Augmenting Artificial Development with Local Fitness

Taras Kowaliw* and Wolfgang Banzhaf
Department of Computer Science,
Memorial University of Newfoundland
http://kowaliw.ca

Presented at the IEEE CEC, Tuesday May 19th, 2009

* Now with the Centre for Electronic Media Arts, Monash University.

AD Design

Consider these tree designs, the one on the left being a viable means of distributing resources, the one on the right not.

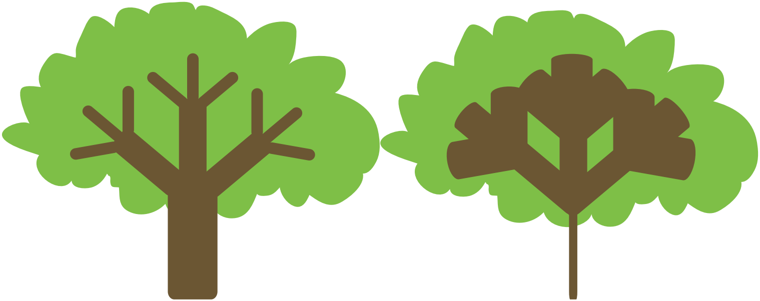

Will some given evolutionary representation have both designs as equally likely phenotypes?

Natural Adaptive Design

"Amphibious plants" (e.g. *myriophyllum heterophyllum*) select leaf type on the basis of external factors about the meristem during growth. The height at which leaf type changes can be determined adaptively, and hence needn't necessarily be specified genetically.

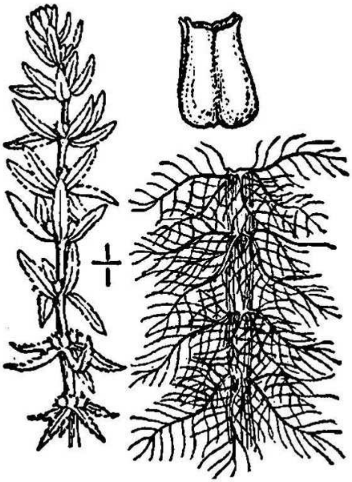

Image from: USDA-NRCS PLANTS Database / Britton, N.L., and A. Brown. 1913. An illustrated flora of the northern United States, Canada and the British Possessions. Vol. 2: 616.

Genetically specified?

During (real world) vasculogenesis,

- Murray's law: the cube of the size of the parent vessel is approximately the sum of the cubes of the child vessels.
- Are the relative widths of the vessels genetically specified, or the result of natural constraints during the developmental process?
- We suspect the latter; That the relative widths emerge as an interplay between genetics and the availability of resources resulting from previous dynamics.

Motivation: An Artificial Analogue?

- The specification of morphological design as an adaptive response to environmental stimulus is an important property of natural embryogenesis
- Continued necessity of viable designs during the developmental process significantly constrains the space of possible phenotypic outcomes, probably increasing evolvability.
- Perhaps functional evaluation during development will increase evolvability in the artificial case.

Potential Extent of Environmental Influence

In nature, it is possible that the entire gross morphology of an organism be specified through external influences

Young and mature Rata vines growing over trees.

Toy Model of Vasculogenesis

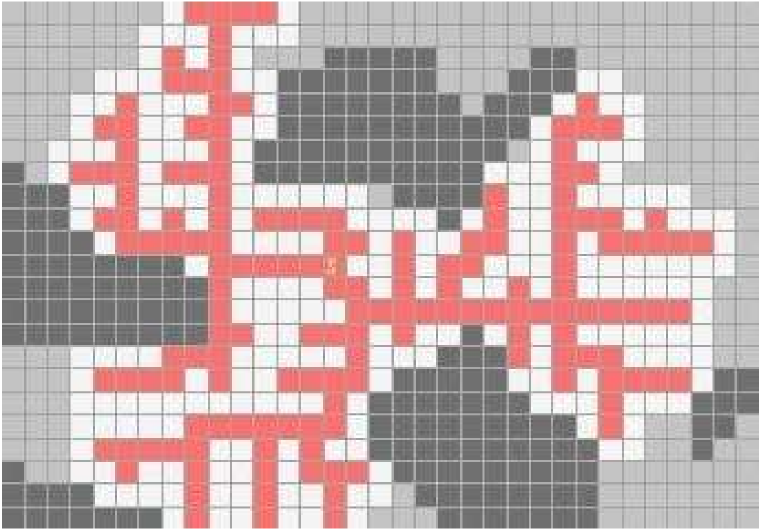

A Highly Simplified Vascular Model

Given some two-dimensional connected environment, we're interested in configurations of cells, where transport cells carry fluid from a central start cell.

