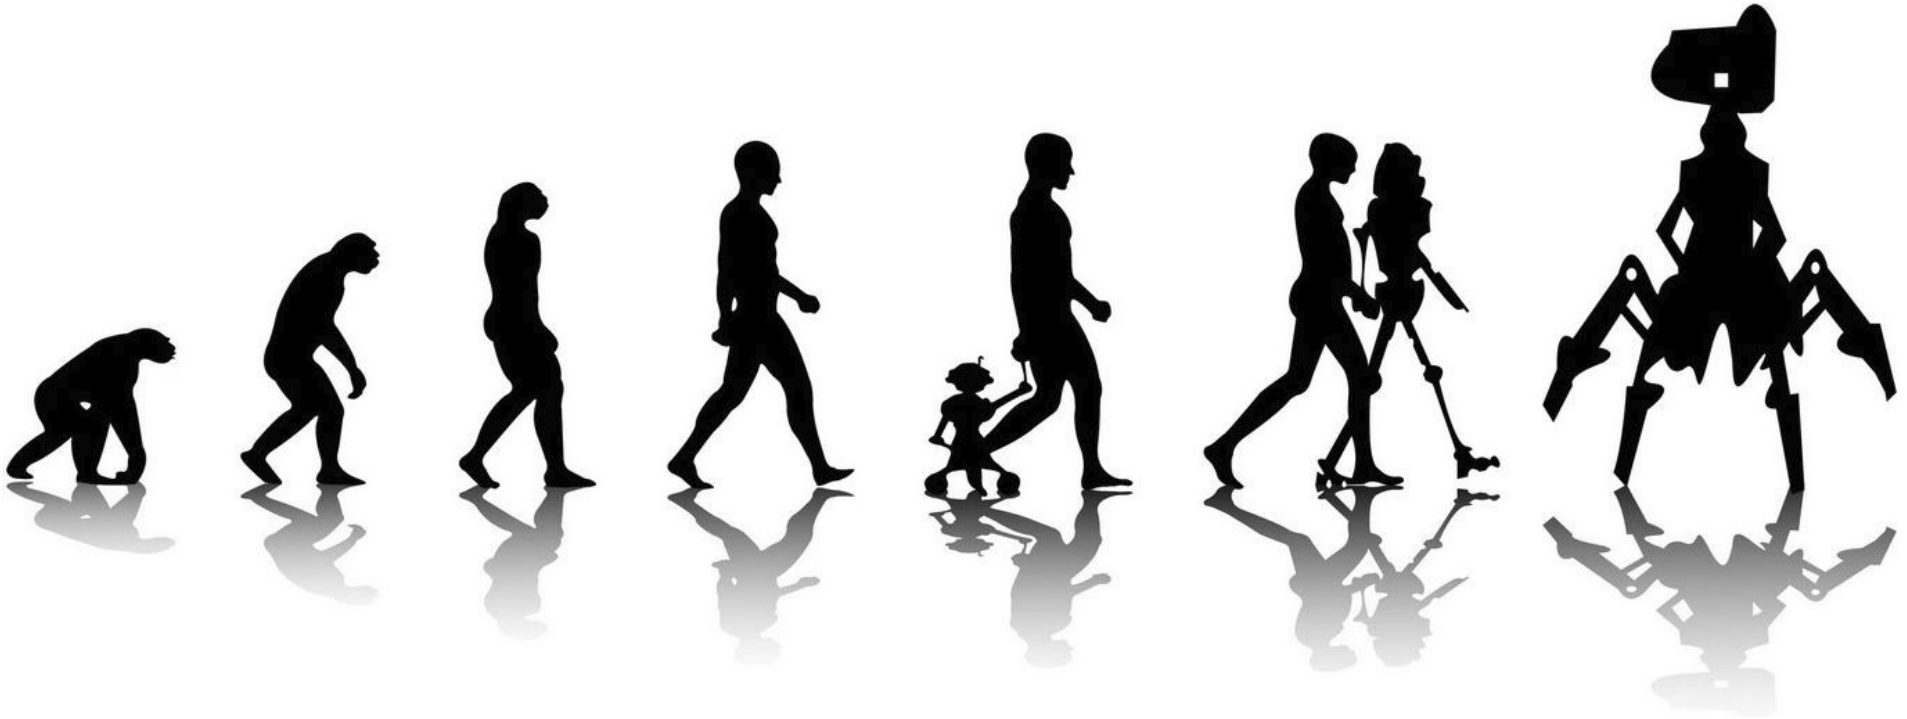

Modern Robotics: Evolutionary Robotics

COSC 4560 / COSC 5560

Professor Cheney

2/19/18

DPM	molecules	physics	evo-devo role	effect
ADH	cadherins	adhesion	multicellularity	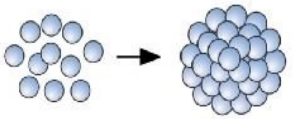
LAT	Notch	lateral inhibition	coexistence of alternative cell states	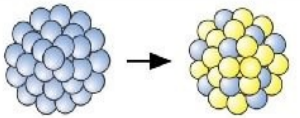
DAD	cadherins	differential adhesion	phase separation; tissue multilayering	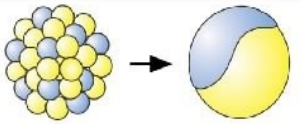
POL_a	Wnt	cell surface anisotropy	topological change; interior cavities	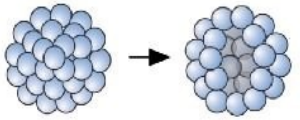
POL_p	Wnt	cell shape anisotropy	tissue elongation	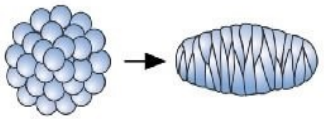
ECM	chitin; collagen	stiffness; dispersal	tissue solidification; elasticity; EMT	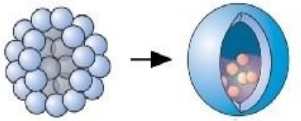
OSC	Wnt + Notch	synchrony of oscillatin	morphogenetic fields; segmentation	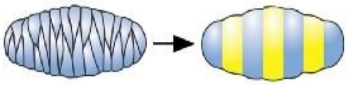
MOR	TGF-β/BMP; FGF; Hh	diffusion	pattern formation	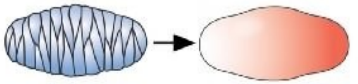
TUR	MOR + Wnt + Notch	dissipative structure	segmentation; periodic patterning	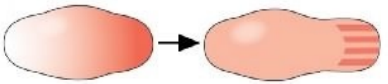

[Newman and Bhat, 2008]

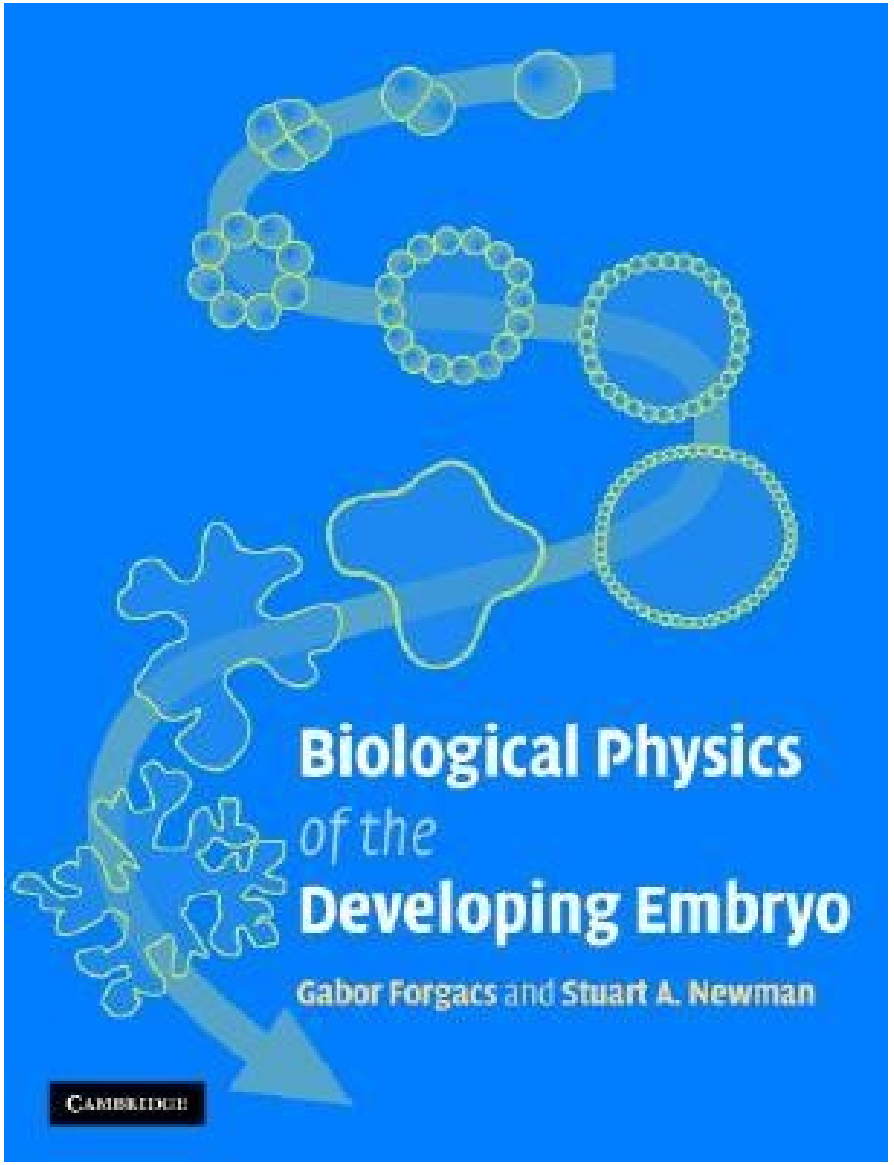
Biological Physics
of the
Developing Embryo
Gabor Forgacs and Stuart A. Newman
CAMBRIDGE

We haven't talked about evolving network topologies yet!

NeuroEvolution of Augmenting Topologies (NEAT)

Genetic Operators:

Mutate Weight
Add Edge
Remove Edge

Add Node
Remove Node
Modify Activation Function

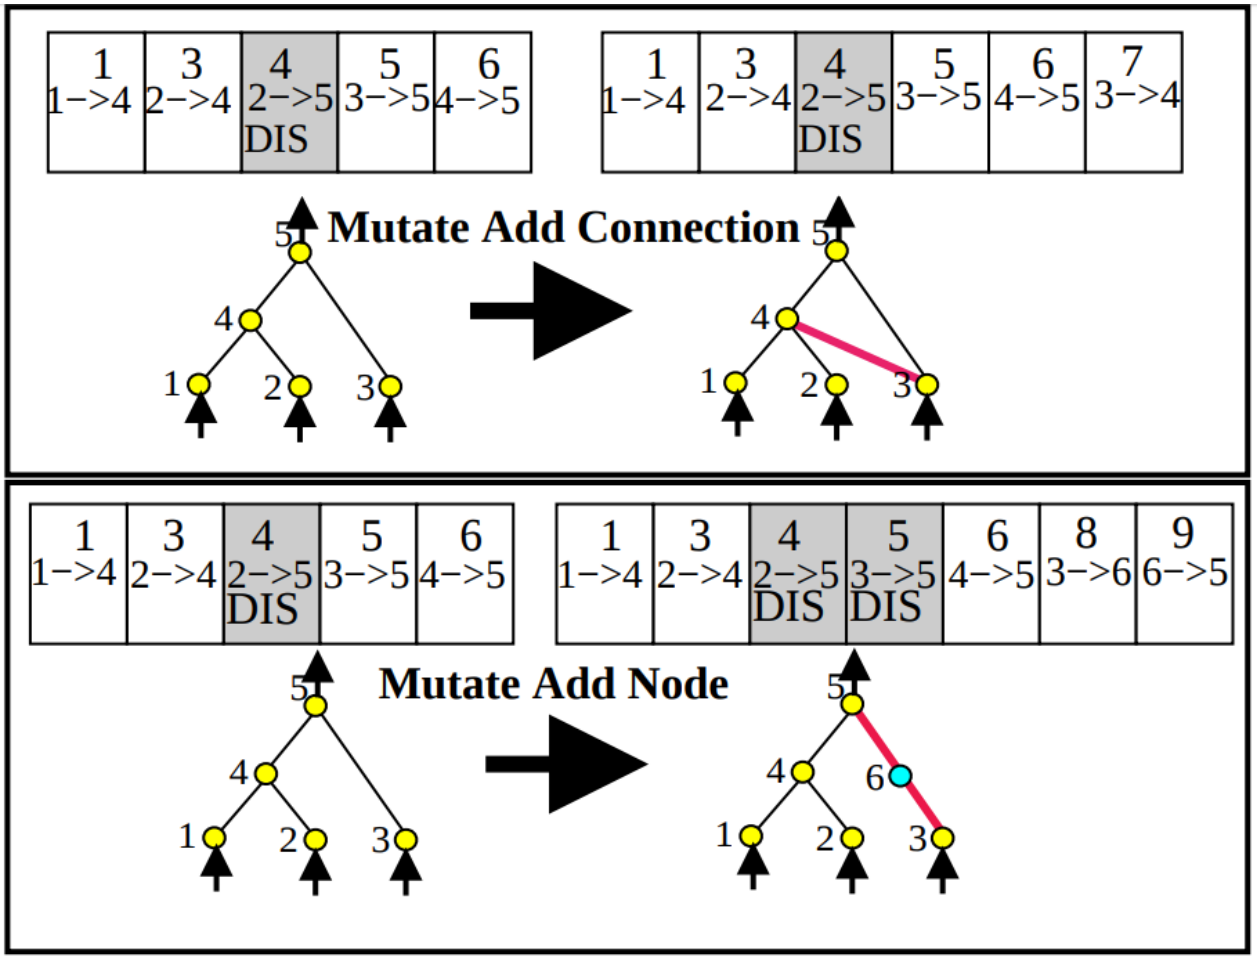

Figure 3: **The two types of structural mutation in NEAT.** Both types, adding a connection and adding a node, are illustrated with the connection genes of a network shown above their phenotypes. The top number in each genome is the *innovation number* of that gene. The innovation numbers are historical markers that identify the original historical ancestor of each gene. New genes are assigned new increasingly higher numbers. In adding a connection, a single new connection gene is added to the end of the genome, and given the next available innovation number. In adding a new node, the connection gene being split is disabled, and two new connection genes are added to the end the genome. The new node is between the two new connections. A new node gene (not depicted) representing this new node is added to the genome as well.

