Life in the Immaterial World

Matthew Caldwell

CoMPLEX, University College London

Supervisors: Buzz Baum, Mark Miodownik

Body: 4400 Words

Appendix: 640 Words

January 26, 2007

Abstract

Multicellular organisms acquire their structure and behaviour incrementally in the course of development from a single cell. This process is enormously intricate and may give rise to attributes that confound interpretation. Greatly simplified cellular automaton models can provide useful insights, especially when considered in the light of evolution. We present one such model, *EmbryoCA*, and consider some of its implications.

Contents

1 Development & Evolvability

1.1 From Egg to Organism

All multicellular organisms develop by division from a single cell.[1] To a good approximation,[2] that cell's genome, along with most of its structure, is faithfully reproduced in all its myriad descendants. Yet those descendants can vary greatly in their form and function. Much of the process by which this occurs remains mysterious. [8, 7]

As the egg divides and divides again, its initial homogeneity breaks down, and the cell mass undergoes *patterning*: a separation into different functional regions. Within those regions, the (still genetically identical) cells begin to take on more specialised functions (*differentiation*). The overall organism starts to develop distinctive features of shape (*morphogenesis*), and exhibits *growth*, becoming larger and more complete.[18, 14]

These are not one-off processes, but continue in a sequence of increasing localisation throughout the development, and to some extent the lifetime, of the organism. Moreover, they are both profoundly *contingent* (tightly bound up with their immediate chemical and physical environment) and remarkably *robust*. We will return to these characteristics repeatedly in the discussion below.

The behaviour of all cells, along with their sub-cellular constituents, is subject to the constraints of material reality. Those constraints are quite different at different scales: a cell's chemical behaviour will be governed by electromagnetic forces, its physical construction by surface tension and viscosity, the gross structure of the organism of which it is a part will most likely be strongly affected by gravity and so on.[11] Thus, development entails an overall *accretion* of constraints, the interaction of which contributes substantially to the final form of the organism.

Individual cells experience a similar increase in behavioural restriction over the course of development. The original fertilised egg has the potential, and indeed the *need*, to divide into every kind of cell that the organism will ever possess. This plasticity diminishes as the organism becomes more complete, until most mature cells have no choices left at all.

1.2 Robustness & Homeostasis

In many organisms, though by no means all, the initial embryonic environment is carefully controlled; nevertheless, there will always be some degree of variability. The world is an uncertain and dynamic place and many external influences—including experimental scientists—may disrupt the environment in which development occurs.

Development is intimately dependent on intricate interactions between the organism and its surroundings, and one might expect such a process to be highly sensitive to environmental perturbation. In fact, it turns out to be remarkably robust, not only to biologically-plausible levels of noise but to considerably more drastic experimental interventions.

The well-documented phenomenon of embryonic *wound healing* is one such example: in marked contrast to adult organisms, developing embryos can quickly and seamlessly repair inflicted damage involving significant tissue loss.[13] Viewed as a distinct phenotypic character, this ability is of doubtful evolutionary value: it is unlikely that an embryo could suffer such a wound

---

[1]Various plants and animals can reproduce by budding or by regeneration from a piece of an adult. Somewhere along the line there must nevertheless have been an original single cell from which the divided adult grew.

[2]It is reasonable in this context to ignore somatic mutations and, to a lesser extent, generational changes such as truncation of telomeres.

outside the laboratory without irrecoverable damage to its environment and consequent death. Any selective pressure for perfect wound healing *per se* would be much more likely to favour it in the adult.

However, evolution operates in the space of the possible; and wound healing may turn out to be just one facet—perhaps a quite incidental one—of the more general feature of developmental robustness.

If the goal of development is reliably to produce a phenotype—to *re*produce, with minor variations, the phenotype of the parents—then robustness is clearly advantageous, since it means the required form will be reached irrespective of (some tolerable range of) environmental disturbances. Moreover, once the form has been reached, it must be *maintained*—it is unlikely that the organism will be able to fulfil its reproductive purpose at the very instant it achieves its proper shape—and most organisms do indeed exhibit such *homeostasis*.

Both of these criteria can be satisfied if the ultimate form constitutes, in the language of dynamical systems, an *attractor*.[3] The corresponding basin of attraction would define the range of perturbations to which the development was robust.

Clearly, any conjectural *phase space* of embryonic development would likely be so complex as to be barely imaginable, let alone amenable to analysis. Were we to try to *engineer* a developmental process producing a viable organism with attractor-like attributes in such a space, we would not even know where to begin.

That, however, isn't the problem that evolution faces.

1.3   Nothing Succeeds Like Success

Like development, natural selection is highly contingent: the reproductive success of a species—its only measure of evolutionary merit—is closely dependent on the details of its environment. Evolution creates—or *finds*—things that reproduce well *in context*. While immediate phenotypic fitness is an ingredient in this success, it isn't the whole story. The reproductive processes themselves are also subject to evolution.

Successful reproduction implies some degree of generational similarity: it would be meaningless to describe an organism as 'reproducing' if all its offspring were wholly unlike their parent and each other. Moreover, in such a case, any reproductive advantages of the organism would most likely be lost in its randomly-differing children.

To properly occupy a region of the fitness landscape, a species must continue to reproduce with reasonable fidelity over many generations. For multicellular organisms, this suggests also that they must *develop* with reasonable fidelity. If the environment is noisy, this may constitute a selective pressure in favour of developmental robustness—at least if developmental variations lead to reduced fitness.

However, the relationship between robustness, evolution and environment is fraught with complications. [9] With too little variation, evolution has little to work with; with too much, advantages may be too easily lost.

For ideal evolvability, an organism must be able to support both incremental variation—exploring possible local improvements in phenotype—and also the background accumulation of *locally neutral* changes—allowing the investigation of longer-range possibilities, and perhaps

---

[3]The use of the term here is extremely informal, but I hope it will be of at least metaphorical value later on.

escape from a local maximum, without sacrificing current fitness.[4] [15, 16, 6]These features typically carry a cost of their own, and so will also be subject to selective pressure.

There is thus a balance to be struck between robustness, evolvability and efficiency, and where that lies will depend both on the phenotypic space and the genotypic: how the organism interacts with its environment, but also what is evolutionarily *possible* given its current genome. Robustness may be one result; but so may pared-down fragility in the right circumstances.