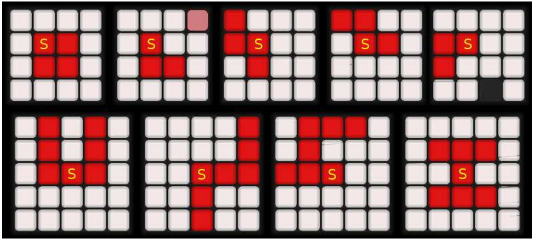

Fitness is the number of served normal cells. i.e. we want a network of transport cells which will neighbour as much of the space as possible, but take as little space as possible. Global optimum around $\frac{2}{3}$.

Random Environments

We can work with the empty environment,

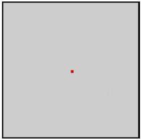

Or we can work with *random environments*, where connected 2D environments are created using a system of random points and Voronoi diagrams. (Visual) complexity is controlled through a simple parameter.

Code: http://kowaliw.ca/envs.html.

Vasculogenesis Problem

Our vasculogenesis problem is:
Given an environment with a central start cell, design a (good) vascular system; that is, gradually create a design which efficiently distributes fluid.

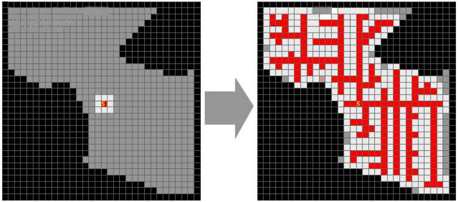

Local Fitness-enhanced Growth Strategies

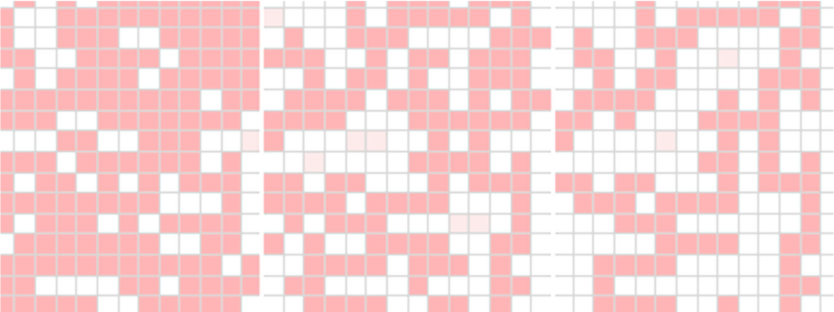

Global Fitness-enhanced AD (GF-AD)

- Given some AD process AD and a fitness function f.
- At every proposed action during development, assume we compute the fitness of the original structure and the action-affected structure.
- An action is only undertaken if f increases.
- Development becomes hill-climbing
- Time requirements become large: $O(GF\text{-}AD) = O(AD) \times O(f)$.

Local Fitness Functions

- we can define some local versions of our fitness function, defined only in neighbourhoods of individual points (i.e. constant time).
- these can augment the developmental process with no or minimal change to (asymptotic) running time

Blind Local Fitness (BLF)

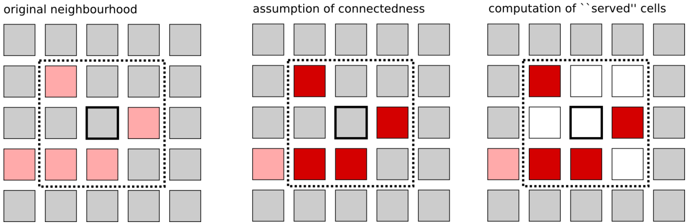

Takes a local neighbourhood of cells. Assumes that any transport cells are connected to the start cell. Computes fitness in the local neighbourhood. Does not affect the computational complexity of the developmental process (asymptotically).

Random Greedy Growth Strategies

How well does random growth do, when guided by our fitness functions?