Complexification:

Start with small network

higher rate of adding node/edge

than of removing a node/edge

(search the space of small networks first – similar idea of iterative deepening search)

Intelligent crossover with “historical markings”

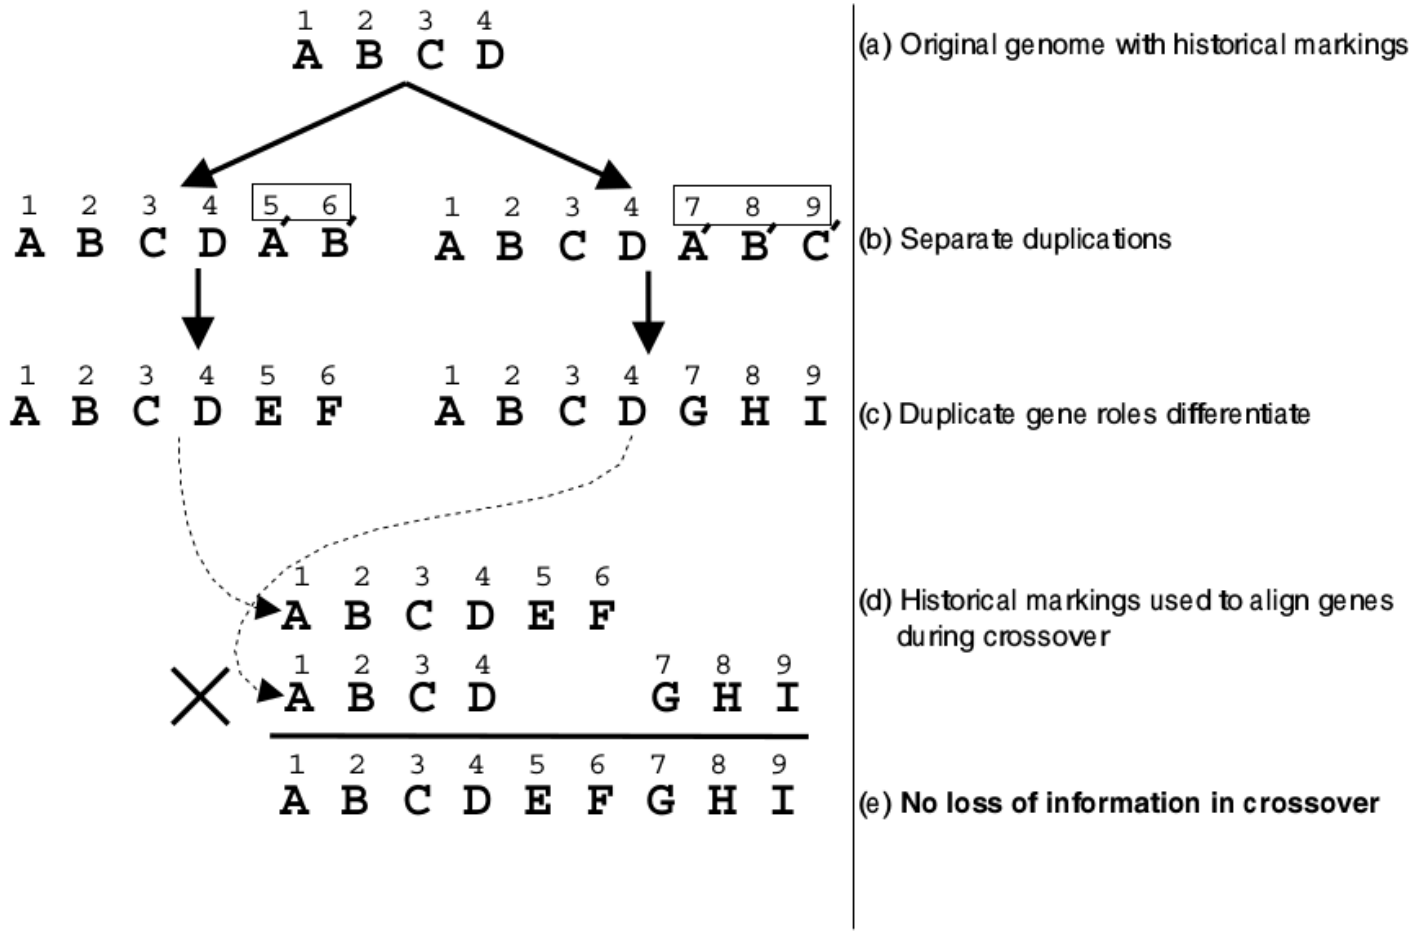

Figure 10. Solving the variable-length genome problem with historical markings. Historical markings are numbers assigned to each gene that represent the order in which new genes appeared over evolution. (a) The original genome contains four genes—*A*, *B*, *C*, and *D*, assigned historical markings 1 through 4. (b) When new genes appear through duplication, they are assigned numbers in the order in which they appear. Assuming the duplication on the left happened before the one on the right, the new genes—*A′* and *B′*, and *A′*, *B′*, and *C′*—are assigned the numbers 5 through 9. (c) As the products of the duplicate genes differentiate, their historical markings continue to serve as a record of their origins. (d) During crossover, those genes that have matching historical markings are aligned, while those that are disjoint are purposely not aligned. (e) The result is that any kind of crossover can preserve the information and relationships between all the genes in variable length genomes by utilizing the historical markings. Historical markings are an abstraction of synapsis, the process used in nature to match up alleles of the same trait during crossover [63, 70].

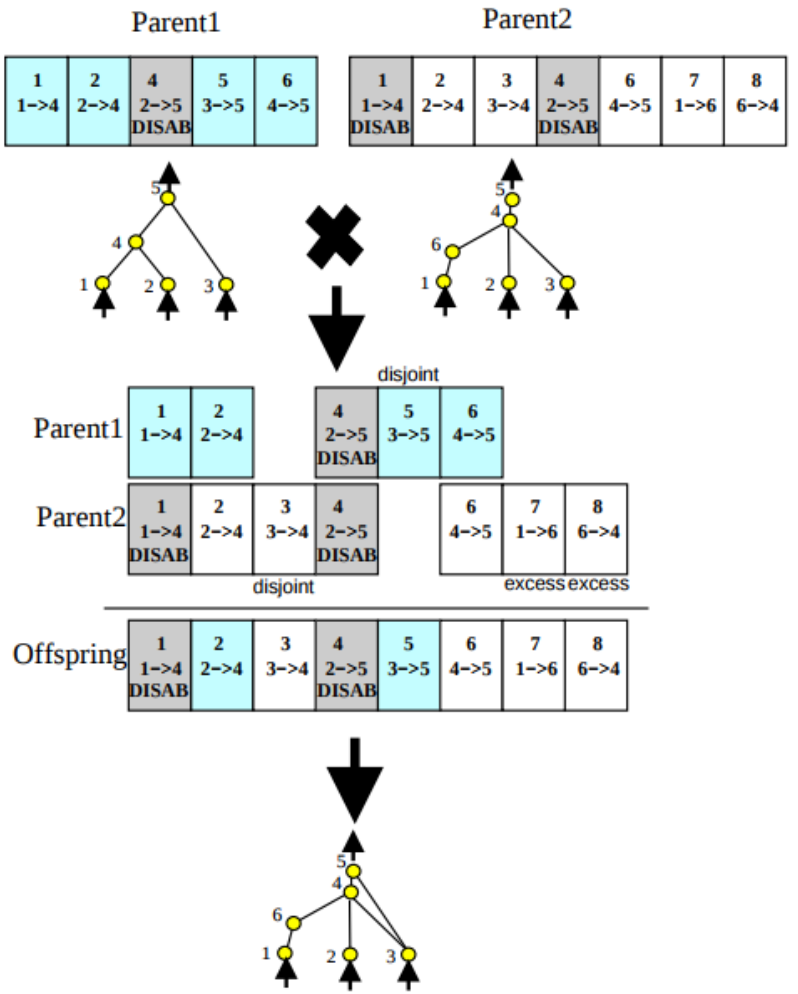

Figure 4: **Matching up genomes for different network topologies using innovation numbers.** Although Parent 1 and Parent 2 look different, their innovation numbers (shown at the top of each gene) tell us which genes match up with which. Even without any topological analysis, a new structure that combines the overlapping parts of the two parents as well as their different parts can be created. Matching genes are inherited randomly, whereas disjoint genes (those that do not match in the middle) and excess genes (those that do not match in the end) are inherited from the more fit parent. In this case, equal fitnesses are assumed, so the disjoint and excess genes are also inherited randomly.

Protecting Innovation through Speciation

Speciating the population allows organisms to compete primarily within their own niches instead of with the population at large. This way, topological innovations are protected in a new niche where they have time to optimize their structure through competition within the niche. The idea is to divide the population into species such that similar topologies are in the same species. This task appears to be a topology matching problem. However, it again turns out that historical markings offer an efficient solution.

The number of excess and disjoint genes between a pair of genomes is a natural measure of their compatibility distance. The more disjoint two genomes are, the less evolutionary history they share, and thus the less compatible they are. Therefore, we can measure the compatibility distance δ of different structures in NEAT as a simple linear combination of the number of excess (E) and disjoint (D) genes, as well as the average weight differences of matching genes ($\overline{W}$), including disabled genes:

$$\delta = \frac{c_1 E}{N} + \frac{c_2 D}{N} + c_3 \cdot \overline{W}. \tag{1}$$

The distance measure δ allows us to speciate using a compatibility threshold δ_t.

Protecting Innovation through Speciation

As the reproduction mechanism for NEAT, we use *explicit fitness sharing* (Goldberg and Richardson 1987), where organisms in the same species must share the fitness of their niche. Thus, a species cannot afford to become too big even if many of its organisms perform well. Therefore, any one species is unlikely to take over the entire population, which is crucial for speciated evolution to work. The adjusted fitness f'_i for organism i is calculated according to its distance δ from every other organism j in the population:

$$f'_i = \frac{f_i}{\sum_{j=1}^{n} \mathrm{sh}(\delta(i,j))}. \tag{2}$$

The sharing function sh is set to 0 when distance $\delta(i,j)$ is above the threshold δ_t; otherwise, $\mathrm{sh}(\delta(i,j))$ is set to 1 (Spears 1995). Thus, $\sum_{j=1}^{n} \mathrm{sh}(\delta(i,j))$ reduces to the number of organisms in the same species as organism i. This reduction is natural since species are already clustered by compatibility using the threshold δ_t. Every species is assigned a potentially different number of offspring in proportion to the sum of adjusted fitnesses f'_i of its member organisms. Species then reproduce by first eliminating the lowest performing members from the population. The entire population is then replaced by the offspring of the remaining organisms in each species.[2]

The net effect of speciating the population is that topological innovation is protected. The final goal of the system, then, is to perform the search for a solution as efficiently as possible. This goal is achieved through minimizing the dimensionality of the search space.

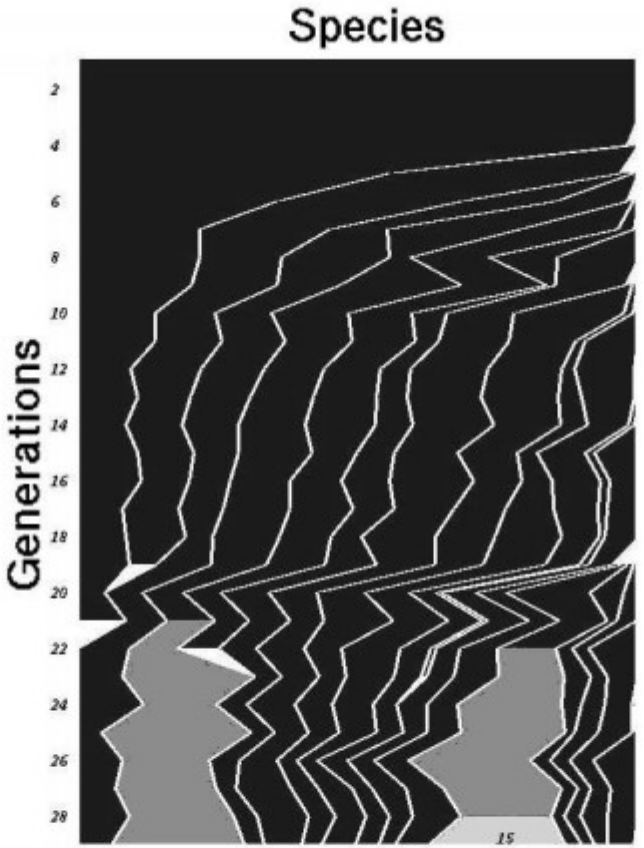

Figure 7: **Visualizing speciation during a run of the double pole balancing with velocity information task..** Two species begin to close in on a solution soon after the 20th generation. Around the same time, some of the oldest species become extinct.

Compositional Pattern Producing Network (CPPN)

+

NeuroEvolution of Augmenting Topologies (NEAT)

CPPN-NEAT

Evolving Mophologies with CPPNs

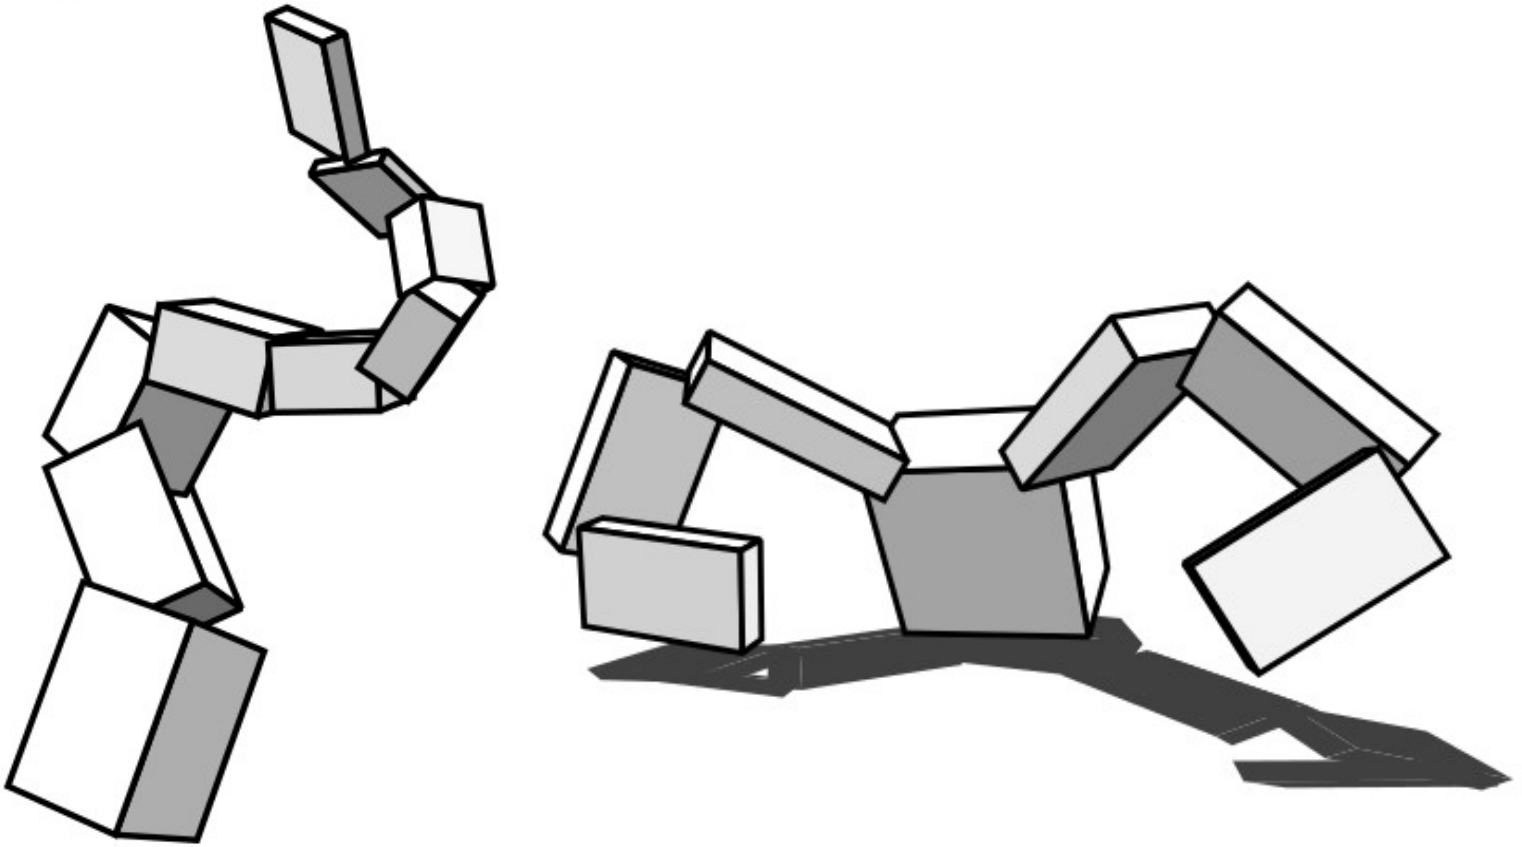

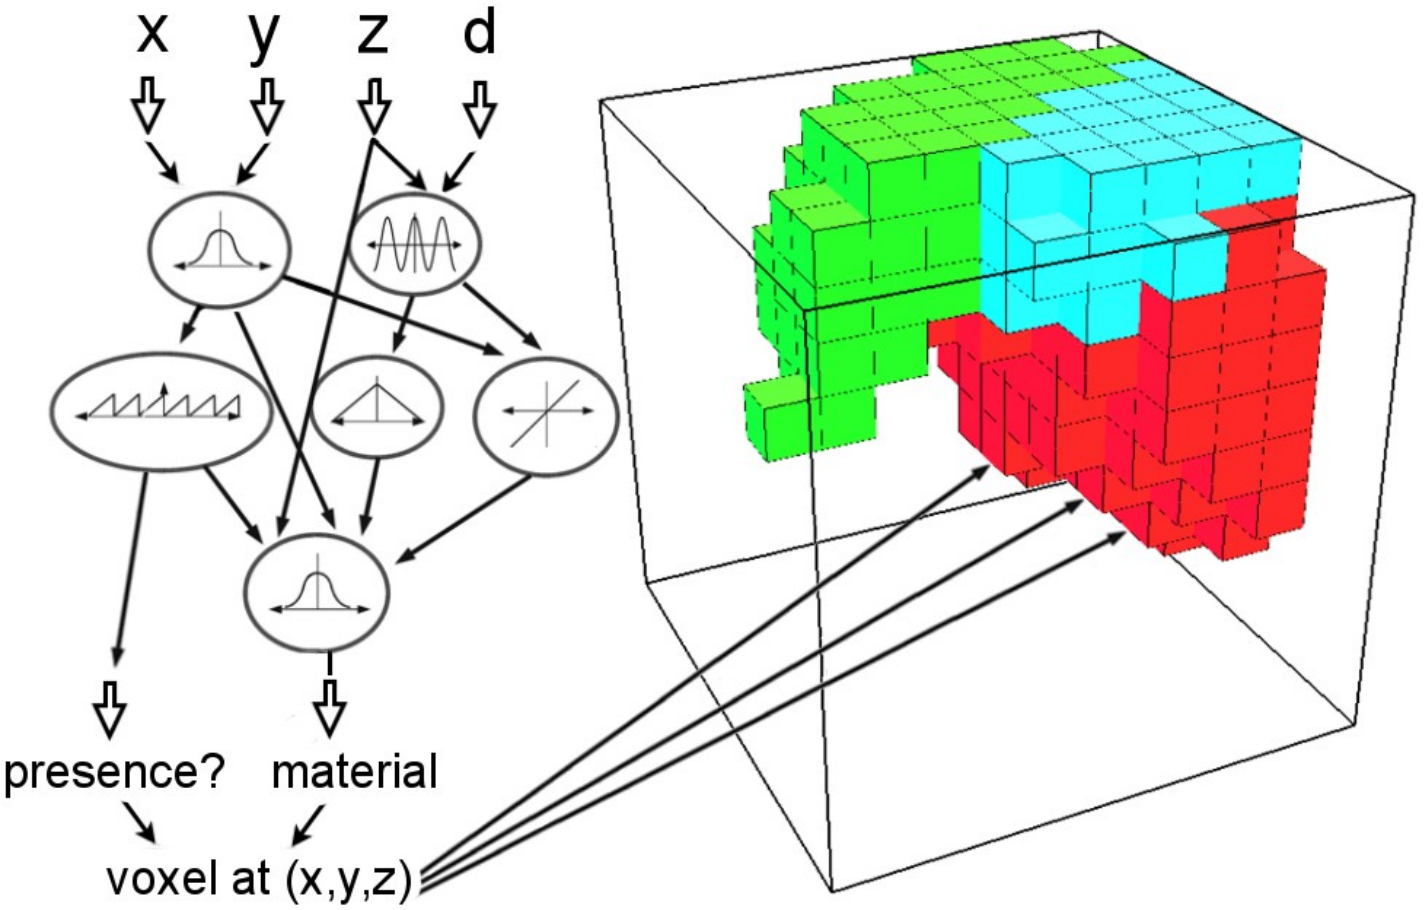
x
y
z
d
presence?
material
voxel at (x,y,z)