2 Computational & Modelling Background

2.1 Cellular Automata

A *cellular automaton* (CA) is a discrete model in which *cells* arranged in a regular array each have an associated *state* drawn from a finite set, and an associated *neighbourhood* of connected (usually adjacent) cells. Some initial condition defines the start states of all cells. At each subsequent time step, the state of each cell is calculated from the previous state of the cell and its neighbourhood by application of some (usually simple and deterministic) rules. In most CAs, all cells are updated in parallel.

There are many variations on this basic scheme (the EmbryoCA model described below is one such), but the characteristics of finite states, local interaction and highly-parallel processing are usually maintained.

Cellular automata can be considered as discrete approximations of systems of partial differential and/or integral equations describing the interdependencies of many variables. As such, they often exhibit highly complex emergent behaviour that is at least qualitatively, and in some cases also quantitatively, comparable to continuous models of real-world dynamic systems. At the same time, they are computationally tractable and can usually be represented graphically in an intuitive way.[5] [5]

CAs have been applied to a wide variety of biological problems, particularly those around pattern formation. Indeed, some specific instances of natural growth patterns are known to be produced by processes that closely approximate CAs.[17]

2.2 Genetic Algorithms

Genetic algorithms (GAs) are one of a number of different *evolutionary computation* approaches attempting to exploit the power of evolution *in silico* for engineering or exploratory purposes. In the former case, the goal may be to solve specific problems: for example, tuning the weights of a neural network to perform some desired calculation. In the latter, the goal is to investigate GA behaviour itself, often in the hope of adding to our understanding of the workings of evolution in the living world.

GAs are distinguished from other evolutionary computation methods by usually including some explicit separation between genotype and phenotype—with variation acting on the former and

---

[4]This latter capability is strongly implicated in the natural evolution of novel characters, and usually implies some redundancy in the genotype.

[5]On the other hand, like many dynamic systems, their behaviour over time may be counter-intuitive and resistant to *ex ante* analysis. Indeed, several simple CAs have been found to be Turing-equivalent: being able to predict the outcome of any arbitrary input to such a CA would be tantamount to solving the Halting Problem. Of course, any number of *particular* initial conditions can still produce readily predictable behaviour.

selection on the latter—and by their use of the *crossover* or *recombination* operator to generate child genomes from pairs of parents. [12]

A GA generally starts from a random population of genomes. These are assessed by a *fitness function* to estimate how well each corresponding phenotype solves the target problem. Parents are selected according to fitness and combined by crossover and mutation to produce a new population. The process is repeated for many generations and will often—although by no means always—discover highly fit individuals.

GAs are usually characterised as a type of search algorithm, conducting a heuristic parallel search of a high-order possibility space by means of *schemata*—combinatorial subsets of proposed solutions.

GAs, along with most other current evolutionary computation approaches, have been criticised on several grounds, in particular their simplistic approach to the genotype/phenotype relationship and the fixed forms of their genomes and fitness functions [1]. However, the proposed alternatives—which incorporate, *inter alia*, ideas of development and co-evolution of the reproductive process—have yet to bear significant fruit and risk losing explanatory power in their welter of complexity. There is much to be said for *simplified* models.

3   EmbryoCA

3.1   Overview & Aims

The local interactions and parallel processing of cellular automata make them appealing as models for embryonic development. A CA can be started from a fixed initial state—typically a single cell—and developed by repeated application of fixed rules (the genome) until it reaches some final phenotype.

However, CA behaviour is typically very *brittle*, depending sensitively on the *letter* of its rules. Such CAs do not lend themselves to evolution because genetic variation tends to result either in no change at all or else a catastrophic alteration of the phenotype. Without middle ground there is little scope for exploration.[10, 3]

EmbryoCA is an *effector automaton*[6] developed by David Basanta and colleagues with the express purpose of providing decent evolvability. Attempts to do this generally demand a significant increase in rule complexity to allow for redundancy and incremental change. In EmbryoCA, such *complexification* is achieved by distributing responsibility across a *population* of rules. The rules themselves—the conditions they assess and the resulting actions—remain very simple, but each contributes only a single *vote* towards the action taken by the cell. Except in the most degenerate cases, no individual rule will be able to dominate the CA behaviour and thereby present an obstacle to evolution.[4]

Rule application and the electoral process remain deterministic and uniform—no attempt is made to model gene regulation or cell differentiation, nor to model environmental perturbation. The system is equivalent to having a single fixed rule with many conditions and possible results, but is both easier to understand and more readily mapped to a genome for the purposes of a genetic algorithm: one rule, one gene.

---

[6]EAs are similar to CAs but treat the cell lattice as a *geography* within which the smaller number of stateful and rule-driven automata operate. Wolfram [17] calls these *generalised mobile automata* and demonstrates their functional equivalence to fixed CAs, so the distinction should be thought of as a matter of optimisation: for suitable problems, the EA will be more efficient because fewer cells are considered at each time step.

Having first established that this model is indeed evolvable, the researchers then applied a GA to evolve solutions to a simple, highly abstracted, developmental task, that of generating an organism with distinct developmental phases: a period of vigourous growth, followed by a longer period of homeostasis. (See Appendix A.)

This is an extremely general problem statement, making almost no assumptions as to the mechanisms that might be used or the purposes they might serve. Further, the CA environment includes no attempt to represent any of the material details—forces, thermodynamics, chemical properties—on which real world development so depends.[2]

Nevertheless, the task captures something essential for development and its very abstraction makes it potentially illuminating.

3.2   Initial Results

The model successfully evolved a number of organisms exhibiting a good level of homeostasis. Despite the simplified environment, the organisms hit on different strategies for doing so: *static* strategies, in which the cells reach a final form and subsequently maintain it by inaction; and *dynamic* strategies, in which form is maintained by an ongoing balance between cell division and death. Between these two poles were organisms exhibiting a combined strategy, with both static and dynamic regions.

These strategies are comparable to those found in living organisms.

A particularly interesting analogy is between the stratified structure of evolved organism #18 and many animal tissues such as those of the skin and gut (see Figure 1). In both cases, a base layer of ever-dividing *stem cells* replenishes a region of transit amplifying cells that migrate towards a surface region in which they die. That this strategy should occur in such an abstracted simulation as well as so widely in living organisms suggests it may be a natural—which is to say, *attracting*—solution to the problem of homeostasis.