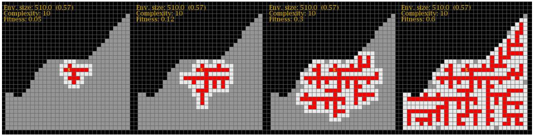

Random blind local fitness-enhanced growth (RLG)

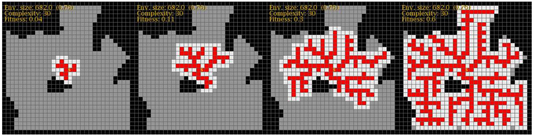

Random global fitness-enhanced growth (RGG)

Random Knockout Strategies

If we instead try knockout strategies, our blind local fitness function is useless; The random knockout tends towards evenly spaced unconnected transport cells.

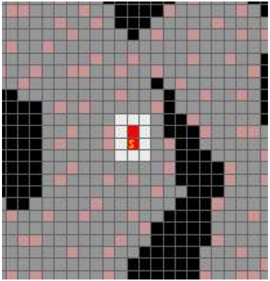

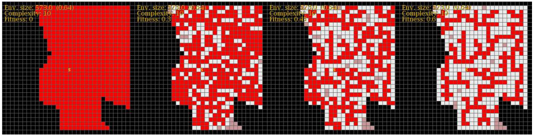

The global fitness-enhanced version works more effectively.

Cellular Automata (CA)

As a control, we can evolve cellular automata to solve the vasculogenesis problem. Development takes time $O(|E|)$, where $|E|$ is the size of the environment.

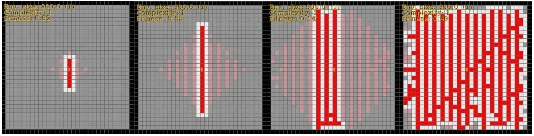

CA growth

Blind Local Fitness-enhances Constructive CAs (BLF-CCA)

BLF-enhanced CAs do terribly, for the same reason as the Random Local Knockout strategy. If we restrict to Constructive CAs (i.e. CAs which allow change from cells of type “normal” to “transport” but not vice versa), we obtain better results, with development time $O(|E|)$.

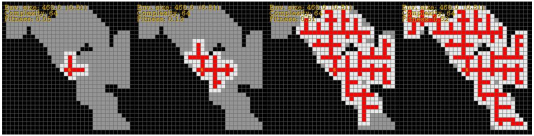

Blind Local Fitness-enhances Constructive CA (BLF-CCA)

Sighted Local Fitness (SLF)

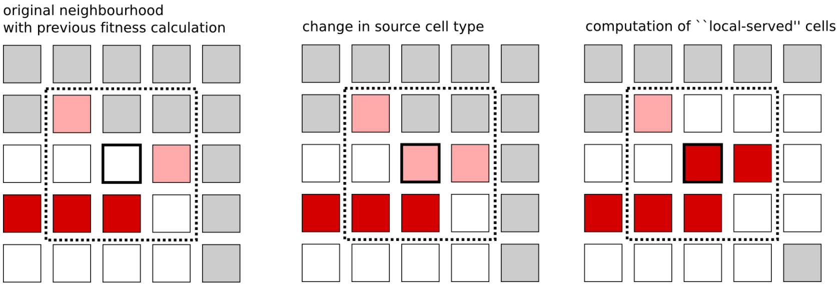

Takes previous fitness computation as a guide for whether or not existing transport cells are connected. If the developmental process can be decomposed into stages (e.g. discrete time over cells), adds a linear factor to computational complexity.

Sighted Local Fitness-enhanced CAs (SLF-CA)

SLF-enhanced CAs can be evolved with development time $O(|E|^2)$.

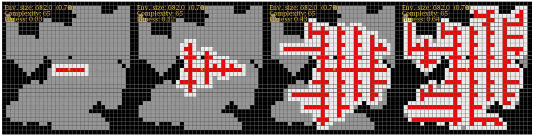

Sighted Local Fitness-enhanced CA (SLF-CA)

Global Fitness-enhanced CAs (GL-CA)

GF-CAs can also be evolved, with highest computational expense $O(|E|^3)$.