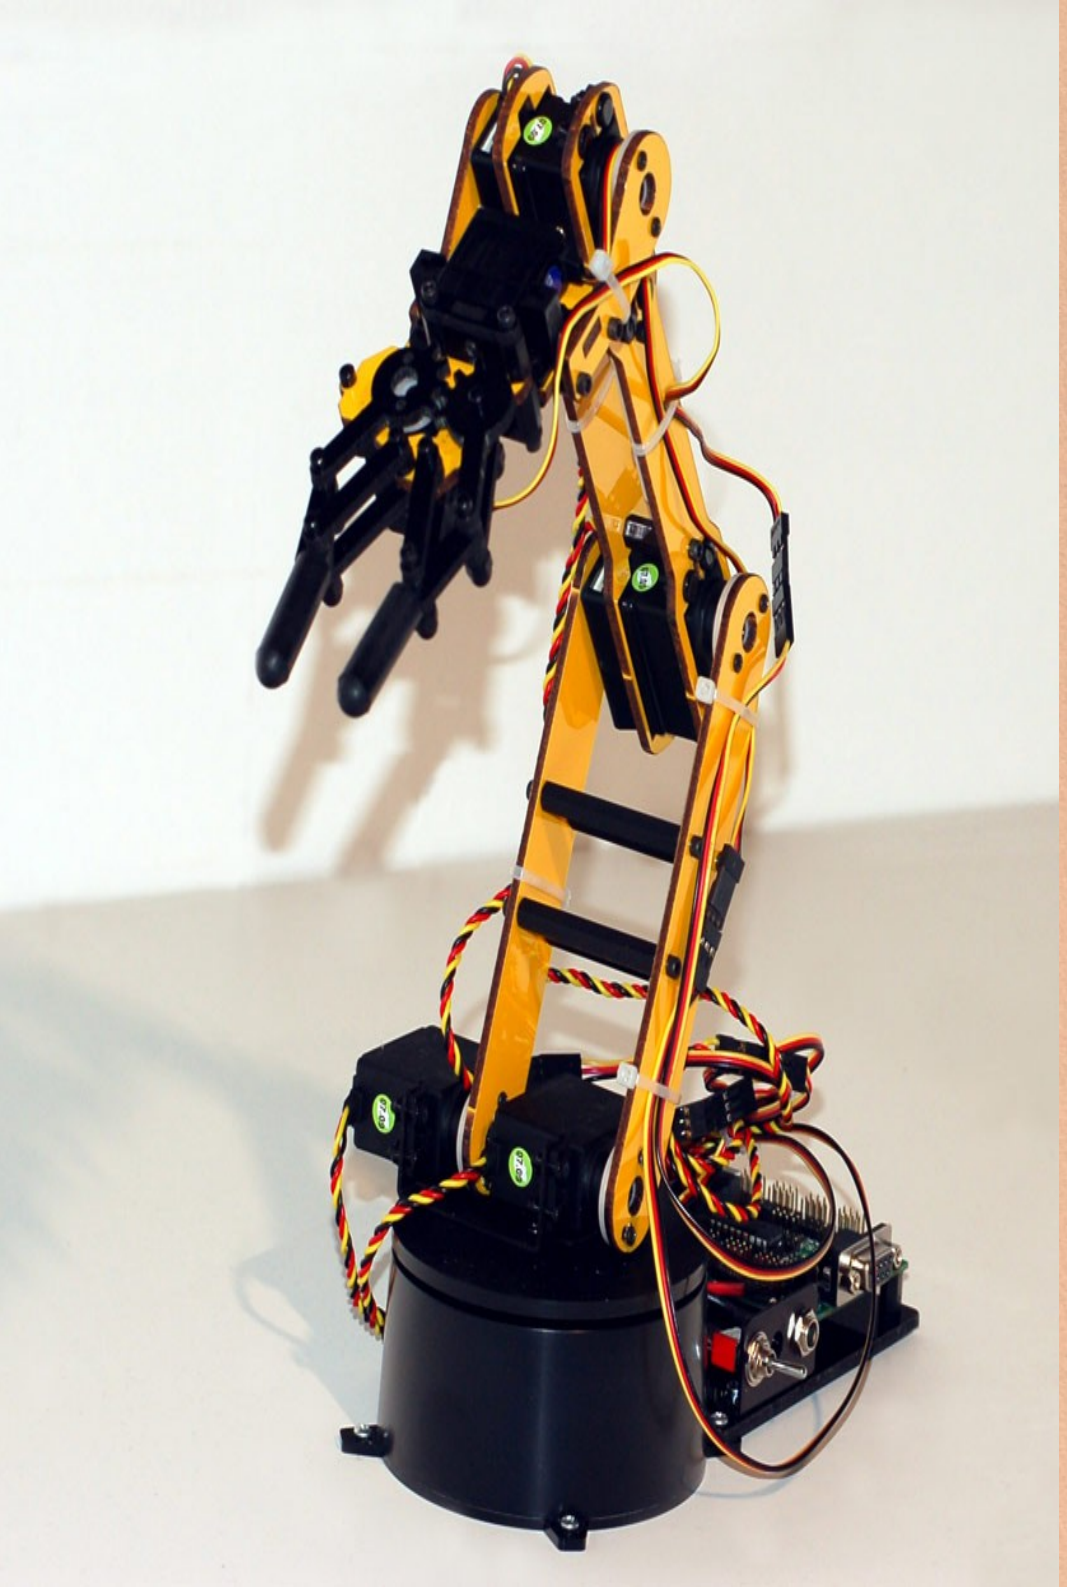

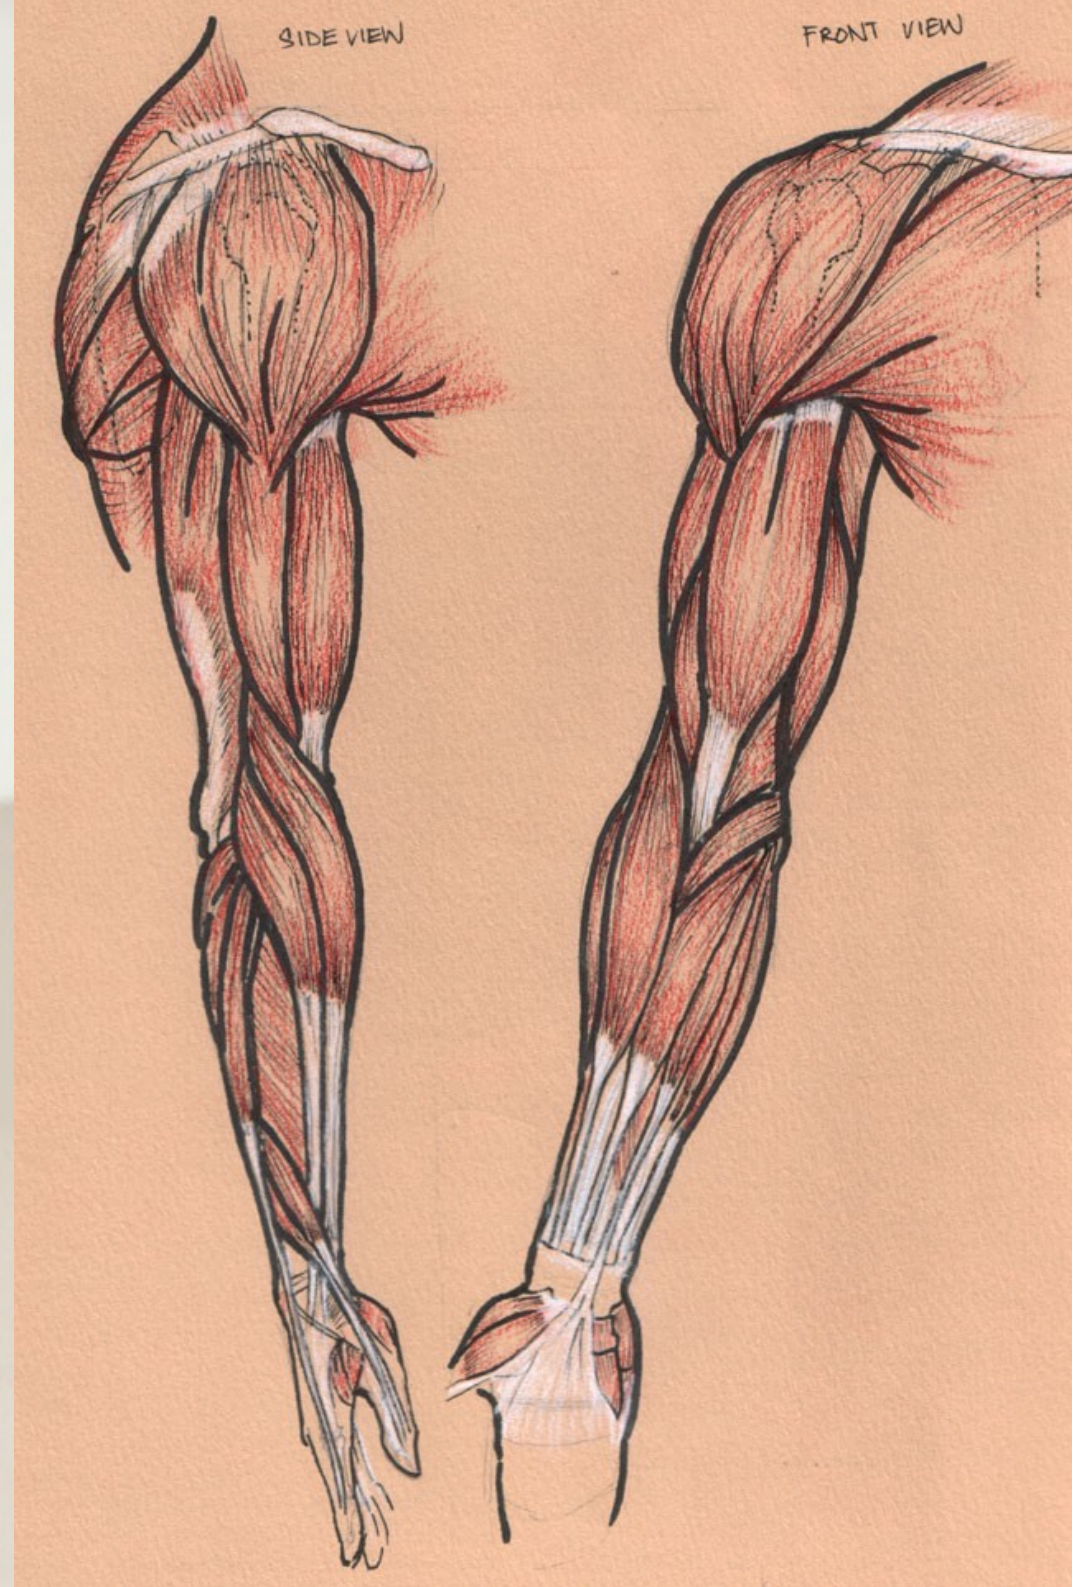
SIDE VIEW
FRONT VIEW

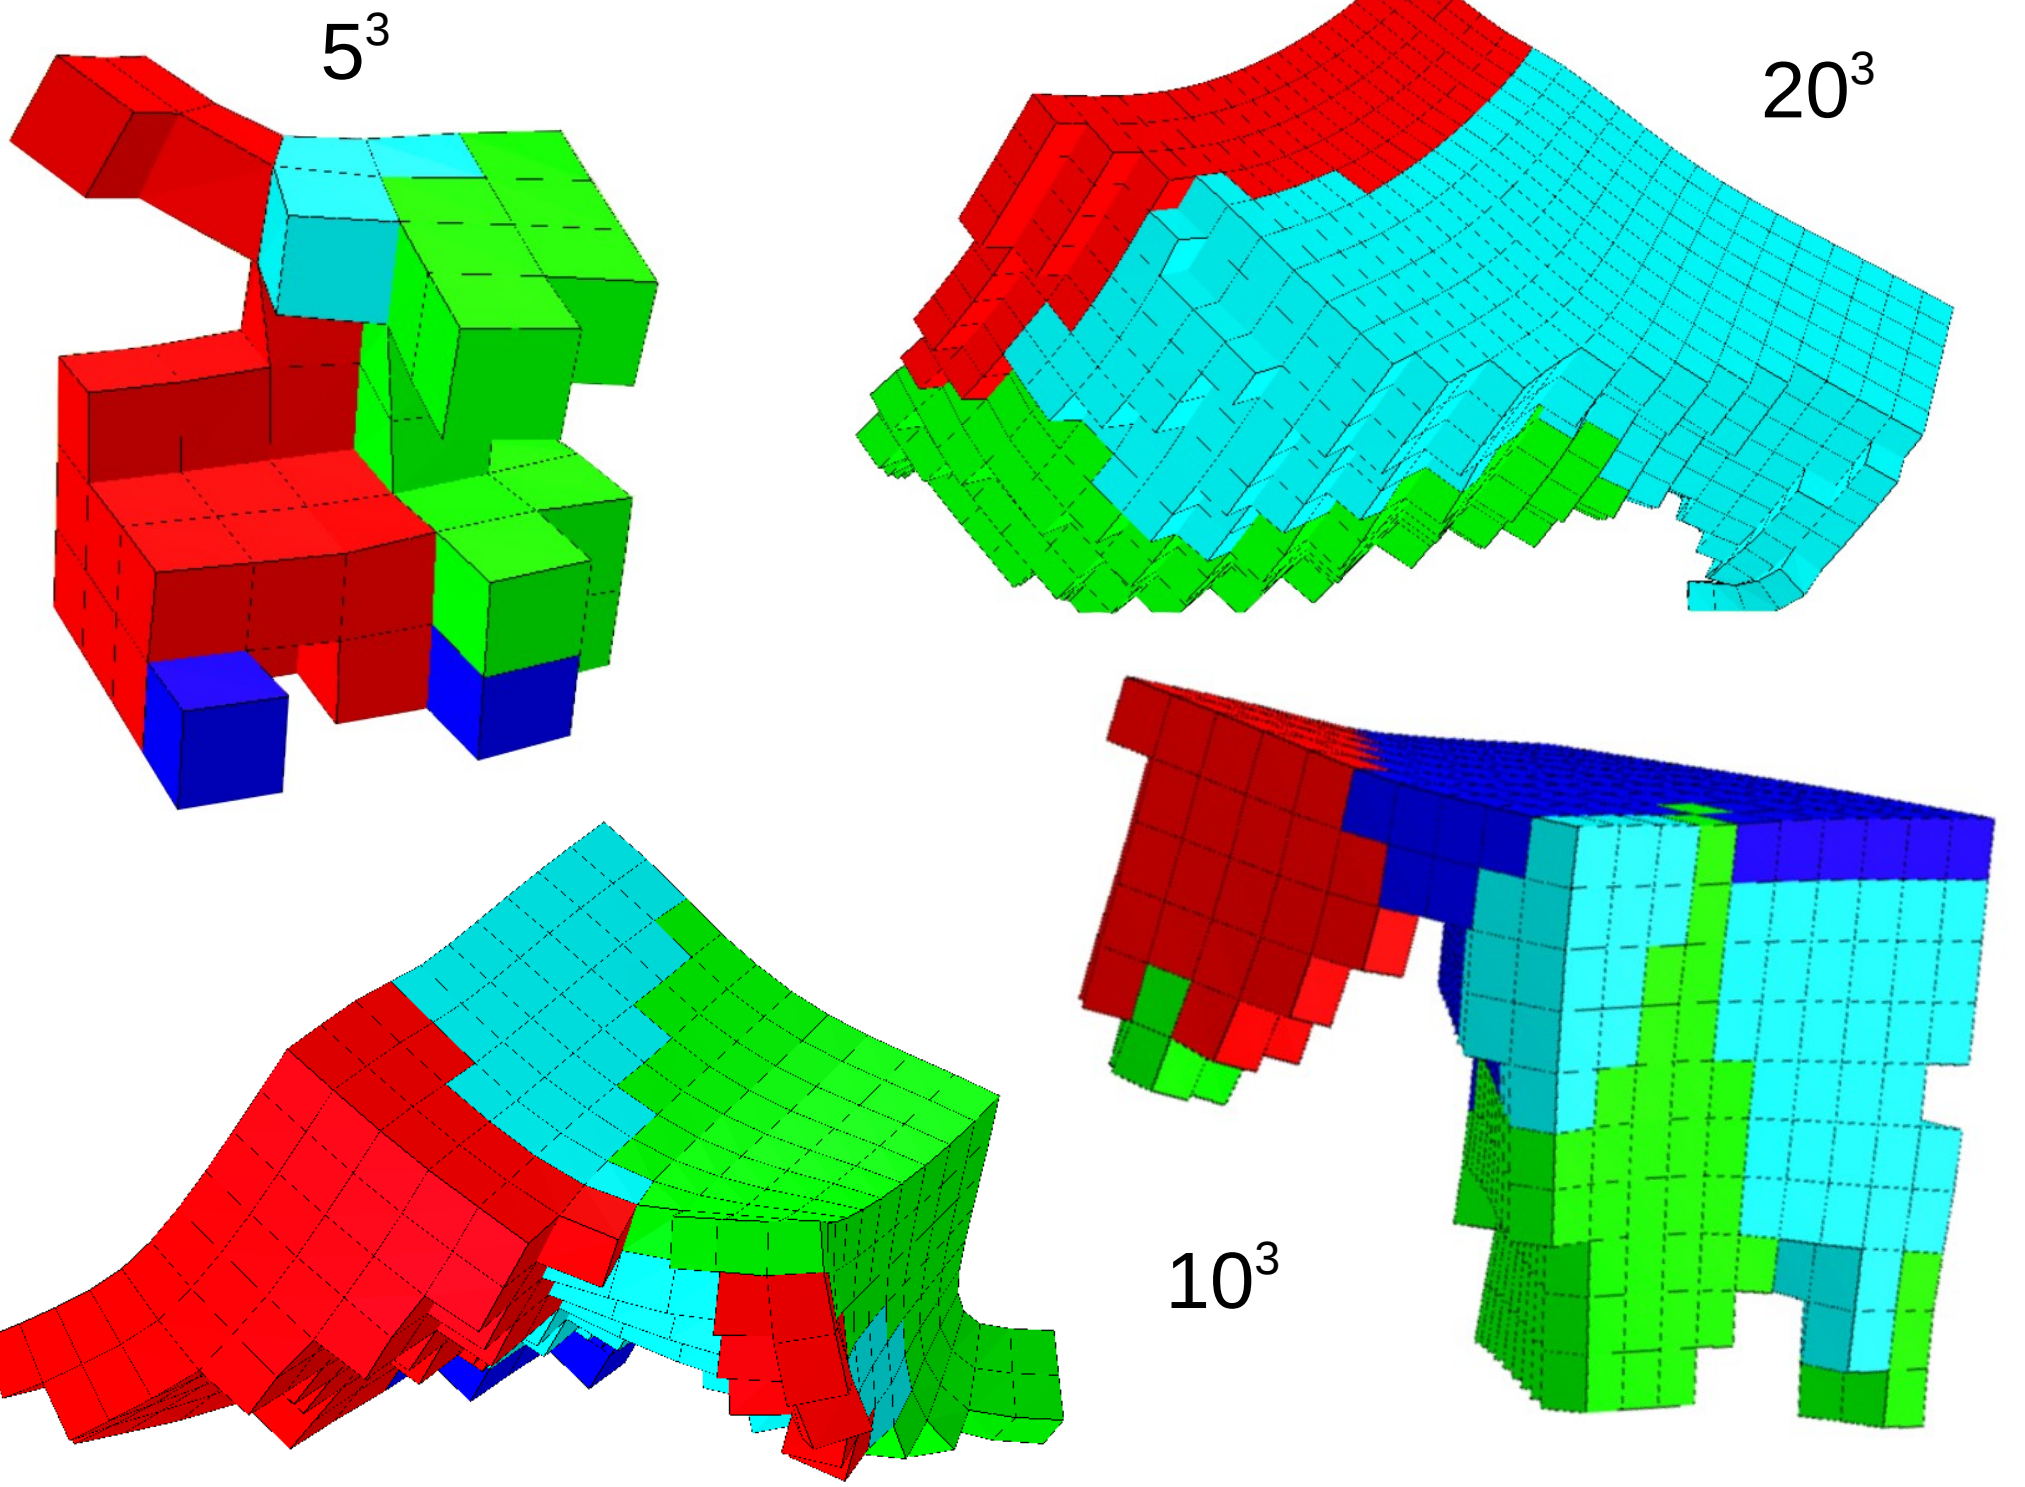
5^3
20^3
10^3

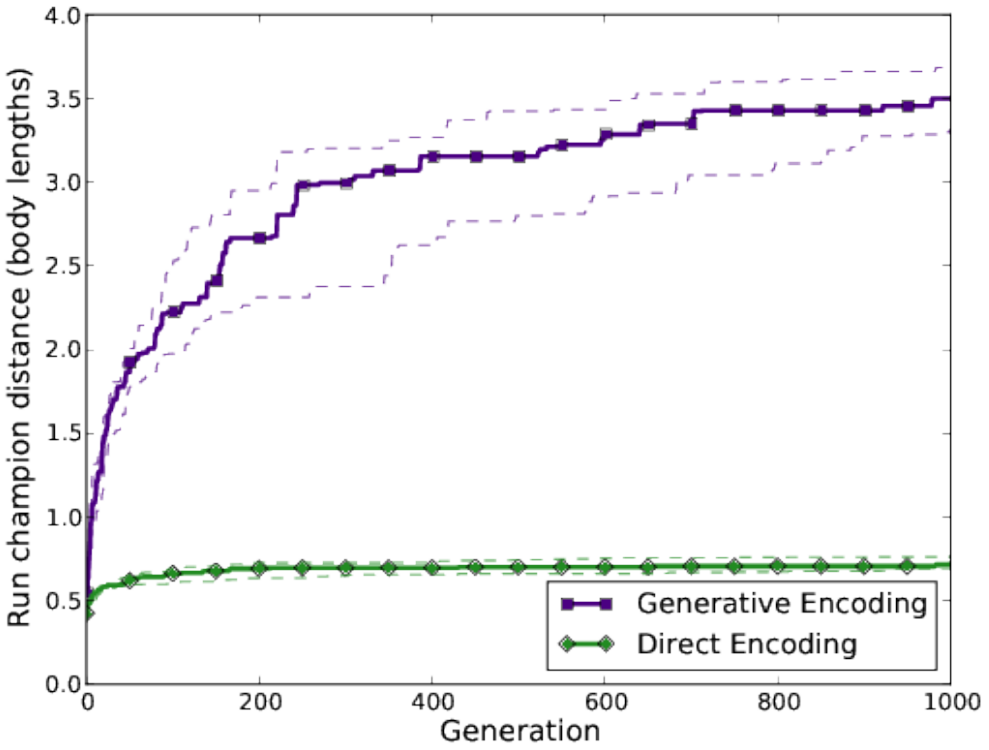
Run champion distance (body lengths)
Generation
Generative Encoding
Direct Encoding
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0
0
200
400
600
800
1000

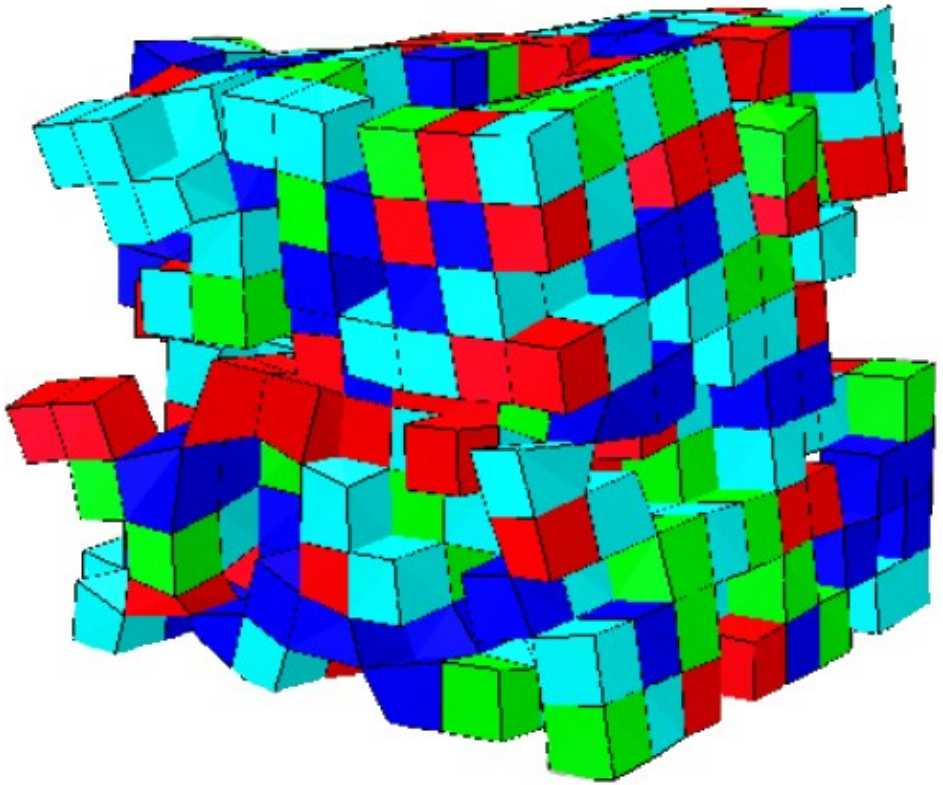