3.3   Subsequent Experimentation

Having evolved such organisms, a natural extension was to investigate their behaviour under systematic perturbations. Two principal tests were performed:

- To identify the single genes with the greatest impact on homeostasis, a systematic loss of function analysis was applied, generating phenotypes from variant rulesets in which each
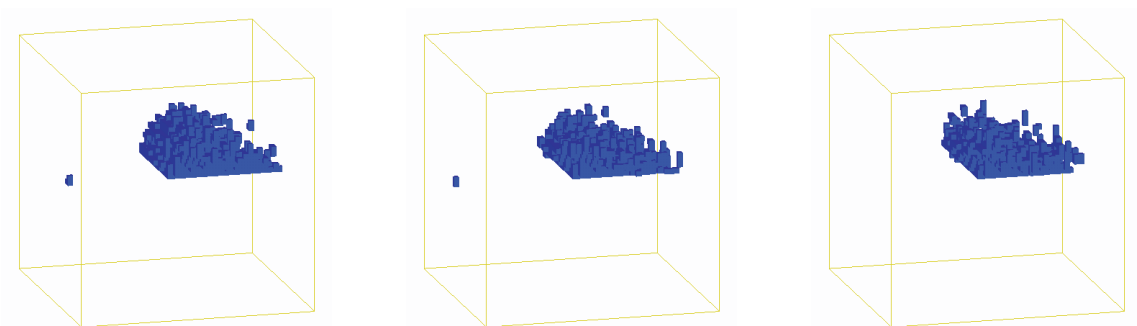

Figure 1: Organism #18 at timesteps 50, 100 and 150. Cells propagate up from the flat base to die at the surface.

of the 100 rules was deleted in turn and comparing those to the unmodified original.

- To test robustness to environmental perturbation, the original organisms were subjected to 'gunshot wounds'—the instantaneous removal of a body of cells at timestep 100—and their form at step 150 compared to that in the unwounded version.

In the first test, a small number of rules were found to play a disproportionately large role in homeostasis. In particular, for organisms whose homeostatic strategy involved a dynamic component, the critical rules were all found to be promoters of *apoptosis*. Deletion of these genes led to loss of growth regulation and consequent disruption of form. While it would be reckless to take the analogy too literally, there is a clear correspondence here with cancerous growth in living tissue: controlled cell death is evidently an important component of health in artificial systems, just as it is in living ones.

All the homeostatic organisms also proved remarkably robust to wounding, despite that being an event to which they had never been exposed during their evolutionary history and thus for which there would never have been an explicit selective pressure. While such a result might be expected in organisms with a dynamic strategy—which would, in any case, be constantly renewing their form—it is more surprising in the static organisms, in which the circumstances for such self-repair would not normally occur.

In the specific case of the tissue-like organism #18, the wounding experiment showed evidence of something interpretable as functional differentiation between the cells in separate strata: while a wound to the transit and outer layers was quickly repaired, a substantive wound to the generative 'stem cell' layer at the base led to progressive tissue loss and even the death of the entire organism.

A subsequent loss of function investigation of the wound healing response in this organism revealed that the rule with greatest impact on wound healing was one that was *not* crucial for homeostasis. Examining the evolutionary history of the organism, however, showed that in early antecedents *not* exhibiting significant wound healing, the same rule *did* contribute somewhat to homeostasis. It appears that, as more efficient mechanisms for the primary purpose of homeostasis evolved, the redundant gene acquired this inadvertent secondary role.

3.4   Conclusions

At least in this highly abstract universe, there is clearly some relationship between the primary traits for which an explicit selective pressure exists and secondary behaviour that, although not selected for, seems consistently to co-evolve. Some degree of robustness may turn out to be an inevitable corollary of the evolution of homeostasis, which itself is a common feature of most living organisms.

The fact that artificial organisms evolve similar strategies in the solution of this simple problem to those found in the natural world is *at best* circumstantial, but it is nonetheless *suggestive* that there may be patterns in the evolutionary landscape toward which evolution tends to gravitate.

It is possible that the apparent correlation between embryonic wound healing and more general developmental robustness represents one such attractor. But even if that turns out not to be the case, the results from this model should at least sound a note of caution about too readily interpreting particular behaviours as evolutionarily selected in and of themselves rather than just coming along for the ride.

4 Further Investigation

4.1 Speculation

The wounding experiment demonstrates that these evolved EmbryoCA organisms exhibit some robustness to acute environmental perturbation, even though such perturbation was never explicitly part of the selective landscape in which they evolved. Their developmental environment during evolution, and thus the mapping from genome to fitness, was entirely deterministic and error-free.

Nevertheless, from the discussion in §§1.2-1.3 above, some degree of robustness is not a wholly unexpected companion to success. The random variations by which the genome evolves constitute a kind of *noise*, and small genomic changes may manifest unpredictably during development. We might expect that selection would favour organisms capable of holding on to their advantages despite such minor mutations, and which include sufficient redundancy to allow exploration of more distant evolutionary possibilities.

If this *were* the case, the successful organisms ought also to exhibit some level of robustness to imperfections in *development*: faults individually less extreme than the gunshot wounds—which incur a major recovery cost—but which, unlike those, are *chronic*.

4.2 Introducing Stochastic Effects

In order to test this hypothesis, the operation of the CA has been altered to allow a parameterised level of stochastic imperfection when applying the ruleset, analogous to errors in transcription or translation. Since there is no single obvious definition of 'imperfection' in this context, two different versions are implemented. In the less specific variant, there is a probability that any rule might be replaced by a randomly concocted other—on a one-off basis—before being tested and invoked. In the more directed, the probability is for the existing rule to be applied *in reverse*: promoting the action it would normally inhibit, or vice versa.

In both cases, we expect a degradation of phenotypic behaviour as the error probability increases. However, the exact nature of the degradation will depend on the presence and nature of developmental robustness. If development *is* robust, we would expect the effect of noise to be relatively low up to some threshold; if not, then even small chronic perturbations should be enough to derail the final phenotype.

