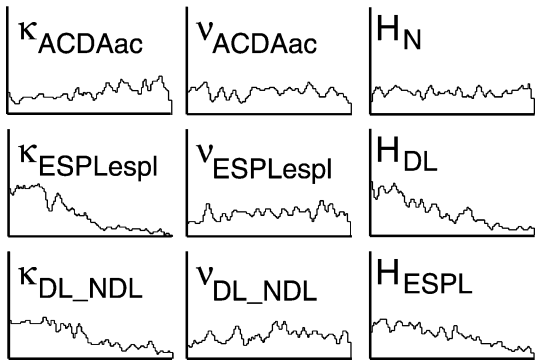

Figure 3. Selected Histograms of Successful Parameter Values from Solutions to the Augmented Neurogenic Network Using the 2-Cell Lateral Inhibition Test

Each histogram shows the values of just one of the 69 parameters in the model. Hs are half-lives, κs are half-maximal activations, and νs are cooperativity (Hill) coefficients. The horizontal axes span the range within which we sample each parameter (ranges: H = 1–1000 min; κ = 0.001–1 proportion of max concentration; ν = 1–10). We sample Hs and κs on a log scale. We selected these histograms to show the variety of clustering among different parameters (see the full suite in the Supplementary Material). We found the best range for each parameter by choosing values from the smallest subrange covering at least an order of magnitude and containing 50% of the values used in solutions. We also forced all transcriptional activation and repression in these sets to be at least slightly cooperative.

part of their range, there do not appear to be any absolute restrictions. That is, solutions permeate the entire huge box within which we randomly chose values (with one trivial exception of AC half-life, explained in the Supplementary Material). This apparent lack of restrictions could mask compensation between pairs or groups of parameters (e.g., parameter A could be high as long as parameter B was low and vice versa). We tested for but found no strong pair-wise cross-correlations, so any compensation must be relatively weak or, more likely, involve more than two parameters.

The Augmented Network Succeeds Almost Everywhere Inside at Least One Large Neighborhood in Parameter Space

Although only a single neurogenic network parameter showed absolute restrictions in its range, many parameters tended to cluster in a subset of the full range. Using the augmented network, we restricted each parameter to the best order-of-magnitude subrange with respect to the 2-cell solutions (see Figure 3 legend). We then conducted another random parameter search inside these new restricted ranges. The success rate increased almost 1,000-fold inside these (still very broad) ranges (Table 1). Using the same restricted range found with the 2-cell pattern, the success rate with the 7-cell test also increased more than 100-fold. We found a similar result with our model of the segment polarity network (G.v.D., unpublished data). Thus, both the neurogenic and segment polarity networks contain at least one large neighborhood of parameter space (10-fold-wide in all directions) in which solutions are very common. Furthermore, this neighborhood is shared between at least two of the patterns the neurogenic network makes.

The Neurogenic Network Can Act as a General Inhibitor that Outside Activators Must Overcome

One explanation of proneural cluster function invokes the classic Delta/Notch lateral inhibition switch. Experiments in grasshopper embryos [5], wherein ablation of the presumptive NB results in its replacement by another cell in the cluster, supports this explanation. A defining feature of this switch is that Dl is expressed at high concentrations only in proneural cells (as indeed occurs in wing discs [10]) and is downregulated in cells that lose *ac/sc* expression. In all the models above, DL is expressed only in the proneural cluster, consistent with the wing disc paradigm, and remains expressed forever only in the winner. In some solutions, however, DL expression attenuates very slowly, persisting in all cells of the cluster long after only the winner expresses AC/SC. This latter behavior is more consistent with experimental reports of Dl expression in embryos: Dl concentrations differ little among the NB, its neighbors in the cluster, and cells outside the cluster [11], and *Dl* mRNA concentrations begin to differ only after NB selection is complete [12]. These observations support a different view of neural specification in which neurogenic genes inhibit neural fate equally in all cells; prepatterning genes overcome this general inhibition only in the winning cell, with lateral inhibition playing no strong role (reviewed in [2]). Supporting this view, Seugnet et al. [13] replaced endogenous Dl with ubiquitously expressed Dl and found a fairly normal pattern of NBs in embryos, although 20% of clusters formed an extra NB.