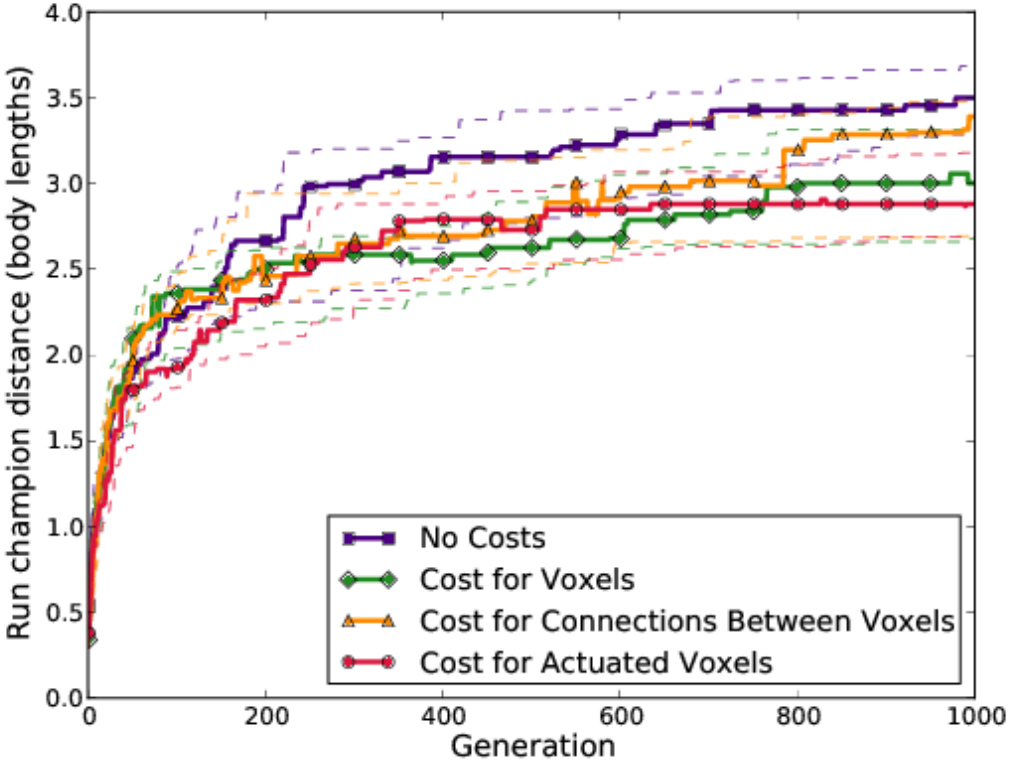
Run champion distance (body lengths)
Generation
No Costs
Cost for Voxels
Cost for Connections Between Voxels
Cost for Actuated Voxels

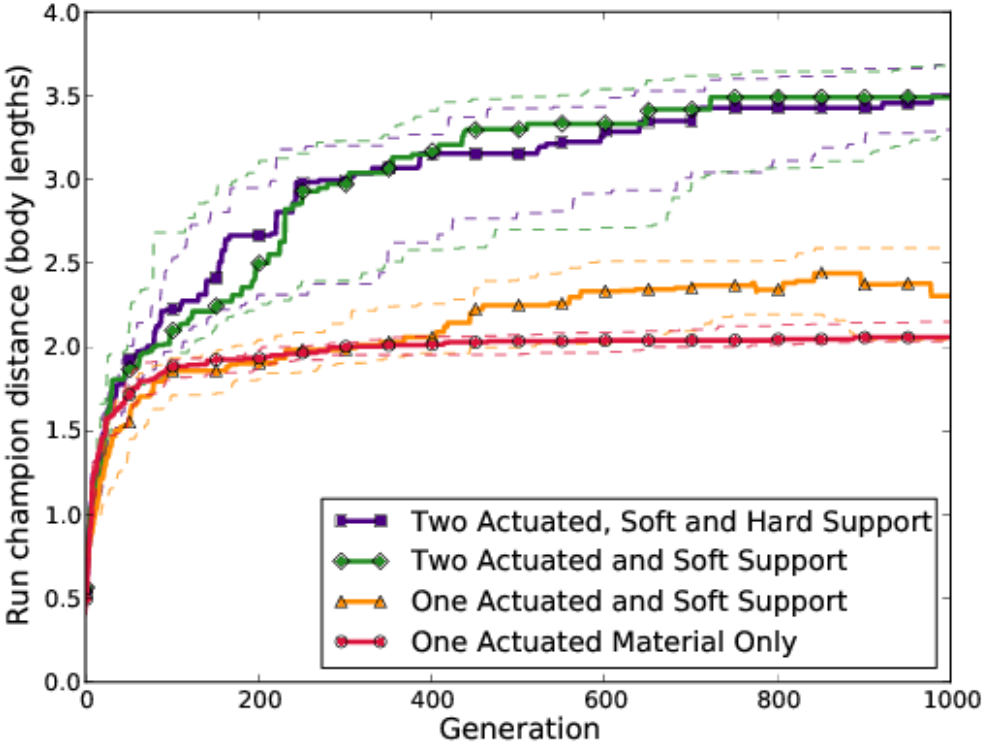
Run champion distance (body lengths)
Generation
Two Actuated, Soft and Hard Support
Two Actuated and Soft Support
One Actuated and Soft Support
One Actuated Material Only