With the introduction of stochasticity, it is no longer possible to get a single 'correct' answer from the model: different runs produce different results. Instead, a statistical approach must be used, in which the results of many runs are taken together, and their *distribution* assessed.

The fitness functions applied in the evolutionary stages are a natural choice for quantifying the effects of developmental noise, at least in the first instance. However, they are unlikely to tell the whole story.

4.3 Results

Summary results from a series of trials on the tissue-esque organism #18 are shown in Figures 2 and 3. While far from conclusive, they exhibit several interesting features. Perhaps surprisingly, there is a qualitative difference between the effects of the two alternative noise implementations as their level increases.

Effect of Noise on Fitness

Fitness Distributions with Inversion
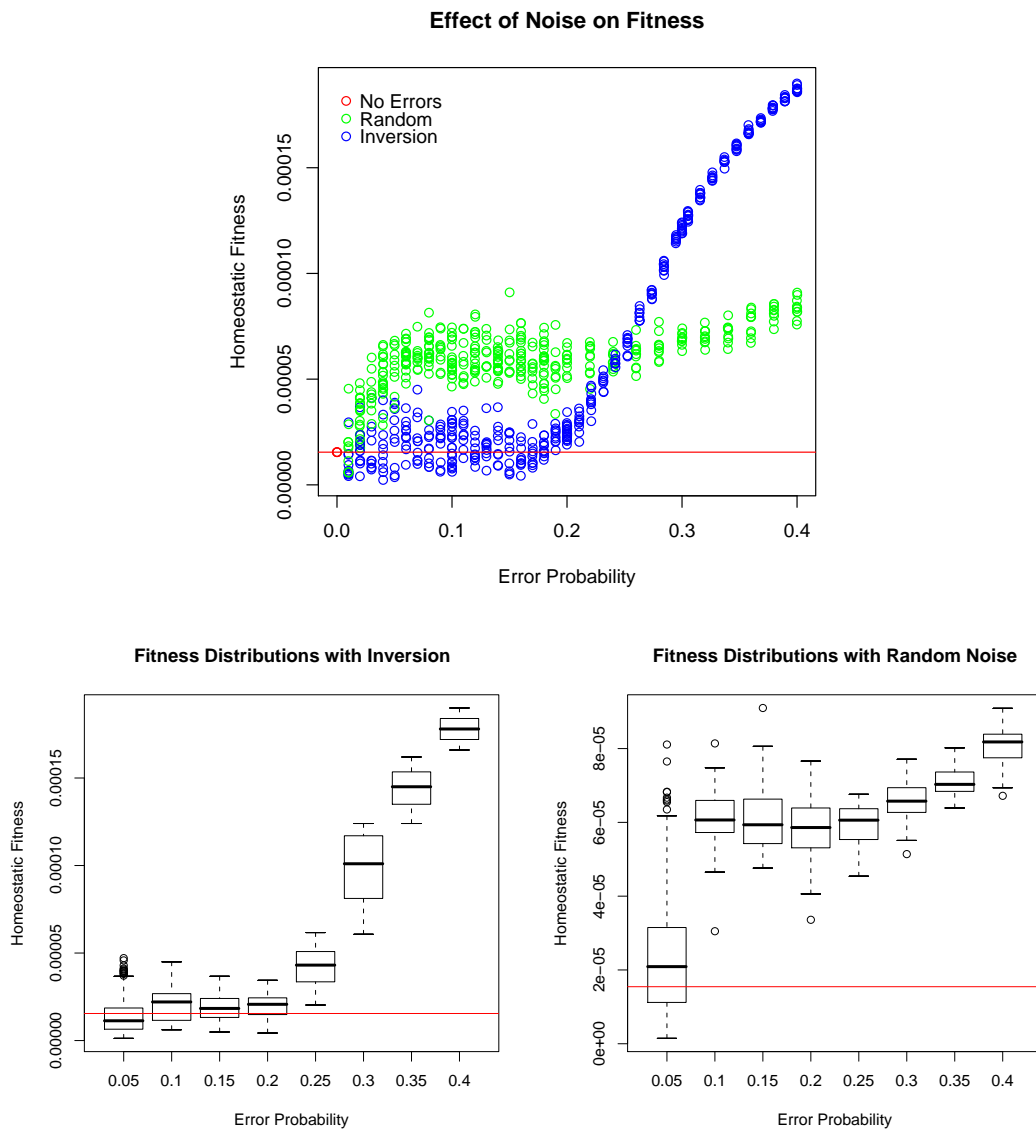
Fitness Distributions with Random Noise

Figure 2: Fitness results from simulations of organism #18 with developmental noise; lower values are more fit. Red lines indicate the noise-free fitness level.

In both cases, even low levels of noise introduce considerable variance to the homeostatic fitness of the phenotype. However, with *inversion* this seems to remain roughly centred on the deterministic value over a considerable range of error probabilities—up to just below 0.2, by which point we must expect multiple malfunctions in almost every cell at every timestep—and a significant proportion of the final phenotypes in this range achieve a *better* homeostasis score. Beyond the threshold, fitness degrades very rapidly.

By contrast, the purely *random* noise leads to deteriorating homeostasis almost immediately; only for very small error probabilities is there any accidental improvement in fitness score. But this deterioration does not continue at the same rate: the fitnesses soon settle down to a broad plateau, never approaching the extreme values reached by high levels of inversion.

Of course, the fitness scores provide only a narrow window onto the behaviour of these organisms; important, as the express criteria of selection, but nonetheless incomplete. To get a qualitative idea of the effects of our introduced stochasticity, we should also look at the actual phenotypes.