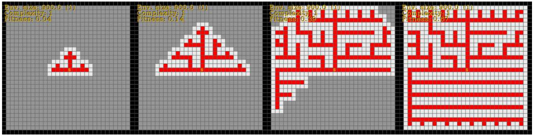

Global Fitness-enhanced CA (GL-CA)

Analysis

Comparison in the Trivial Environment

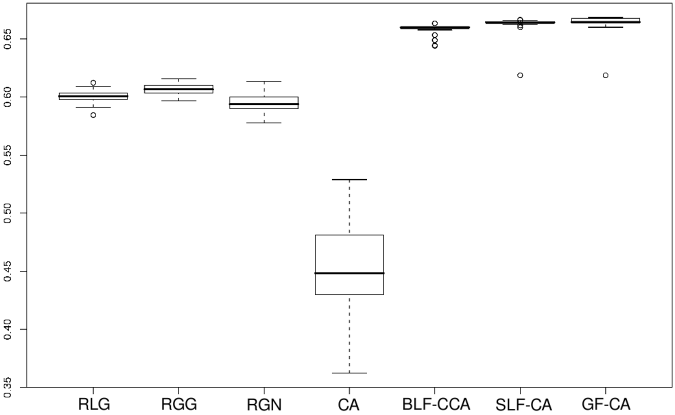

Comparison in Random Environments

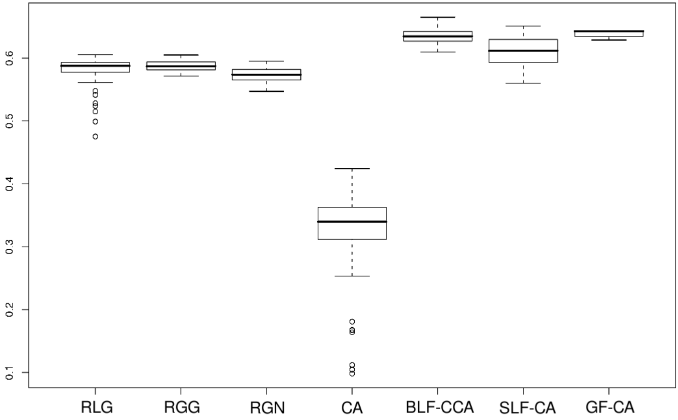

Comparison in Random Environments

Correlation analysis shows that there is, for the CA control group, an inverse correlation between the complexity of the environments and the performance of the CA evolution. This correlation existed for the BLF-CCA strategy as well, but was insignificant for all other strategies.

Course of Evolution: Trivial Environments

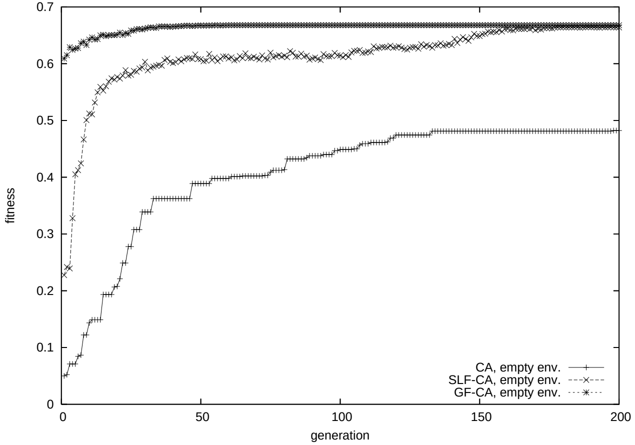

Course of Evolution: Random Environments

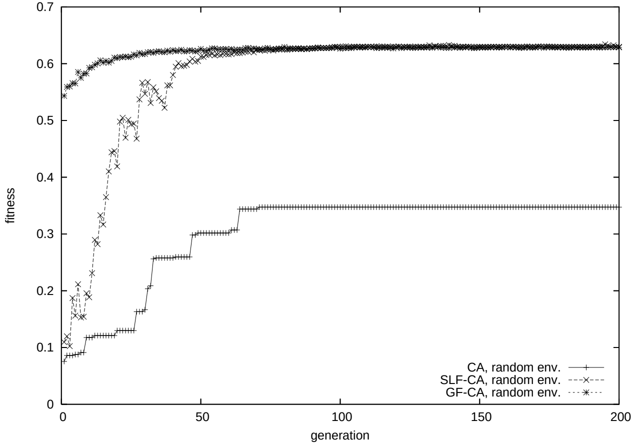

Conclusions

- the inclusion of local fitness quickly increases the efficacy of the developed solutions (even the random growth strategies outperform the evolved CA!)
- use of global fitness or a sighted local fitness effectively eliminates dependency on problem complexity
- sighted local fitness serves as a means of including some global information in a primarily local development

Using a local fitness function (or the global fitness function) does indeed improve evolvability in the vasculogenesis problem domain, as (we suspect) the viability requirement improves evolvability in natural embryogenesis.

“Augmenting Artificial Development with Local Fitness”

Taras Kowaliw and Wolfgang Banzhaf
http://kowaliw.ca