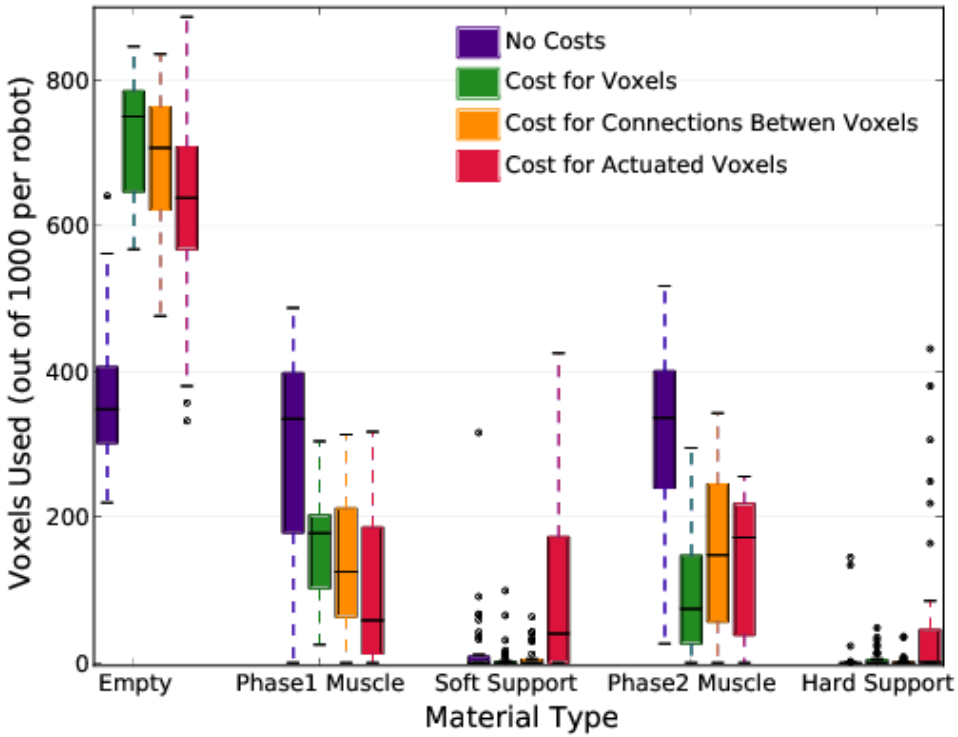
No Costs
Cost for Voxels
Cost for Connections Betwen Voxels
Cost for Actuated Voxels
Voxels Used (out of 1000 per robot)
0
200
400
600
800
Empty
Phase1 Muscle
Soft Support
Phase2 Muscle
Hard Support
Material Type

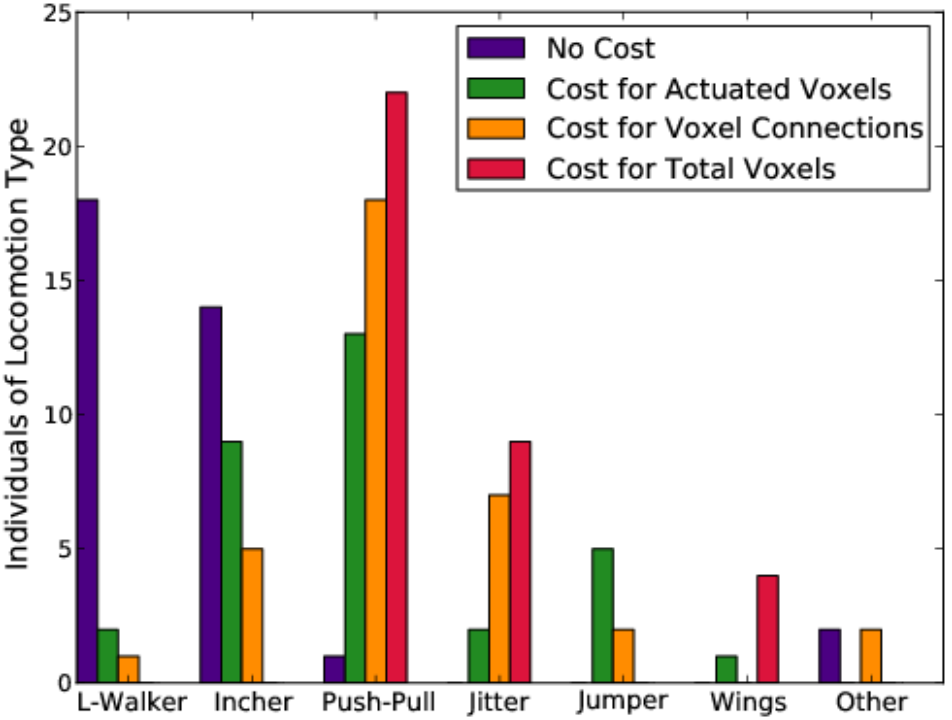
No Cost
Cost for Actuated Voxels
Cost for Voxel Connections
Cost for Total Voxels
Individuals of Locomotion Type
0
5
10
15
20
25
L-Walker
Incher
Push-Pull
Jitter
Jumper
Wings
Other

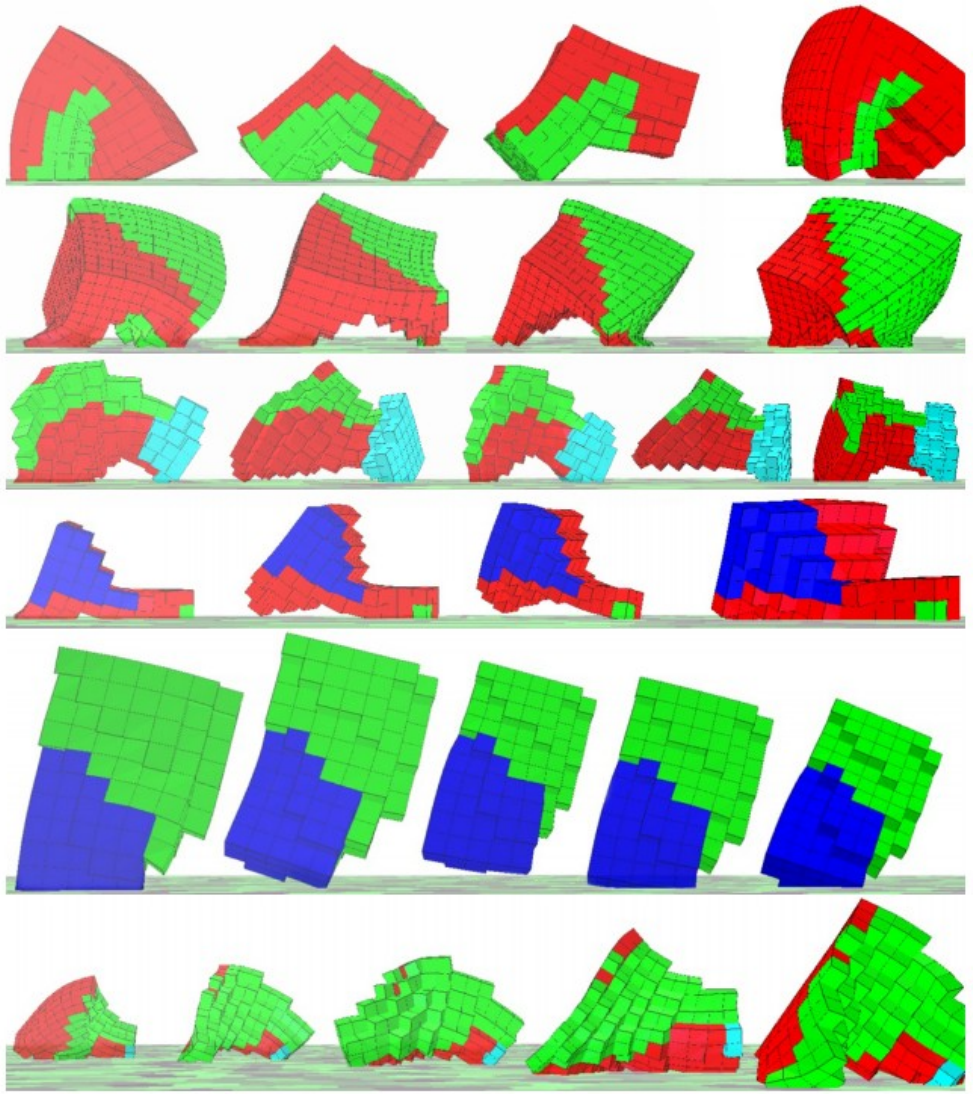

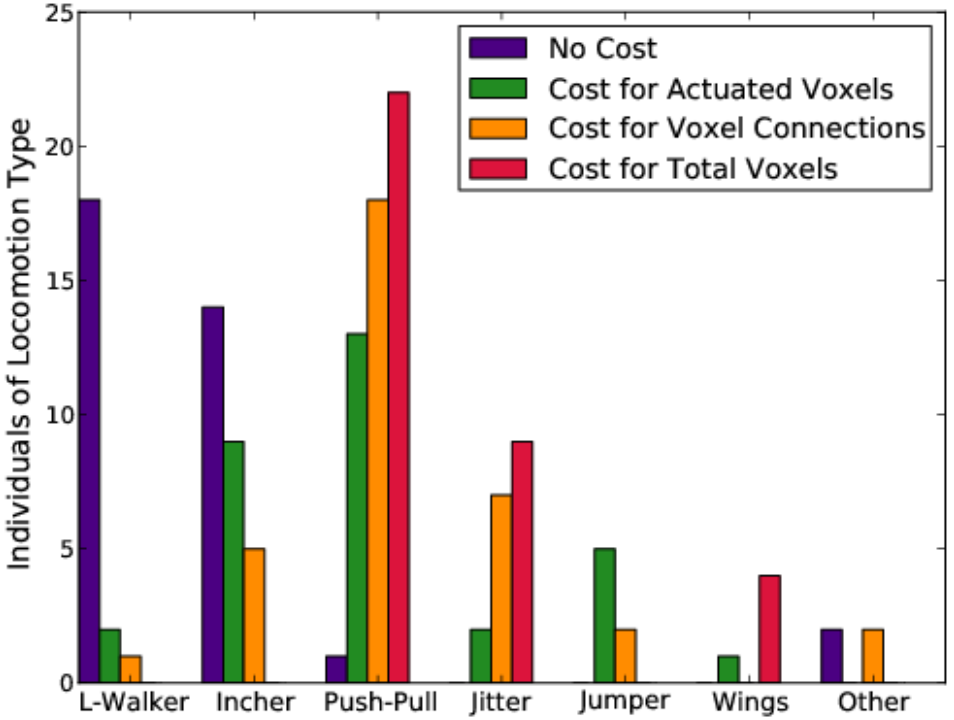

Figure 12: Time series of common soft robot behaviors as they move from left to right across the image. From top to bottom, we refer to them as L-Walker, Incher, Push-Pull, Jitter, Jumper, and Wings. Fig. 11 reports how frequently they evolved.