Effect of Noise on Fitness (Close Up)

Fitness Distributions with Inversion (Close Up)

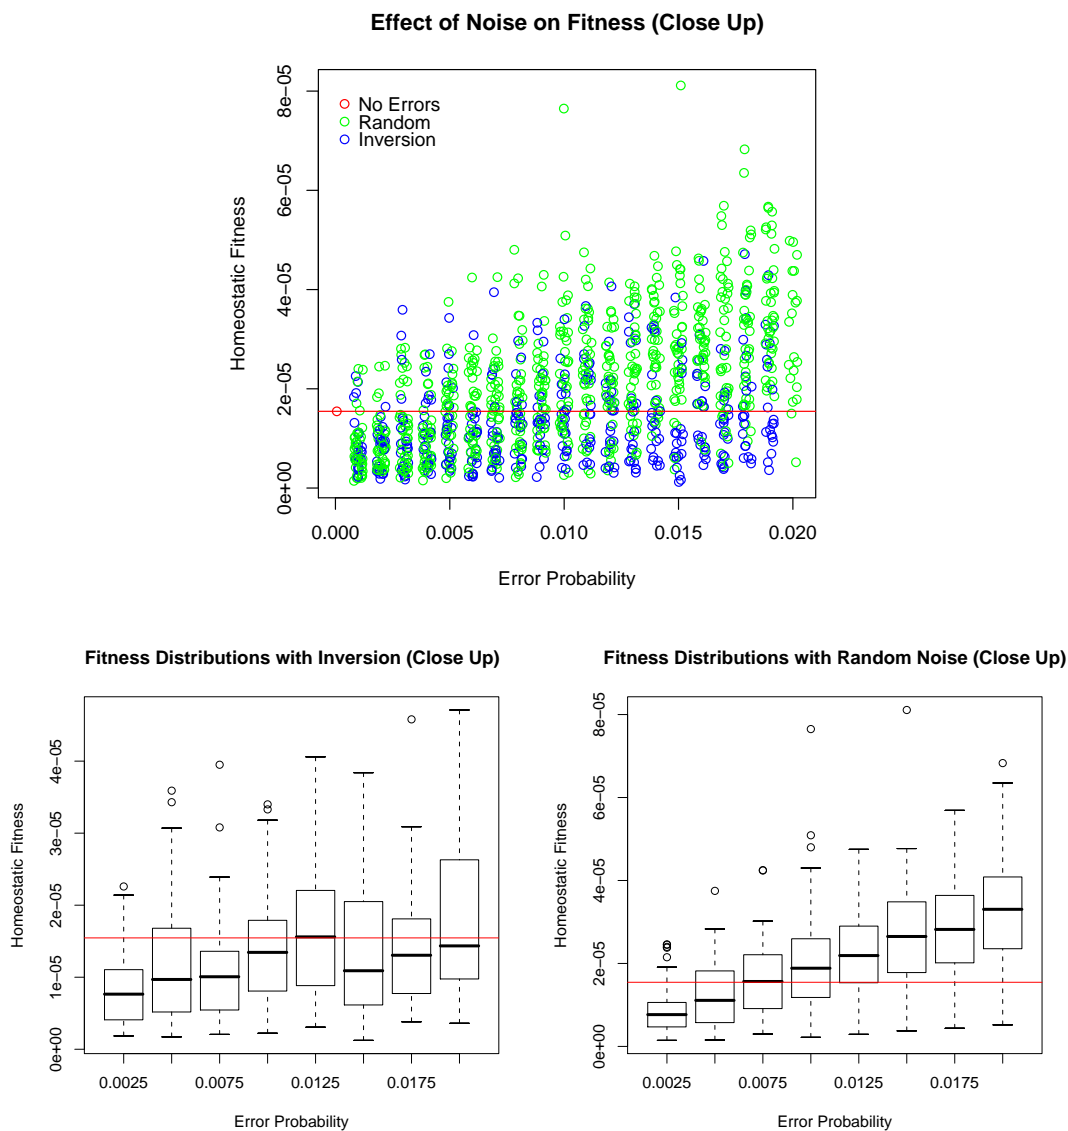
Fitness Distributions with Random Noise (Close Up)

Figure 3: Closer examination of the low noise region.

Several typical examples are shown in Figures 4 and 5.

It is immediately apparent that the main cause of loss of homeostasis at higher levels of noise is unrestrained growth. This is consistent with the original loss of function analysis.

4.4 The Problem of Interpretation

Taken in isolation, the inversion results might seem to indicate a high level of developmental robustness in organism #18: 10% is a pretty large sustained error rate. The random results, on other hand, are much less supportive. Certainly, if the ruleset *has* evolved to exploit a basin of attraction in the phenotypic landscape, it is a relatively limited one, whose edges can quickly be discovered by cumulative random errors. Yet the organism has previously been shown to be robust to acute perturbation.