Seugnet et al.'s experiments are so compelling that we think it useful to explore the conflict between the two apparently opposite views just outlined. Many of the solutions in our 7-cell test exhibited similar levels of dl (mRNA) and DL (protein) in all seven cells over most of their run, but all had at least some period of time in which dl concentration was at least twice as high, and often much higher, in the presumptive neural cell than its neighbors. If lateral inhibition and proneural regulation of *Dl* are indeed involved in neural specification, our models indicate that, for at least some brief period of time, Dl or some related protein (such as different isoforms of N [14]) must have higher concentrations in the winner. Dl dynamics are complex [12] and difficult to measure accurately due to the concentration of both *Dl* protein and mRNA in membranes ([11, 12]; E.M. and D. Lehman, unpublished data), so it is still unclear whether there are in fact heterogeneous concentrations within proneural clusters. But if differences in *Dl* levels are important to the mechanism, how can the process still succeed in most cases despite constitutive *Dl* expression [13]?

We added a constitutive input (with tunable parameters determining that input rate) to the dl promoter in our augmented model to produce the broad distribution of DL observed in early embryos. We found solutions in which this model still passed our 7-cell test (Table 1), in one of which the final DL concentrations differed by less than 40% between the winner and all other cells. When we removed the upregulation of dl by AC and SC in one of these solutions, leaving the constitutive expression, the network passed the test for lower initial AC concentrations, but all cells became neuronal with higher initial AC concentrations. This suggests a resolution of conflicting experimental results. In cases in which the prepattern is well tuned, the core neurogenic network, without lateral inhibition, can select a single winner through cell autonomous processes. But if there are small changes in the prepattern, or developmental noise, lateral inhibition enables the network to continue to function where it would otherwise fail.

The Model Neurogenic Network with a Single Set of Parameter Values Can Form Several Different Patterns

We might expect that the neurogenic network, or some ancestral version, evolved originally to perform just one patterning behavior. Over time, the evolutionary process coopted it to perform novel functions. This path would seem easier to travel if the network with a single set of parameter values could pass each of our tests simply by starting from different prepatterns (initial conditions). We found this to be true for the augmented network. Many of the 2-cell solutions also enable the network to pass the 7-cell or line tests; a large fraction could form all three patterns (Table 3). Many 7-cell solutions similarly enabled 2-cell, line, or all three patterns. This reveals a fundamental relationship between robustness and evolvability (see Figure 4 legend), as anticipated by Kirschner and Gerhart [15].

Table 3. Proportion of Solutions to One Test that Can Pass Other Tests in the Augmented Network

Original Pattern	2-Cell	7-Cell	Line	All Three
2-Cell	1.0	0.25	0.14	0.12
7-Cell	0.8	1.0	0.33	0.25

The Shape of the "Working Region" of Parameter Space

Clearly, simple but realistic representations of the core neurogenic network can accomplish classic lateral inhibition. Given the large number of parameters in each model, random sampling within a huge box in parameter space produces solutions with extraordinarily high frequency, indicating that the neurogenic network is robust to changes in parameters. From a biological perspective, however, this statistical measure is a crude characterization of parameter space. We could get such results from a single broad basin of good parameters, a skinny crevasse that snakes around parameter space, many small, disconnected neighborhoods, or anything in between. These differences would have important biological consequences, both for robustness of the network as it exists today and for the ability of evolution to navigate parameter space. We explored these shapes in three ways. We first used a traditional sensitivity analysis to show that most parameters can vary by over an order of magnitude around the typical solution (see the Supplementary Material). We now propose two new measures to quantify genetic network robustness in more biologically meaningful terms and reveal the shape of the functional territory.

In nature, different solutions are blended through recombination during sexual reproduction. We tested robustness by recombining randomly found solutions to create offspring parameter sets. Our "recombination" means randomly selecting each offspring parameter value from one or the other of two randomly chosen successful parental parameter sets (Figure S3). Success of offspring sets at passing our tests generally correlates with the frequencies of finding solutions through a random search of parameter space (Table 1). As with random sampling, the recombination success rate was much higher among solutions found within the restricted parameter ranges. These results confirm the robustness of the networks and hint that the restricted parameter ranges contain a region in which almost all recombinations would be successful.