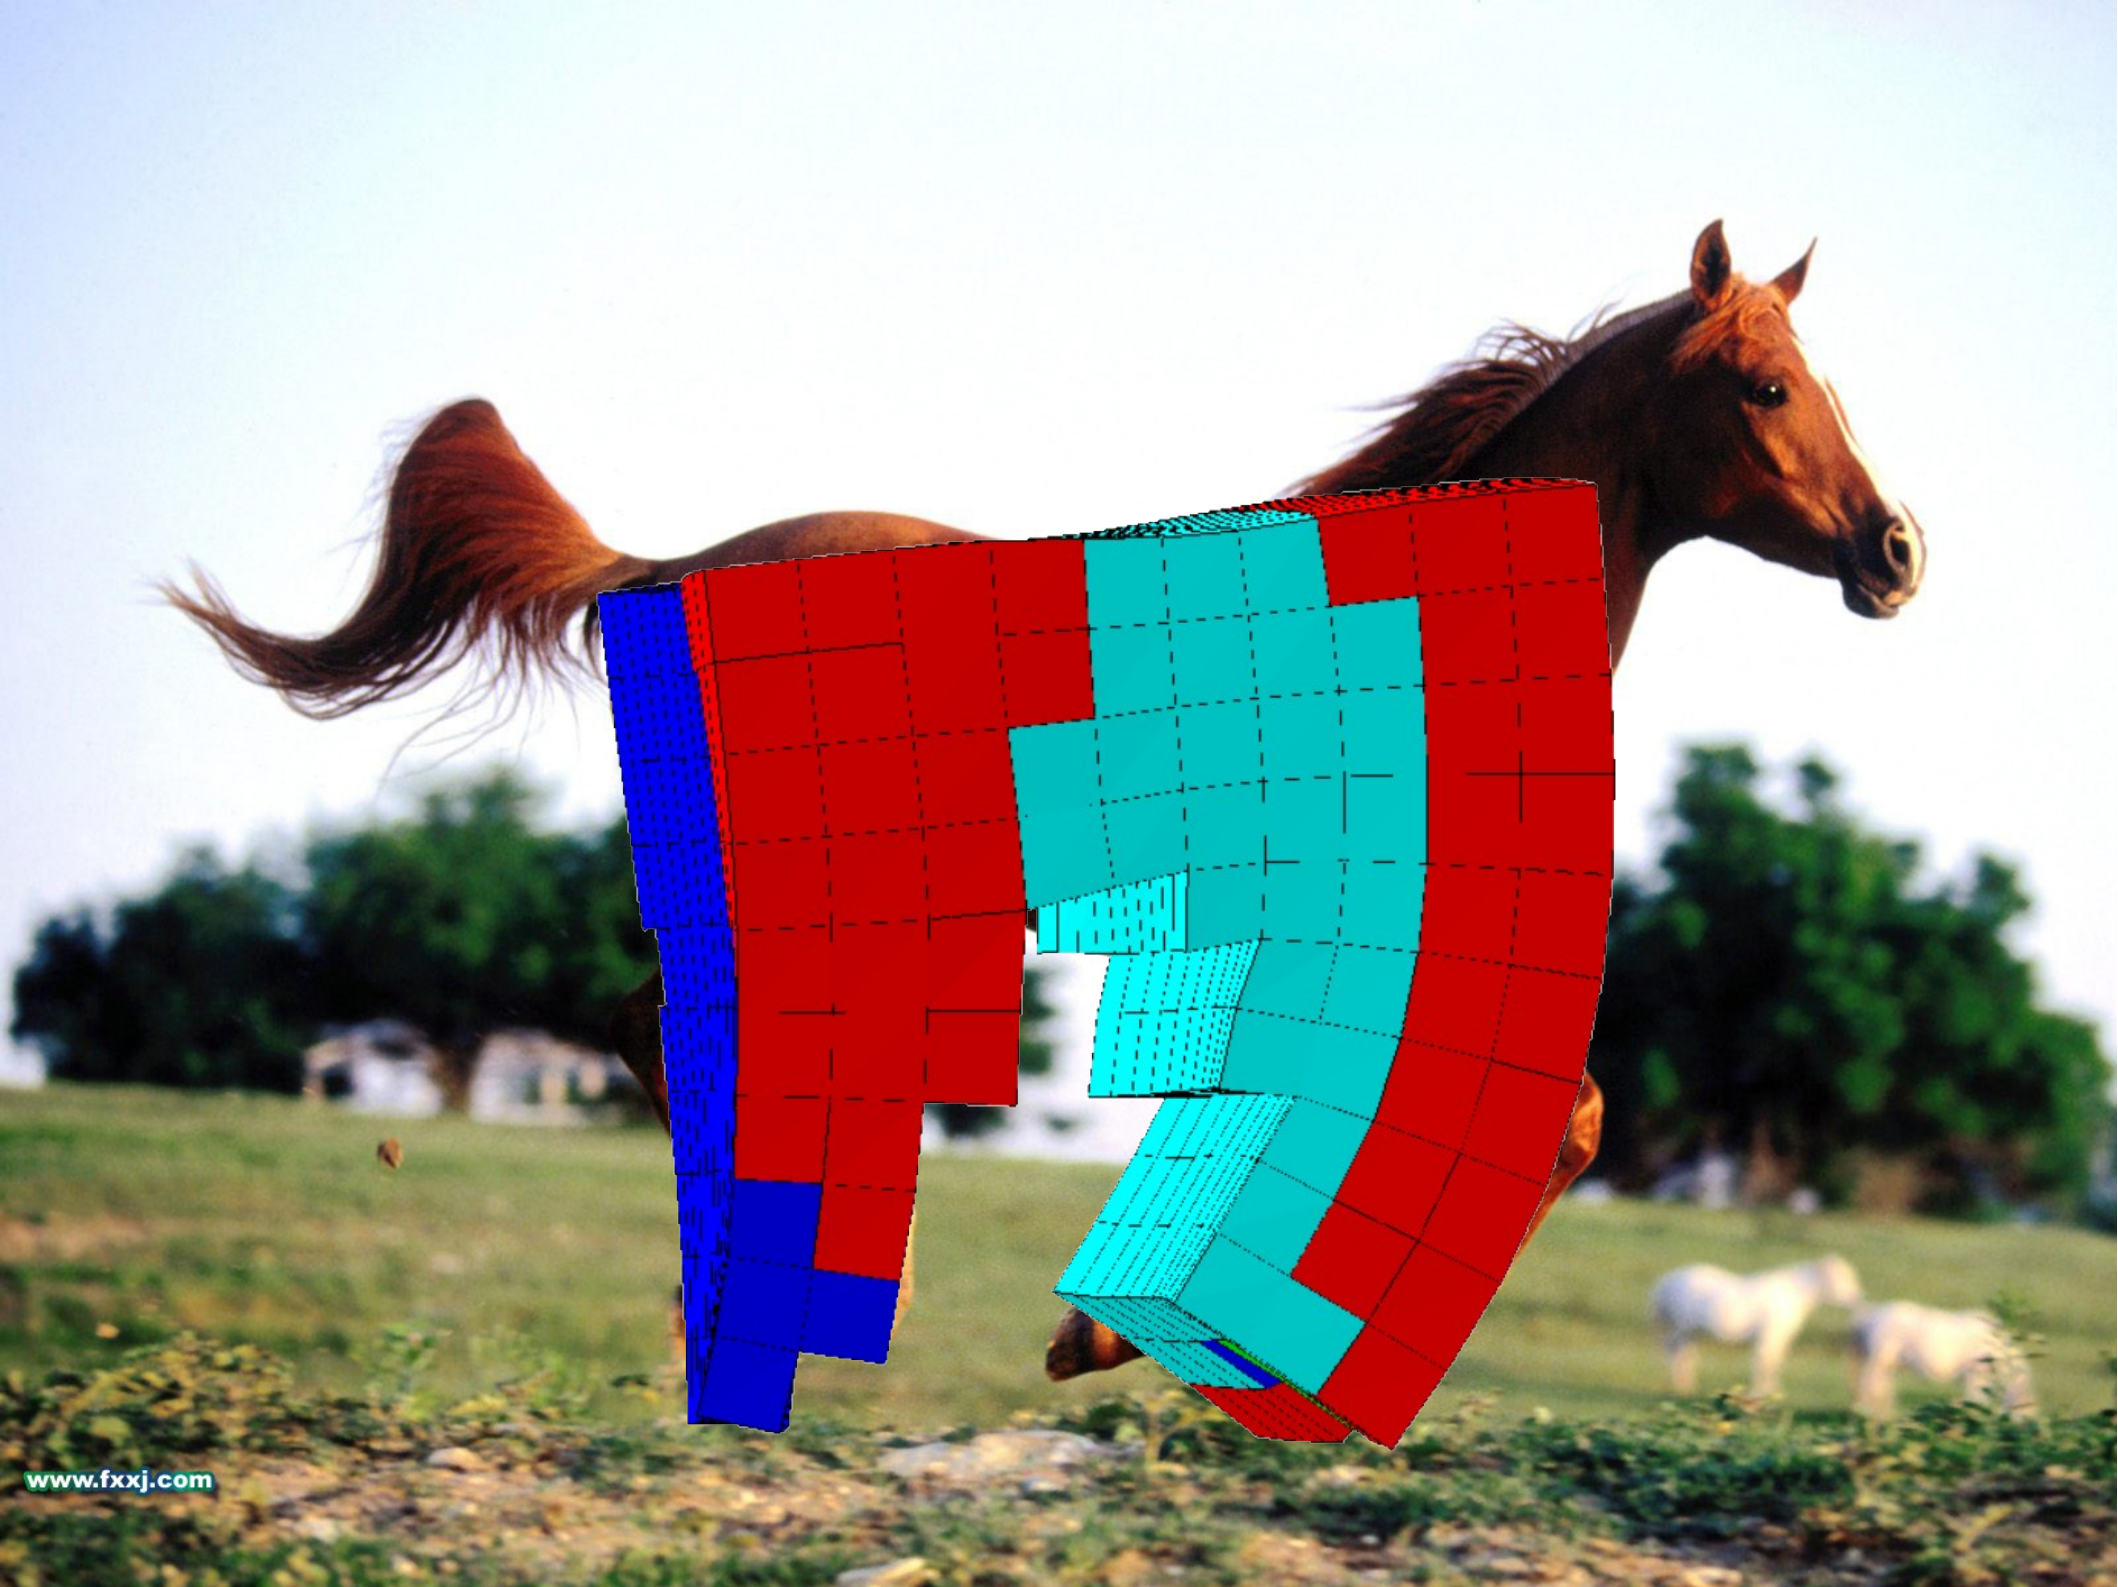
www.fxxj.com

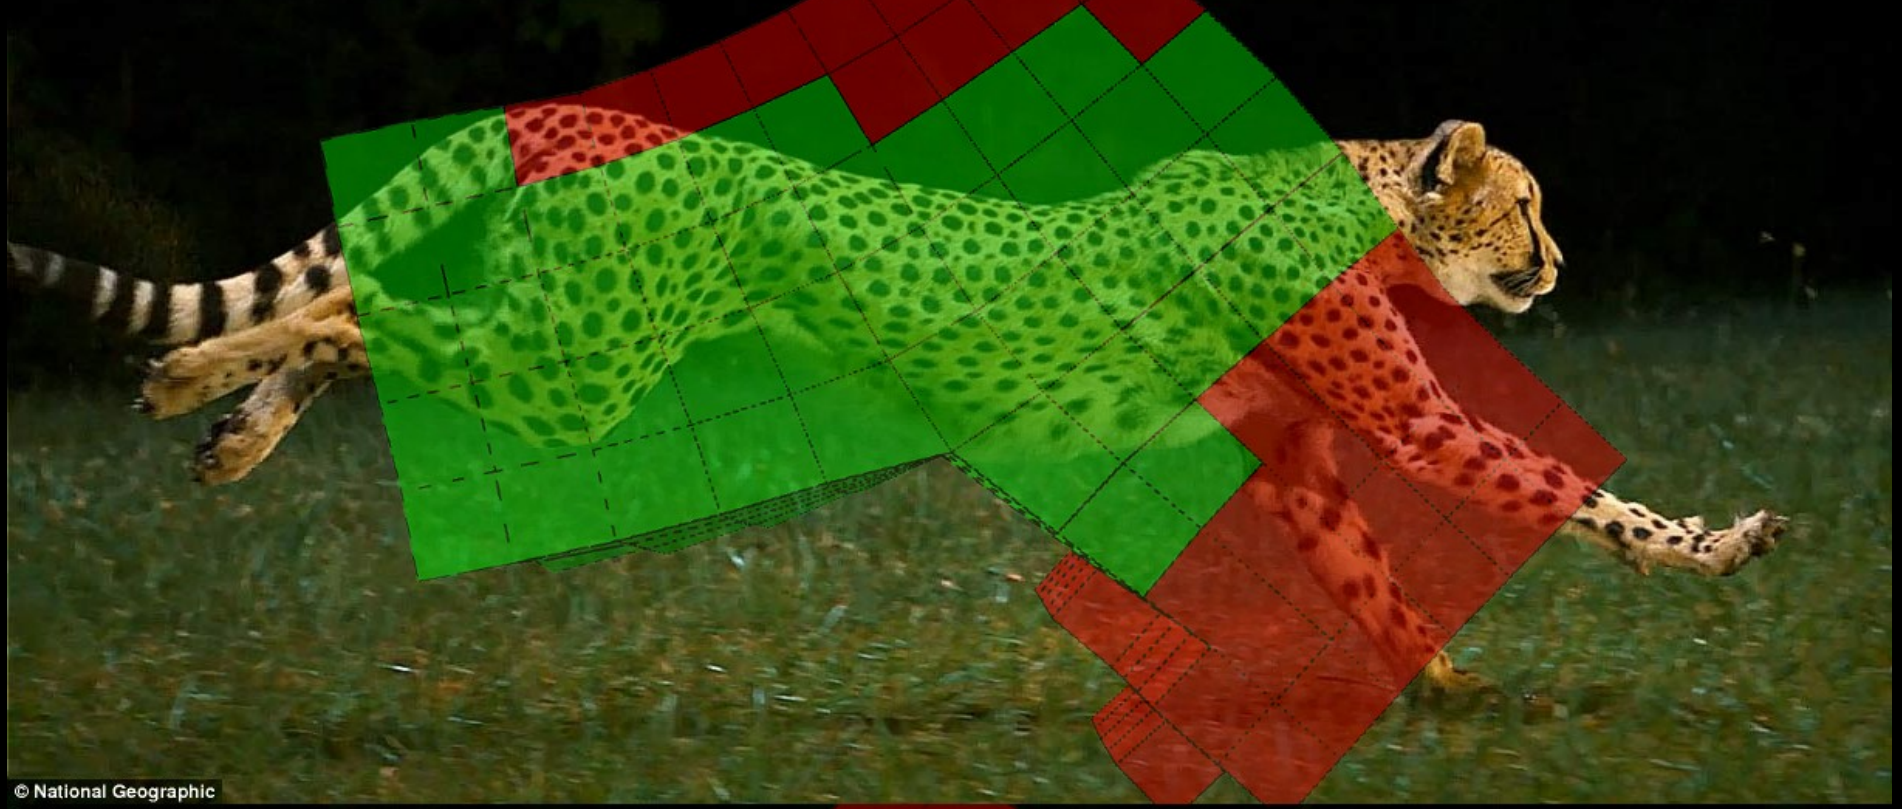
© National Geographic

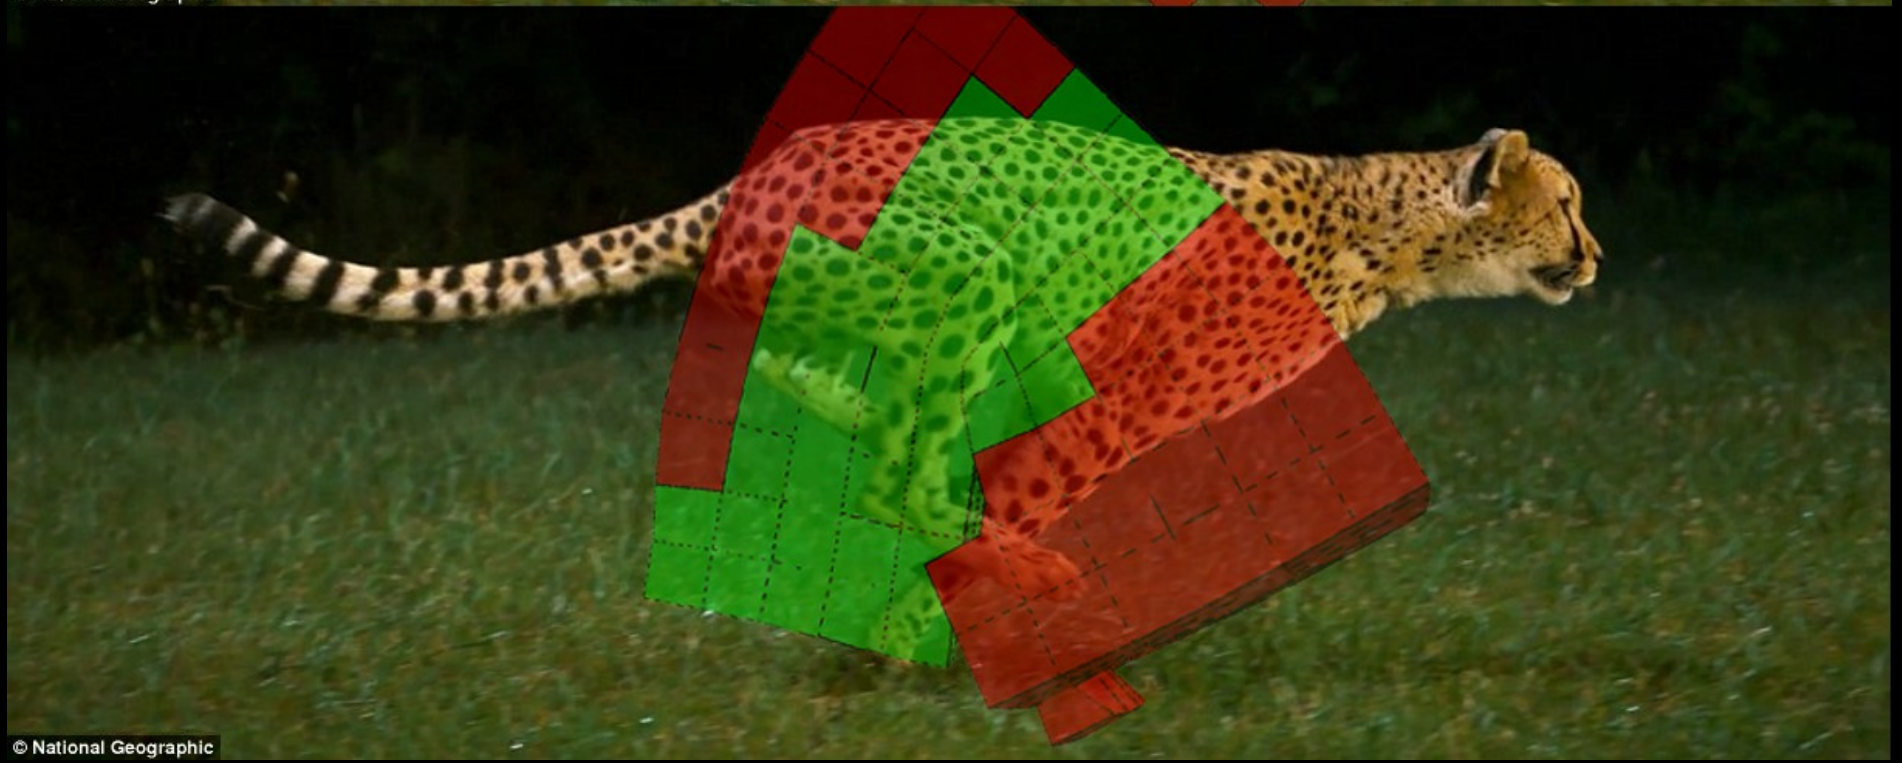
© National Geographic

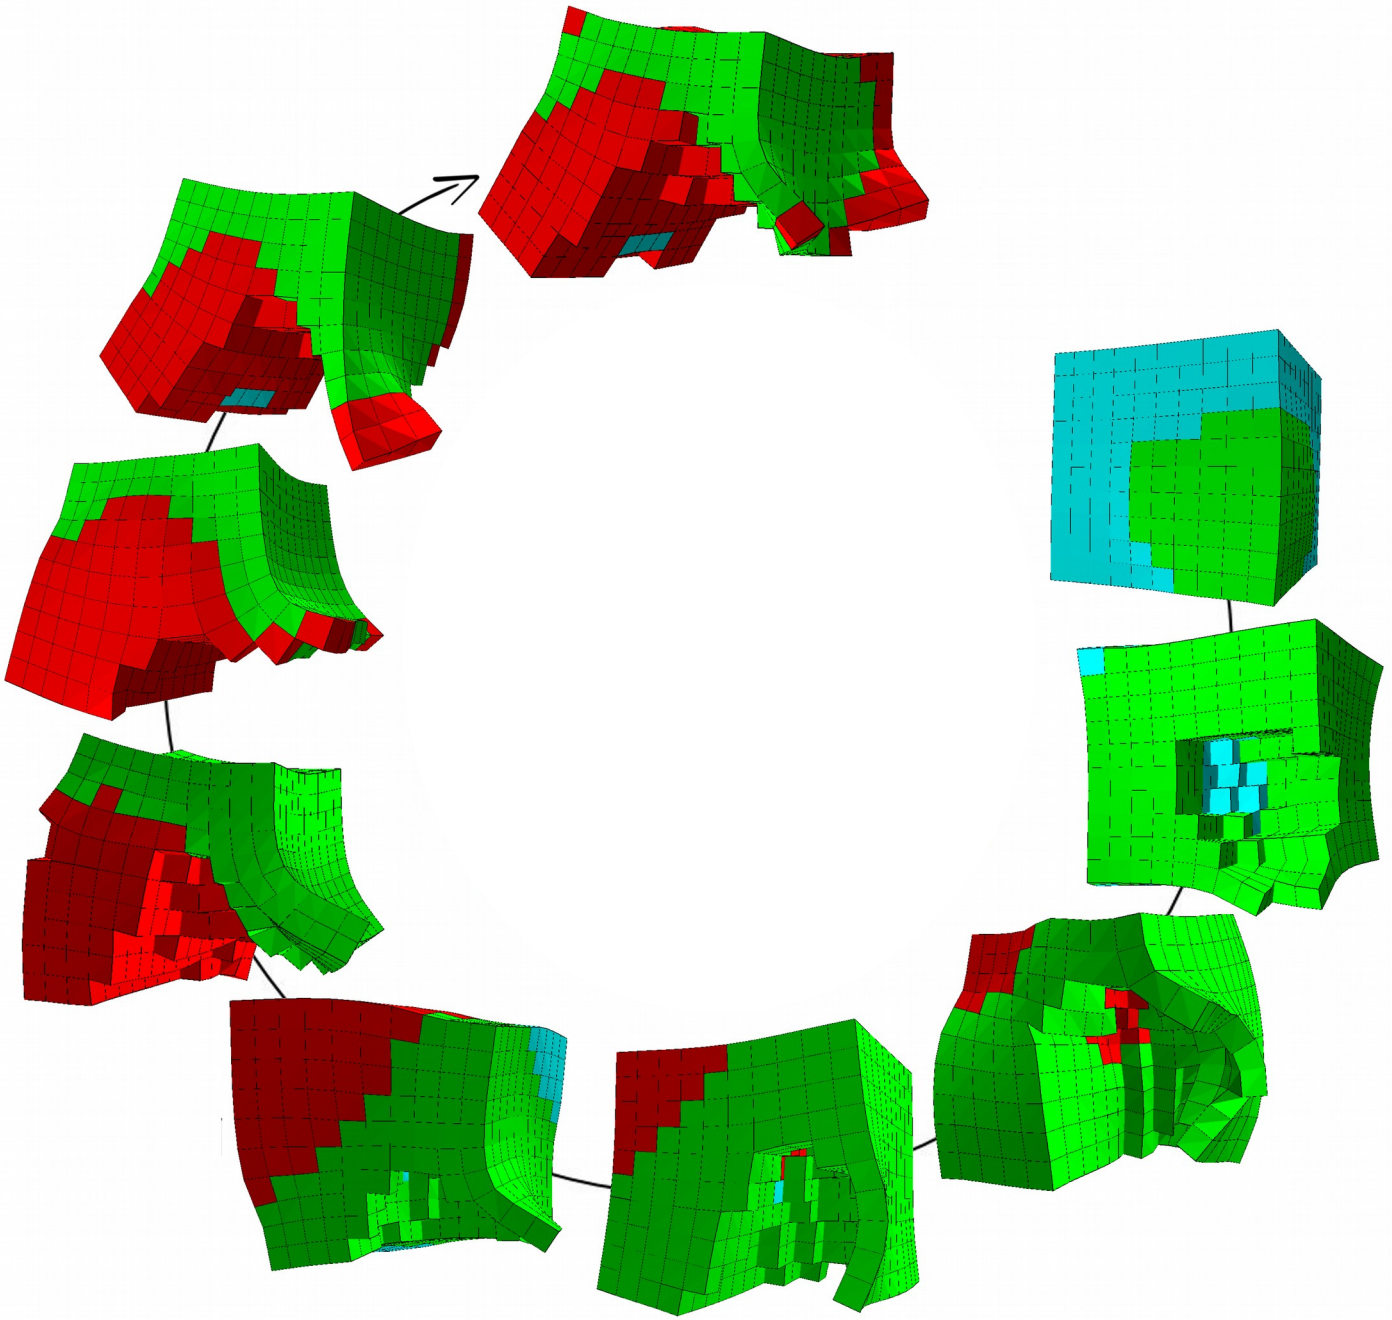

Evolving Complete Robots with CPPN-NEAT: The Utility of Recurrent Connections

Joshua E. Auerbach
Morphology, Evolution and Cognition Lab
Department of Computer Science
University of Vermont
Burlington, VT 05401
joshua.auerbach@uvm.edu

Josh C. Bongard
Morphology, Evolution and Cognition Lab
Department of Computer Science
University of Vermont
Burlington, VT 05401
jbongard@uvm.edu

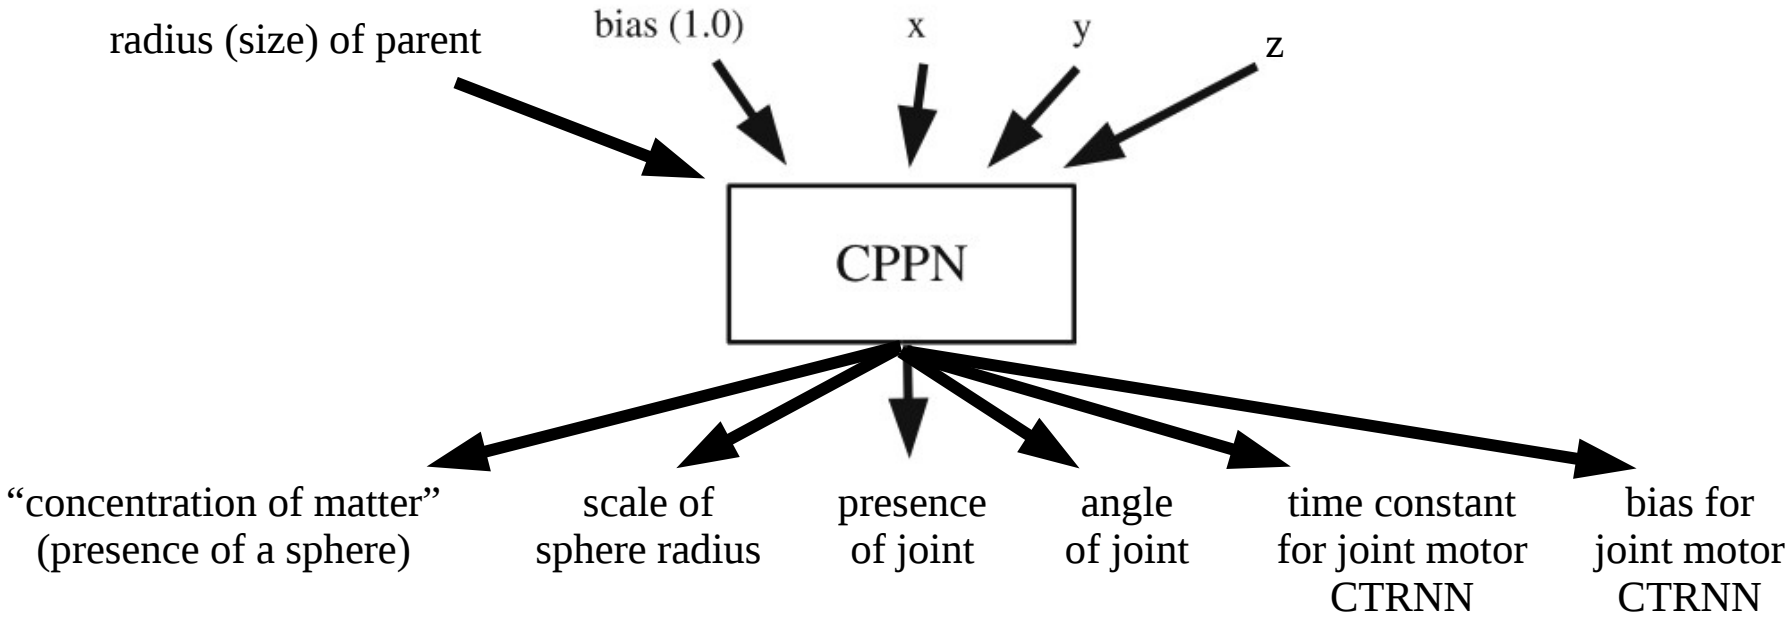