At this stage, there is insufficient information to justify a full explanation of this apparent contradiction. Nevertheless, some features can be discerned—even if they tell us more about
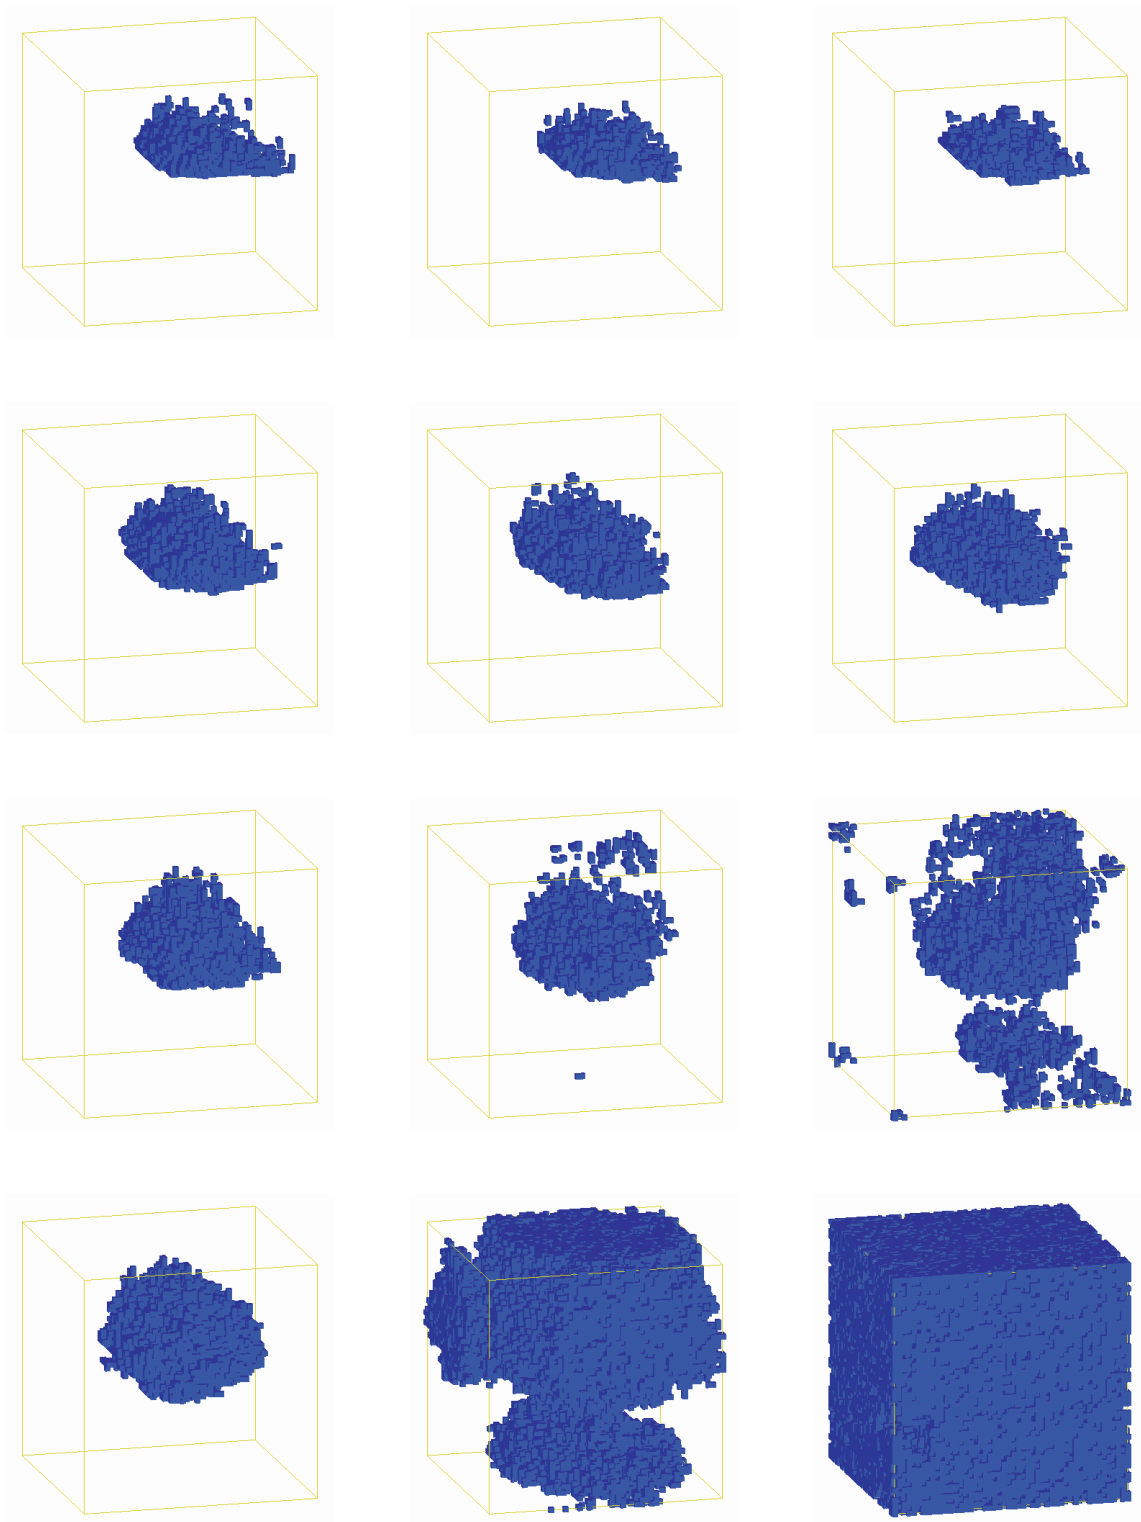
Figure 4: Phenotypes at steps 50, 100 and 150 for inversion probabilities (from top to bottom) 0.01, 0.1, 0.2 and 0.4.
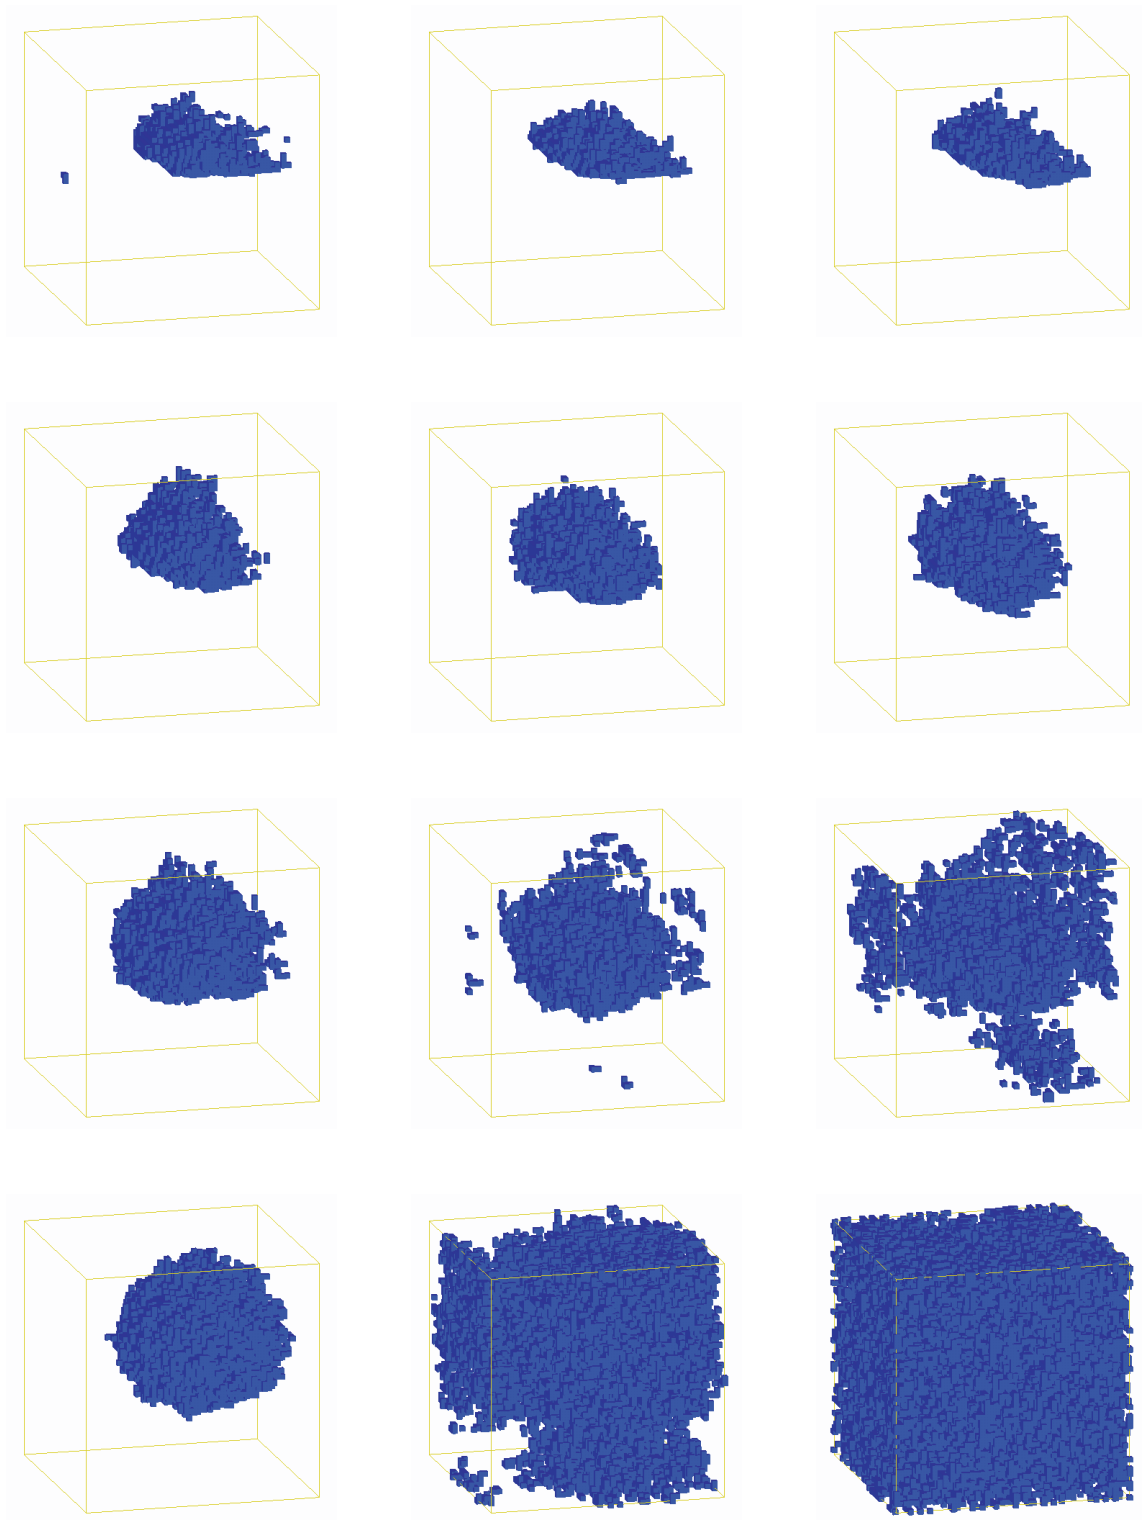
Figure 5: Phenotypes at steps 50, 100 and 150 for random noise probabilities (from top to bottom) 0.01, 0.1, 0.2 and 0.4.