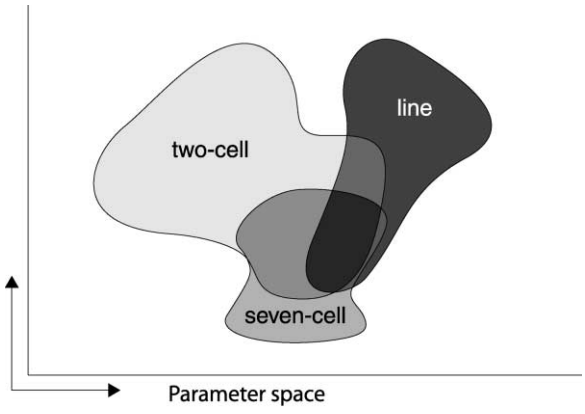

Figure 4. Robustness to Parameter Variation Confers Evolutionary Flexibility in the Pattern Produced by the Network

The cartoon shows a hypothetical cross-section through the parameter space with outlines for regions where the network can form different patterns. The regions overlap approximately as shown. If the neurogenic network was originally tuned to make any one of the three patterns, evolution could modify parameter values, wandering through neighborhoods of working territory, until the network could make the other patterns as well. This is only possible because of the network's robustness to parameter variation; if the pools were too small, they would not overlap, and natural selection would have to leap from one function to another across nonfunctional zones.

To test that latter prediction, we devised another measure of robustness, which we term "mutational expansion", by adding an evolutionary component on top of recombination. We imagine that a single founder solution set of parameters expands by simulated mutations into a population of individuals who all recombine with each other randomly to produce successive generations. As more generations go by, more mutations arise, and the spread of values for each parameter in the population will grow larger. We want to find out how large a spread in parameter values the population can tolerate before the within-population recombination success rate drops too low (below 90% mating success) (see the Supplementary Material).

We tried this procedure on approximately 25 randomly selected "founder" parameter sets found within the full parameter ranges ("widely selected sets") and 25 found within the restricted parameter ranges ("narrowly selected sets") (cartooned in Figure 5A). After completing the mutational expansion procedure on each, we measured the average ratio between the highest and lowest value for each parameter within the population. This measure will be highest when the volume in parameter space covered by the population is wide in all dimensions (along all parameter axes). Using this measure, we found that parameter sets from the full parameter ranges generally expanded into populations that spanned much smaller volumes than those found in the restricted range (Figures 5B and 5C; cartooned in Figure 5D). The largest volumes spanned over an order of magnitude on average for each parameter, a truly enormous range considering the difficulty of the test.

How do these populations (parameter space volumes) relate to each other? To answer this, we recombined parameter sets from each population with those from the other populations. The mating success between narrowly selected populations often approached the mating success within a single population (Figure 5D caption). The mating success between a narrowly selected population and the widely selected populations was quite a bit lower, and the success among widely selected populations was even lower, though still higher than the success rate for completely random parameter selections. This leads us to a new view of the functional territory for this network in parameter space; it has the shape of a high-dimensional octopus (Figure 5E), with at least one large central area (the body) in which solutions are very common and long tentacles of good solutions snaking away from this area within which solutions