radius (size) of parent
bias (1.0)
x
y
z
CPPN
“concentration of matter” (presence of a sphere)
scale of sphere radius
presence of joint
angle of joint
time constant for joint motor CTRNN
bias for joint motor CTRNN

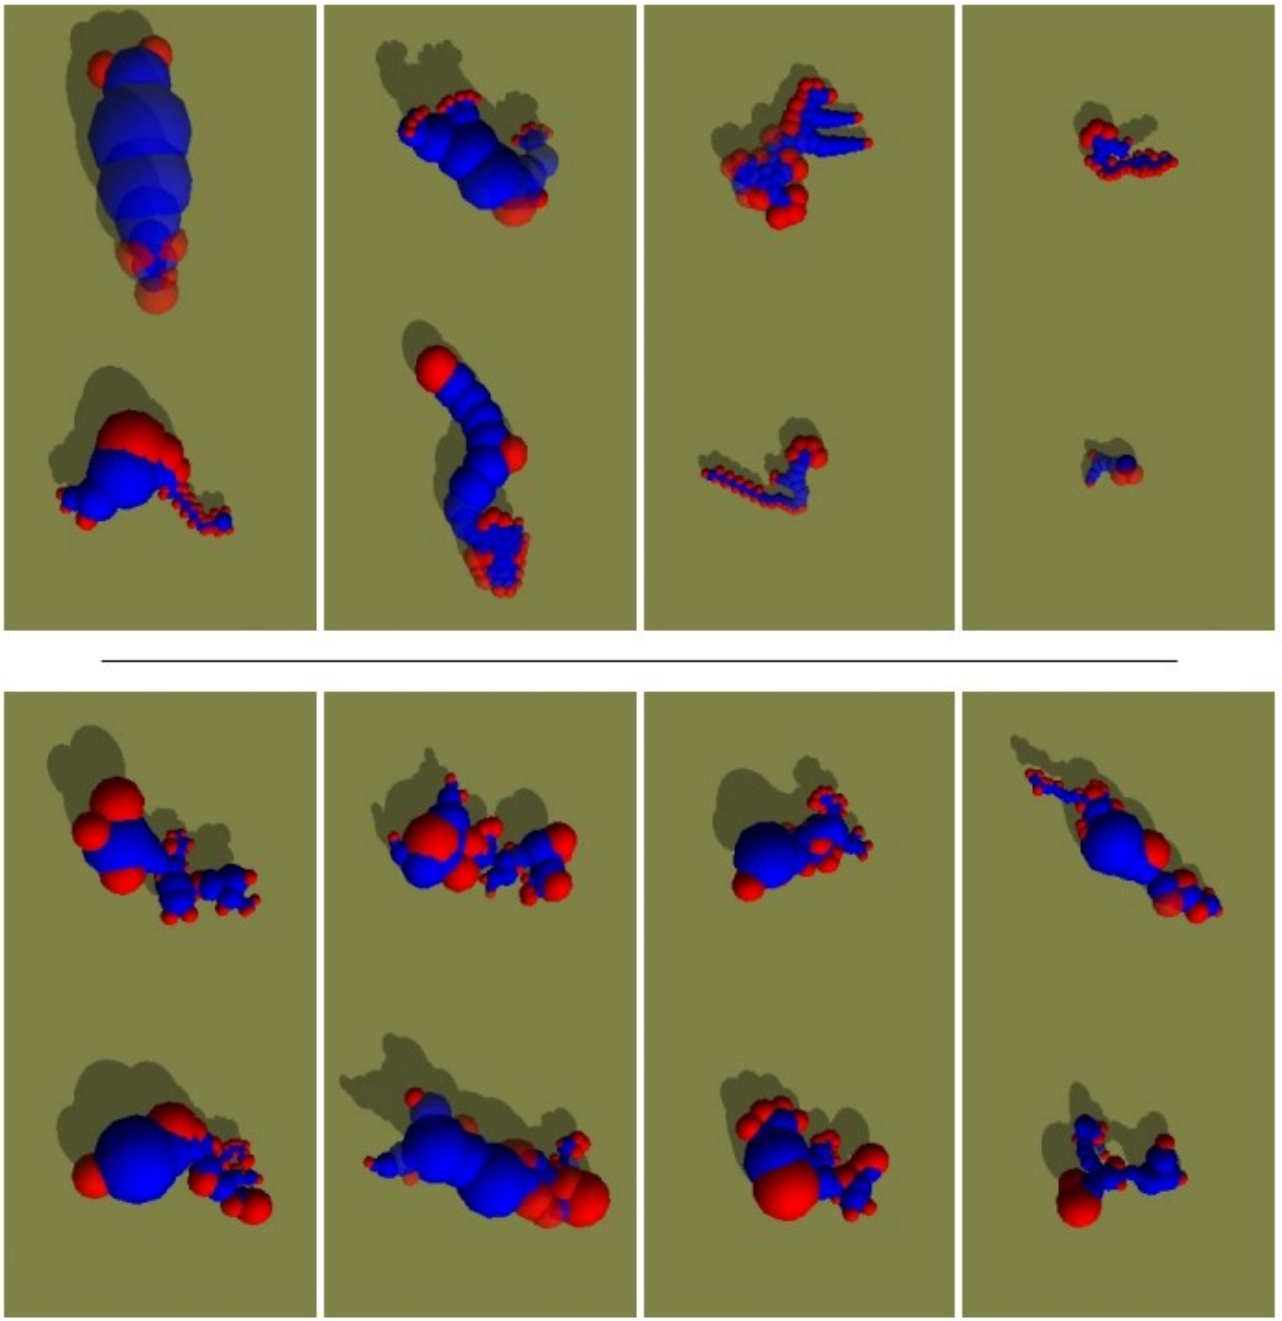

Figure 4: The Zoo: Pictures of the top eight best of run individuals from the control regime (top) and experimental regime (bottom). Leaf spheres are colored red while all other spheres are colored blue. Videos of these robots in action are available at `http://www.cs.uvm.edu/~jauerbac`

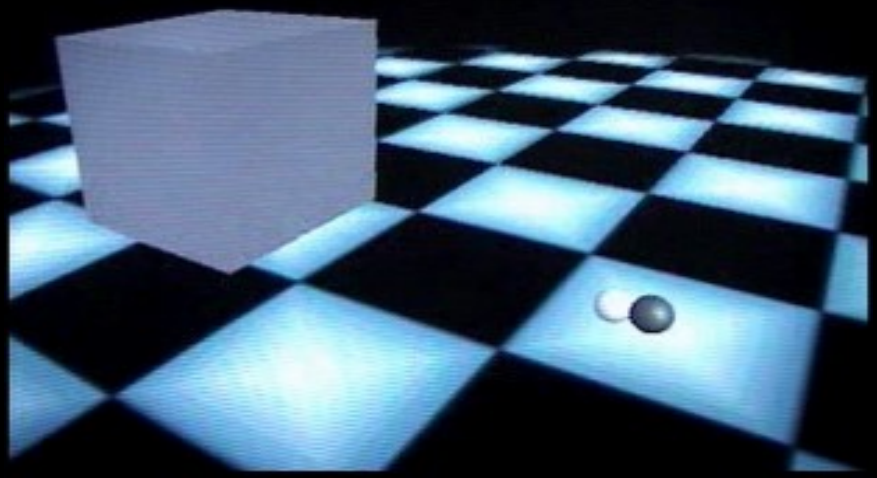

(same author, different model)