our experiment than the evolution of robustness.

The *directed* nature of inversion promotes a specific kind of deterioration. At high levels, inevitably, this compromises the phenotype—that is, in a sense, its purpose. However, inversion can only subvert the rules already present; it can never introduce new ones. The impact of inversion will therefore depend on the particular distribution of rules in the genome.

While much of the behaviour of a ruleset is intricately unpredictable, it need not all be. In the case of organism #18, almost one third of the rules have conditions that can *never* be matched in the given simulation environment and are thus guaranteed not to be invoked. While these rules have no effect one way or another on the phenotype, they do constitute a sort of redundancy. Inversion will not affect them, so they can 'mop up' a certain amount of perturbation.[7] During evolution, these may provide a pool of mutable genes and hence improve evolvability.

Random changes, on the other hand, can introduce new types of rule which were not present before. Moreover, they are *undirected*. Broadly speaking, a mutant is as likely to act for or against any action, so the cumulative effect of such changes should be approximately—*pace* the structural biases of the CA and fitness functions—*diffusive*. It is perhaps not so surprising, then, that the effect increases more markedly over the short probability range than the long, although it's certainly nothing as clean as $F \propto \sqrt{P}$.

Taken together, the results suggest that, while this organism exhibits a high degree of robustness to specific kinds of perturbations during development, it cannot be described as robust *in general*. Which begs the question: how, if at all, do those particular manifestations reflect the evolutionary process that gave rise to this organism?

5   Conclusions & Future Work

In the study of complex processes such as development and evolution—and particularly their interactions—models like EmbryoCA provide invaluable abstraction, control and repeatability. When phenomena common to natural systems recur in the rarefied context of these models, it can suggest new ways of interpreting their mechanics and help us recognise which factors are more or less important to understanding the processes of life.

However, it is important to be aware of the ways in which the structure and implementation of our models prejudice the results obtained. The inversion results in the previous section illustrate this clearly—taken out of context they could be quite misleading.

The EmbryoCA experiments so far have only begun to explore the model's possibilities, and a great deal of further work would be possible. Having suggested, for example, that the response to environmental perturbation will depend a good deal on the rule distribution, a natural progression would be to run similar tests on a range of other organisms and see how they compare.

One highly suggestive feature of the results obtained above is that a degree of developmental noise—some *non-uniformity* in the application of the rules—frequently gave rise to organisms with improved fitness. Such variability is not available for exploitation by evolution in the current model, but we know it is used extensively by living systems. Despite the inevitable increase in complexity it would entail, the addition of some kind of heritable regulation to the model seems an avenue particularly rich in possibilities.

---

[7]Although the inversion operation is artificial, there is certainly no reason why a rule's condition and consequence should mutate together; consider the operation of transcription factors in real cells.