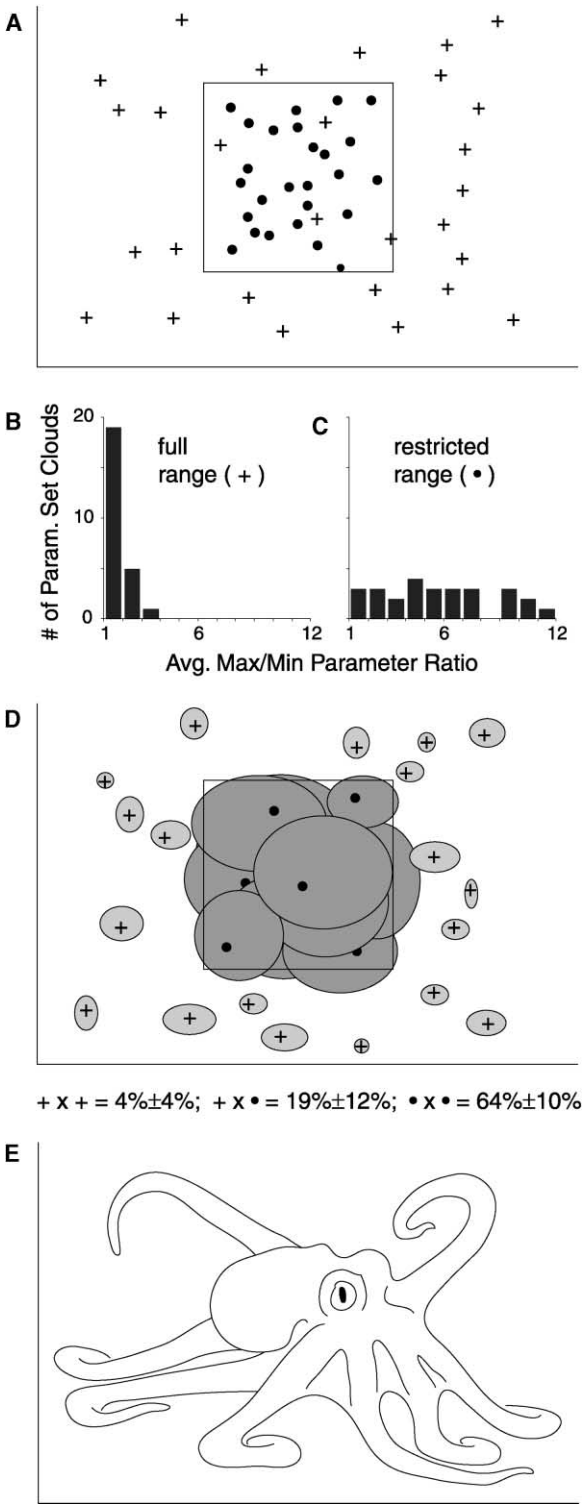

Figure 5. The Parameter Space Is Shaped Like an Octopus with Multiple Possible Network Functions Overlapping

(A–D) A cartoon of randomly selected solutions from the full (plus) and restricted (bullet) parameter regions. After mutational expansion, each of these grows into a population of parameter sets confined within a local solution volume (shaded domains in [D]). The ratio of max/min values for each parameter across all the sets in a population have average values shown in (B) and (C). These ratios indicate that most parameter sets from the restricted range occupy larger solution spaces than those from the full range (cartooned in [D]). The text under (D) shows the average mating success between parameter sets found within the full range (plus × plus), within the restricted range (bullet × bullet), and between one parameter set from the restricted range with parameter sets found in the full range (plus × bullet). This leads us to view the solution space as resembling an octopus, with the solutions found in the full parameter space out in the arms and a large central area around the restricted range.

are also common, but these are separated by regions in which solutions are rare or nonexistent.

Discussion

The Neurogenic Network's Robustness Leads to Evolutionary Flexibility

Most efforts so far to explore realistic models of gene networks show robustness to parameter variation, e.g., [1, 16–18]. The core neurogenic network also exhibits enormous robustness to variation in parameters as well as robustness to noise in initial conditions. Furthermore, the neurogenic network model reveals a fundamental relationship between robustness and evolvability. We tested the model's ability to perform three distinct functions. Some parameter sets enable one or another function; however, the existence of parameter sets enabling all three functions means that once the network is tuned to do one of those tasks, say, the 2-cell choice task, the parameters are free to vary, only *because* of the network's robustness, into the subspace in which the model could also make the 7-cell and line patterns. This would not be the case if the network was not robust for all three functions; thus, in this case, robustness confers flexibility (Figure 4).

We don't claim that this model accounts for every context in which these genes are deployed by developing embryos. In particular, we don't believe the model can account for more complex pattern-formation tasks, such as the production of regularly spaced bristles, in which mutations in components of this network change the spacing. To produce robust longer-range patterns, we think this lateral inhibition model must be joined with either a mechanism for long-range cell-cell communication (such as Dl diffusion [19, 20] or, as suggested by an anonymous reviewer, long-range filopodial contacts) and/or a mechanism by which N/Dl signaling could feed back on a spatially varying prepattern. The model we present here seeks to reconstitute the canonical lateral inhibition phenomenon from the known facts of the coupling between N/Dl signaling and the proneural genes; something is still missing from our picture, since the known facts seem to us inadequate to explain larger-scale patterning processes.

Constitutive Dl Production Combined with Weak Lateral Inhibition Can Resolve the Controversy about the Role of Lateral Inhibition in Neurogenesis

As discussed above, the experimental literature includes both support for, and refutation of, an important role for lateral inhibition in neural determination. Our results can account for both sets of experiments. If the prepattern

that initiates neuroblast selection is well tuned, the prepattern plus a constant level of inhibition could select the winner, absent lateral inhibition. But lateral inhibition buffers the patterning against perturbations in the initial prepatterning (e.g., due to genetic or environmental variation, or "developmental noise"). Seugnet and colleagues [13] reported that, with only constant production of Dl, 80% of proneural clusters developed normally, but 20% produced an extra NB. We interpret these experiments to say that the prepattern is well tuned in most proneural clusters, but in 20%, either a poorly tuned prepattern or noise causes errors in the absence of lateral inhibition. This is a testable idea. One could remove lateral inhibition as Seugnet et al. [13] did. We would then predict that the embryo would be much more sensitive to hyper- and hypomorphs in prepatterning genes such as *extramachrochaete* and *hairy*. We would also predict that they would be more sensitive to mutations in genes within the network itself, such as missing or extra copies of *Dl* or *N*. We make the latter prediction because those mutations should change the threshold to which the prepattern is tuned. In the absence of lateral inhibition, a prepattern that was well tuned to the former threshold could not also be well tuned to the new threshold.