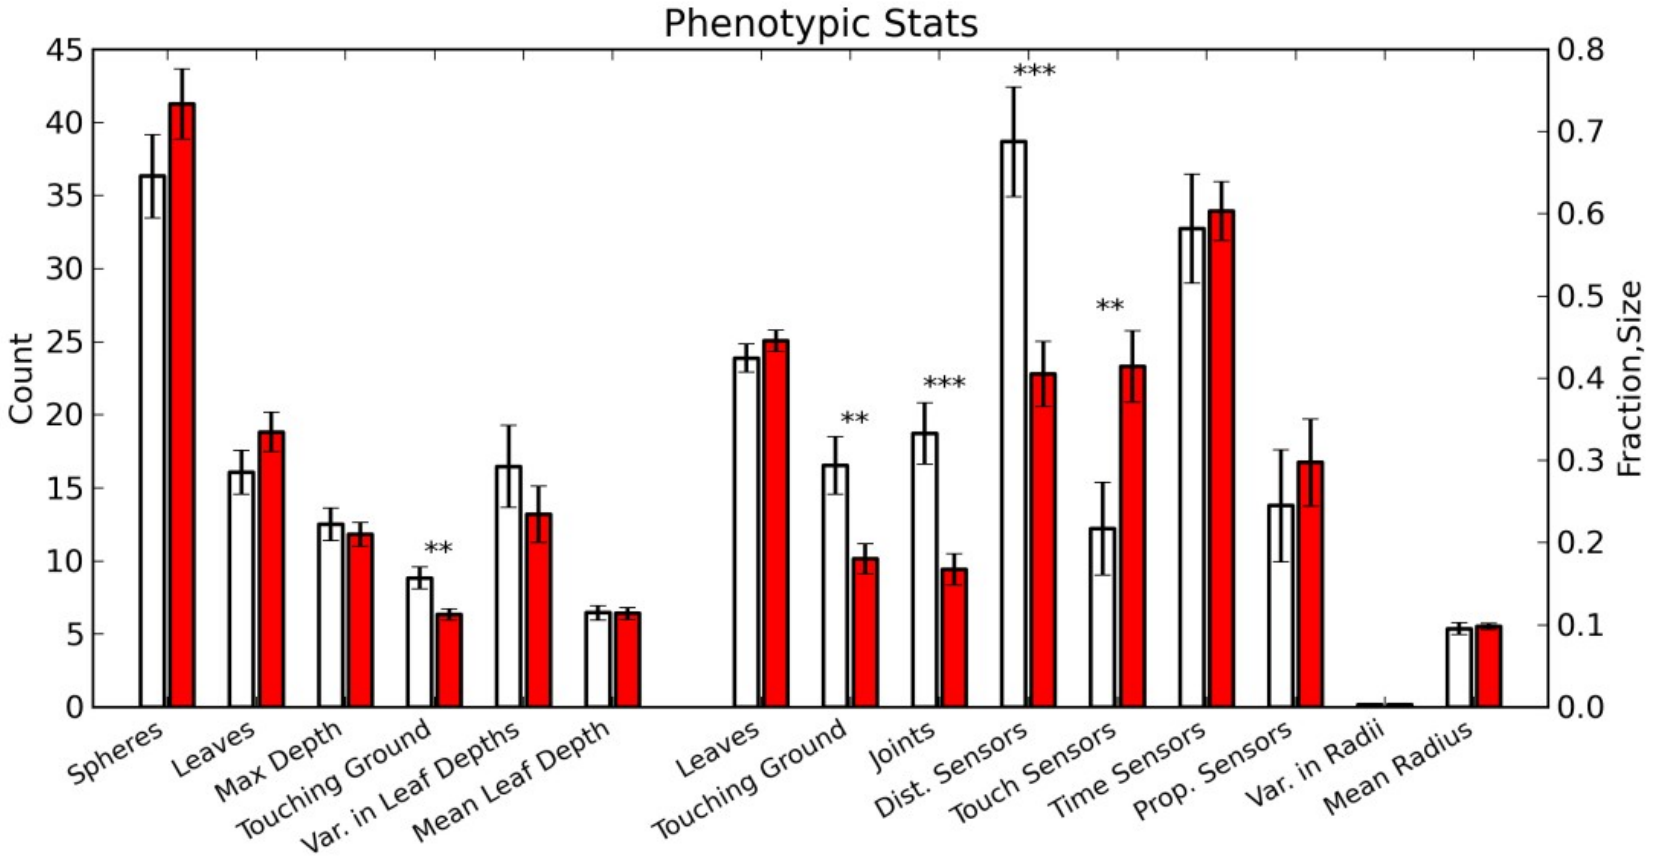

Figure 2: Comparison of several different morphological statistics between the best of run robots produced in the control regime (white) and experimental regime (red). The left hand axis is used for the leftmost six pairs while the right hand axis is used for the other pairs. Asterisks denote statistics that are significantly different between the two regimes: * denotes p-values < 0.05, ** denotes p-values < 0.01, and * denotes p-values < 0.001**

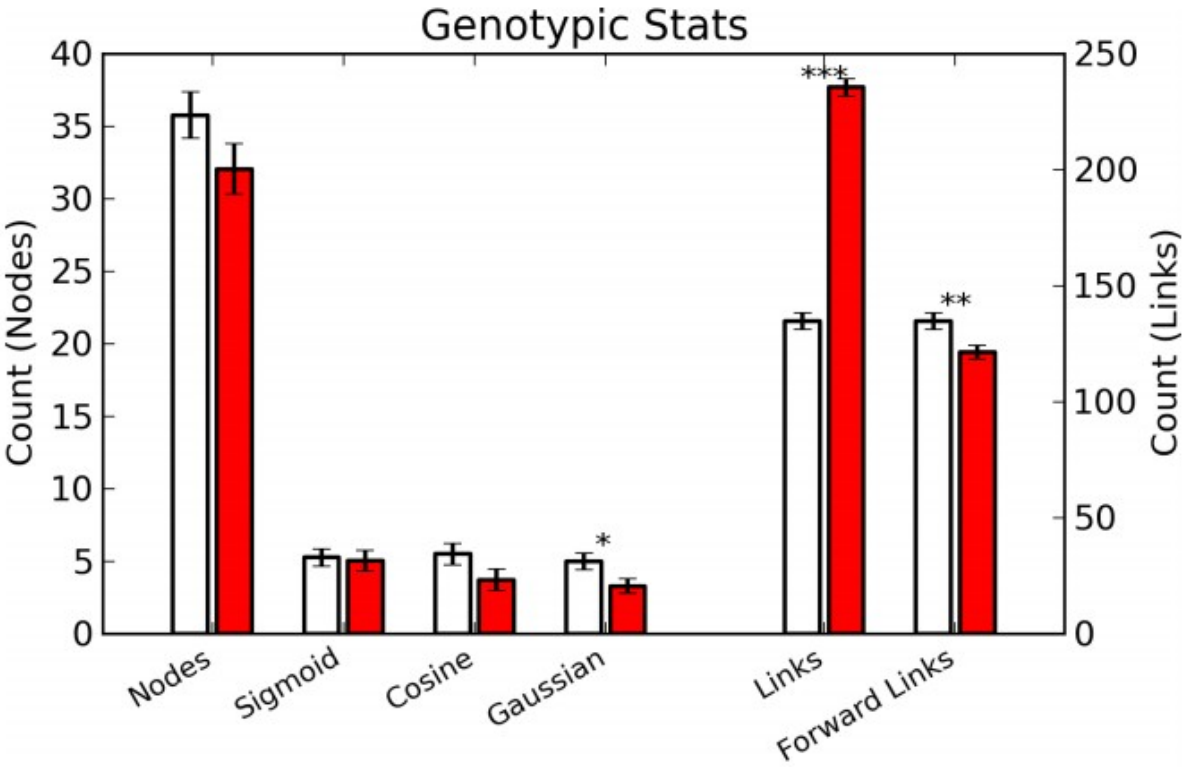

Figure 3: Comparison of genotypic statistics between the best of run CPPNs from the control regime (white) and experimental regime (red). The left hand axis is used for the number of nodes and number of hidden nodes with a given activation function. The right hand axis is used for the number of links.

Evolving CPPNs to Grow Three-Dimensional Physical Structures

Joshua E. Auerbach
Morphology, Evolution and Cognition Lab
Department of Computer Science
University of Vermont
Burlington, VT 05401
joshua.auerbach@uvm.edu

Josh C. Bongard
Morphology, Evolution and Cognition Lab
Department of Computer Science
University of Vermont
Burlington, VT 05401
jbongard@uvm.edu

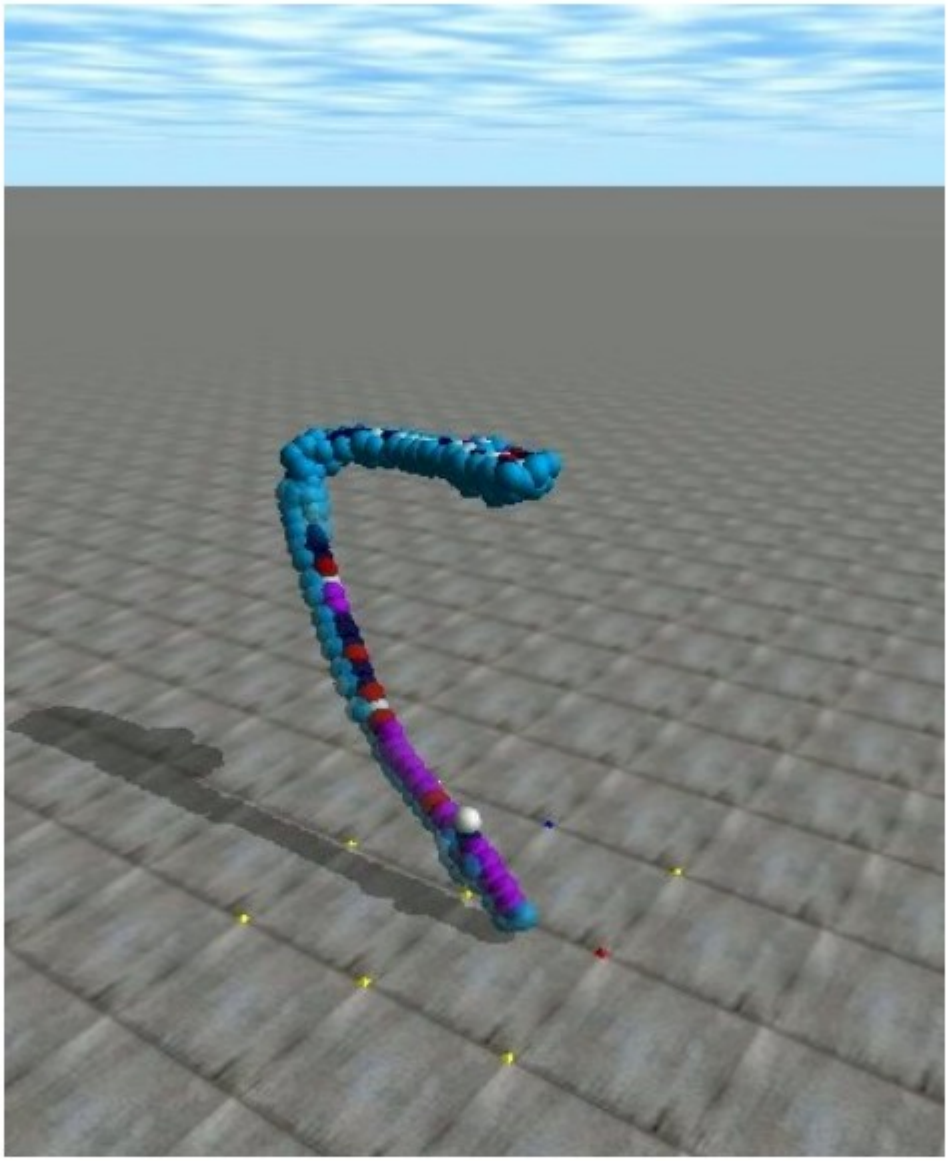

Figure 1: A sample structure evolved for maximum displacement due to gravity.

	CPPN Input Set	Resolution: first 100 generations	Resolution: second 100 generations
Experiment 1	Full	$r = 0.1\text{m}$, $M = 200$	$r = 0.1\text{m}$, $M = 200$
Experiment 2	Restricted	$r = 0.1\text{m}$, $M = 200$	$r = 0.1\text{m}$, $M = 200$
Experiment 3	Full	$r = 0.15\text{m}$, $M = 60$	$r = 0.1\text{m}$, $M = 200$
Experiment 4	Restricted	$r = 0.15\text{m}$, $M = 60$	$r = 0.1\text{m}$, $M = 200$
Experiment 5	Full	$r = 0.08\text{m}$, $M = 391$	$r = 0.08\text{m}$, $M = 391$
Experiment 6	Restricted	$r = 0.08\text{m}$, $M = 391$	$r = 0.08\text{m}$, $M = 391$
Experiment 7	Full	$r = 0.1\text{m}$, $M = 200$	$r = 0.08\text{m}$, $M = 391$
Experiment 8	Restricted	$r = 0.1\text{m}$, $M = 200$	$r = 0.08\text{m}$, $M = 391$

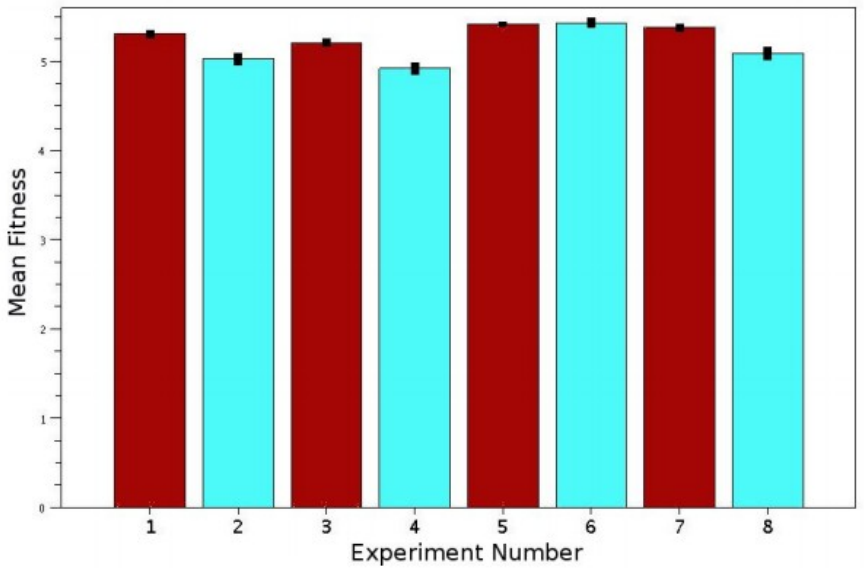

Figure 3: Mean best fitnesses (displacement in meters) in final generation across the 30 independent evolutionary trials with standard error bars for each of the eight experiments. The red bars represent experiments using the full set of CPPN inputs, while the blue bars represent experiments using the restricted set.

Dynamic Resolution in the Co-Evolution of Morphology and Control

Joshua E. Auerbach[1] and Josh C. Bongard[1]

[1]Morphology, Evolution and Cognition Laboratory
Department of Computer Science
University of Vermont
Burlington, VT 05405
joshua.auerbach@uvm.edu

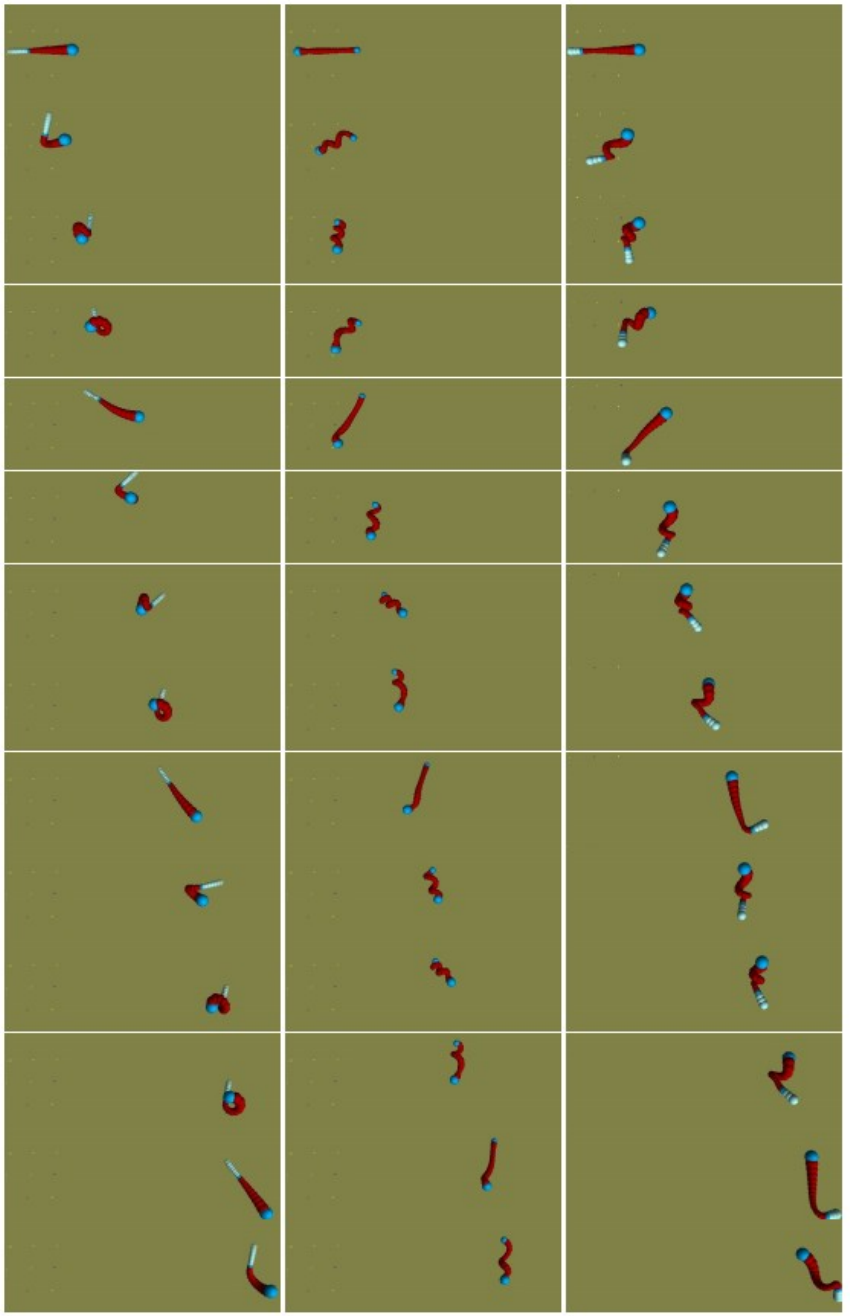

Figure 3: Each column shows the behavior of a different dynamic resolution robot evolved for directed locomotion (with time going from top to bottom). Three different robots are shown. Red cells are attached to two joints while the darker blue cells attach to a single joint. The lighter blue cells all connect rigidly. Enlarged pictures of each of these robots are shown in Fig. 1.