A   Appendix: EmbryoCA Specifics

EmbryoCA is implemented on a 3-dimensional rectangular lattice. At each location, the cell may either be occupied or unoccupied; automata operate only in occupied cells.

A cell's neighbourhood consists of the 26 adjacent cells, including diagonals (a *Moore neighbourhood*). This is subdivided into 6 overlapping sets of 9 cells each in the directions *north, south, east, west, up* and *down*, each of which is assessed *totalistically*; that is, only the number of occupied cells in each set is considered, rather than their arrangement, which greatly simplifies the rules.[8]

In addition, each automaton has an internal state variable that keeps track of the number of times it has divided.

The genome shared by all automata consists of 100 rules of the form:

> if (variable in interval [lower, upper]) then consequence

The variable tested can be the current timestep, the number of divisions the cell has undergone, or the occupancy count for one of the directional subsets of the neighbourhood. The consequence is a vote either *for* or *against* an action. Actions available to the cell are to *move* or *divide* in a particular direction, or to *die*. Inaction is also allowed, though it is not explicitly included as a distinct rule consequence.

At each timestep, all rules whose condition is met by a particular automaton are tallied and the action with the highest excess of *ayes* over *nays* is performed. In the case of a tie, multiple actions may be selected if they are not incompatible. A minimum threshold is required: inaction is preferred over 'least worst' actions.

In order to be computable in reasonable time, the lattice is finite (a $50 \times 50 \times 50$ cube), with periodic boundary conditions; however a fitness penalty is levied for reaching the sides, to discourage strategies that exploit what is essentially an implementation artefact rather than part of the model proper.

The initial condition is a single occupied cell in the centre of the space, and the CA is run for 150 timesteps, with fitness tests for the phenotype—the pattern of occupied cells—at steps 50, 100 and 150.

Fitness is calculated in two distinct ways. The initial growth phase (up to step 50) is scored according to a simple surface/volume ratio test, favouring larger and more compact phenotypes.

The homeostasis phase is tested at steps 100 and 150, comparing the form at each step with that at step 50 using a combination of *two-point correlation* and the *lineal path function*. These are statistical measures of the distribution of phases in a space: two-point correlation gives the probability that two cells are both occupied as a function of the distance between those cells; lineal path measures the probability that a straight line between two occupied cells passes through no unoccupied cells, again as a function of distance. Taken together, these functions serve to quantify the morphological differences of the phenotypes.

The genetic algorithm operates on a population of 1000, initially random. An *elitist* strategy is used: from each generation, the 50 genomes giving rise to the fittest phenotypes are propagated unchanged to the next. The remaining 950 places are produced by mutation and two-point

---

[8]As the sets are not disjoint, the overall effect is only *semi*-totalistic, since the occupancy pattern of each set will affect the populations of those it intersects.

crossover, using parents chosen by *tournament*—a fitness face-off between randomly-selected individuals from the whole generation (including the elite).

Mutation occurs with a probability of 0.05 per rule in the genome, and entails replacing the rule entirely at random: there are no smaller mutable units.

The GA is run for 30 generations.

The system is implemented in the Java language, using the ECJ library for some GA infrastructure. OpenDX is used to visualise the 3D CA phenotypes.

References

[1] Wolfgang Banzhaf, Guillaume Beslon, Steffen Christensen, James Foster, François Képès, Virginie Lefort, Julian Miller, Miroslav Radman, and Jeremy Ramsden. From artificial evolution to computational evolution: A research agenda. *Nature Reviews Genetics*, 7:729–735, September 2006.

[2] David Basanta, Mark Miodownik, and Buzz Baum. The evolution of robust homeostasis and wound-healing in artificial multicellular organisms. *Submitted for publication*, 2006.

[3] David Basanta, Mark Miodownik, Peter Bentley, and Elizabeth Holm. Evolving cellular automata to grow microstructures. In *Genetic Programming, Proceedings of the 6th European Conference, EuroGP*, pages 1–10. Springer, 2003.

[4] David Basanta, Mark Miodownik, Peter Bentley, and Elizabeth Holm. Investigating the evolvability of an embryological model based on CA. In *Proceedings workshops of 9th conference on Artificial Life (ALife9)*, 2004.

[5] G Bard Ermentrout and Leah Edelstein-Keshet. Cellular automata approaches to biological modelling. *Journal of Theoretical Biology*, (160):97–133, 1993.

[6] Diego Federici and Keith Downing. Evolution and development of a multicellular organism: Scalability, resilience and neutral complexification. *Artificial Life*, (12):381–409, 2006.

[7] Sanjeev Kumar and Peter Bentley. *An introduction to computational development*, chapter 1, pages 1–43. In [8], 2003.

[8] Sanjeev Kumar and Peter Bentley, editors. *On Growth, Form and Computers*. Elsevier Academic Press, London, UK, 2003.

[9] Richard Lenski, Jeffrey Barrick, and Charles Ofria. Balancing robustness and evolvability. *PLoS Biology*, 4(12), 2006.

[10] Julian Miller and Wolfgang Banzhaf. *Evolving the program for a cell: from French flags to Boolean circuits*, chapter 15, pages 278–301. In Kumar and Bentley [8], 2003.

[11] Mark Miodownik. *Using mechanics to map genotype to phenotype*, chapter 11, pages 203–219. In Kumar and Bentley [8], 2003.

[12] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, Massachusetts, US, 1996.

[13] Michael Redd, Lisa Cooper, Will Wood, Brian Stramer, and Paul Martin. Wound healing and inflammation: embryos reveal the way to perfect repair. *Philosophical Transactions of the Royal Society B*, (359):777–784, 2004.

[14] Ian Stewart. *Broken symmetries and biological patterns*, chapter 10, pages 181–202. In Kumar and Bentley [8], 2003.

[15] Andreas Wagner. Genetic redundancy caused by gene duplication and its evolution in networks of transcriptional regulators. *Biological Cybernetics*, (74):557–567, 1996.

[16] Andreas Wagner. Robustness, evolvability and neutrality. *FEBS Letters*, (579):1772–1778, February 2005.

[17] Stephen Wolfram. *A New Kind of Science*. Wolfram Media, Champaign, Illinois, US, 2002.

[18] Lewis Wolpert. *Relationships between development and evolution*, chapter 2, pages 47–63. In Kumar and Bentley [8], 2003.