The *E(spl)* Homeostat Might Increase the Network's Robustness to Noise, as in Human-Engineered Circuits

From our results above and others not described here, we deduce that *E(spl)* greatly reduces the percentage of random parameter sets that enable lateral inhibition. We believe this is because *E(spl)* acts as a homeostat. As the expression level of the proneural genes (*ac* and *sc*) rise, their products activate *E(spl)*. E(spl) then downregulates the proneural genes. As with the thermostat in a house, this negative feedback loop tends to keep the proneural genes at an intermediate level rather than allowing them to switch to either a high or low state. Both E(spl) autoinhibition, and to a lesser extent *cis*-Dl inhibition of N activation, help overcome this homeostat. On the face of it, this seems a strange design. The *ac/sc* network itself is a bistable switch that tends to go in the direction it is pushed and remain there. The switch and homeostat mechanisms are exact opposites. We found that removing the homeostat (the reduced model) makes it easier to find parameter sets that pass our tests (which all involve throwing the switch). Why incorporate counteracting mechanisms in the same circuit?

It is, of course, possible that this is simply a vestige of the network's evolutionary history, with no design rationale. But electrical engineering suggests one possible advantage. An op-amp is a famous circuit that amplifies the difference between two inputs. Good op-amps can amplify a voltage difference more than a millionfold. Usually, though, engineers add a negative feedback circuit (that is, a homeostat). This greatly attenuates the gain but makes the amplifier much more stable; noise generated internally inside the op-amp will not affect the output signal. Reducing function to gain stability is common in other electrical circuits as well. These electrical circuits do not make good direct analogies to genetic networks, but the concept of adding negative feedback to increase stability might still apply. Perhaps the *E(spl)* homeostat reduces the network's sensitivity to developmental noise such as stochastic changes in transcription or translation rates, in the prepattern, or in the concentrations of modulators such as Da and Emc.

A related design benefit might be that the *E(spl)* homeostat prevents the network from switching individual cells on or off before the prepattern has a chance to decree the winner. A simple bistable switch consisting of *ac* and *sc* alone could not help but be thrown in one direction or the other by noise (as apparently takes place in *C. elegans* anchor cell specification). Adding *E(spl)* leads to a new, neither-on-nor-off steady state, which could enable the proneural switch to procrastinate until some extrinsic cue forces the system to choose one or the other switched state.

Conclusions

Our results highlight several evolutionarily interesting properties of genetic networks. Like several other recent simulated networks, this model network, inspired by an evolved design, exhibits robustness that would be the envy of any human designer. Perhaps this is a generic feature of genetic organization, but perhaps it reflects a coevolution between evolved networks, biologists, and theorists: modular, robust networks are the easiest to get at experimentally. Thus, they are the best understood and are the best fodder for models. In any case, the model discussed here shows that this robustness can lead to an evolutionary flexibility by which a network originally tuned to one function can mutate, without losing that original function, toward the ability to perform additional functions. Furthermore, were evolution to discover a relatively nonrobust solution (at the tip of an arm in Figure 5E), the network could migrate down the octopus arm to more robust regions. Conversely, the network could migrate back up one of these arms if the characteristics of the central region were no longer selected for. This could lead to reproductive isolation, the precondition for speciation.

Supplementary Material

In the Supplementary Material, we explain how we derived the model network in Figure 1 from relevant experimental facts about the real neurogenic/proneural network. Supplementary Material including the mathematical formulation of the model, further explanation of the algorithms we used to explore parameter space, a comparison of the neurogenic network with our segment polarity model using these methods, and a complete presentation of the results exemplified in Figure 3 is available at http://images.cellpress.com/supmat/supmatin.htm.

Acknowledgments

Thanks to Xantha Karp, Ed Giniger, Jaideep Mavoori, James Defuria, Jose F. de Celis, and Marc Haenlin for informative discussions. Funding for this work came from National Science Foundation (NSF) grants MCB-9732702, MCB-9817081, and MCB-0090835, and we would especially like to remember Delill Nasser at the NSF for encouraging us in early exploratory stages of our work. E.M. received support from an HHMI predoctoral fellowship.

Received: January 16, 2002
Revised: March 5, 2002
Accepted: March 7, 2002
Published: May 14, 2002

References

1. von Dassow, G., Meir, E., Munro, E.M., and Odell, G.M. (2000). The segment polarity network is a robust developmental module. Nature *406*, 188–192.
2. Simpson, P. (1997). Notch signalling in development: on equivalence groups and asymmetric developmental potential. Curr. Opin. Genet. Dev. *7*, 537–542.
3. Perron, M., and Harris, W.A. (2000). Determination of vertebrate retinal progenitor cell fate by the Notch pathway and basic helix-loop-helix transcription factors. Cell. Mol. Life Sci. *57*, 215–223.
4. Wilkinson, H.A., Fitzgerald, K., and Greenwald, I. (1994). Reciprocal changes in expression of the receptor lin-12 and its ligand lag-2 prior to commitment in a *C. elegans* cell fate decision. Cell *79*, 1187–1198.
5. Doe, C.Q., and Goodman, C.S. (1985). Early events in insect neurogenesis. II. The role of cell interactions and cell lineage in the determination of neuronal precursor cells. Dev. Biol. *111*, 206–219.
6. Collier, J.R., Monk, N.A., Maini, P.K., and Lewis, J.H. (1996). Pattern formation by lateral inhibition with feedback: a mathematical model of delta-notch intercellular signalling. J. Theor. Biol. *183*, 429–446.
7. Skeath, J.B., and Carroll, S.B. (1992). Regulation of proneural gene expression and cell fate during neuroblast segregation in the *Drosophila* embryo. Development *114*, 939–946.
8. de Celis, J.F., Garcia-Bellido, A., and Bray, S.J. (1996). Activation and function of Notch at the dorsal-ventral boundary of the wing imaginal disc. Development *122*, 359–369.
9. Meir, E., von Dassow, G., Munro, E., and Odell, G.M. (2002). Ingeneue: a versatile tool for reconstituting genetic networks in silico. J. Exp. Zool., in press.
10. Haenlin, M., Kunisch, M., Kramatschek, B., and Campos-Ortega, J.A. (1994). Genomic regions regulating early embryonic expression of the *Drosophila* neurogenic gene Delta. Mech. Dev. *47*, 99–110.
11. Kooh, P.J., Fehon, R.G., and Muskavitch, M.A. (1993). Implications of dynamic patterns of Delta and Notch expression for cellular interactions during *Drosophila* development. Development *117*, 493–507.
12. Haenlin, M., Kramatschek, B., and Campos-Ortega, J.A. (1990). The pattern of transcription of the neurogenic gene Delta of *Drosophila melanogaster*. Development *110*, 905–914.
13. Seugnet, L., Simpson, P., and Haenlin, M. (1997). Transcriptional regulation of Notch and Delta: requirement for neuroblast segregation in *Drosophila*. Development *124*, 2015–2025.
14. Wesley, C.S., and Saez, L. (2000). Analysis of notch lacking the carboxyl terminus identified in *Drosophila* embryos. J. Cell Biol. *149*, 683–696.
15. Kirschner, M., and Gerhart, J. (1998). Evolvability. Proc. Natl. Acad. Sci. USA *95*, 8420–8427.
16. Tyson, J.J., Novak, B., Odell, G.M., Chen, K., and Thron, C.D. (1996). Chemical kinetic theory: understanding cell-cycle regulation. Trends Biochem. Sci. *21*, 89–96.
17. Barkai, N., and Leibler, S. (1997). Robustness in simple biochemical networks. Nature *387*, 913–917.
18. Laub, M.T., and Loomis, W.F. (1998). A molecular network that produces spontaneous oscillations in excitable cells of *Dictyostelium*. Mol. Biol. Cell *9*, 3521–3532.
19. Klueg, K.M., Parody, T.R., and Muskavitch, M.A. (1998). Complex proteolytic processing acts on Delta, a transmembrane ligand for Notch, during *Drosophila* development. Mol. Biol. Cell *9*, 1709–1723.
20. Qi, H., Rand, M.D., Wu, X., Sestan, N., Wang, W., Rakic, P., Xu, T., and Artavanis-Tsakonas, S. (1999). Processing of the notch ligand delta by the metalloprotease Kuzbanian. Science *283*, 91–94.