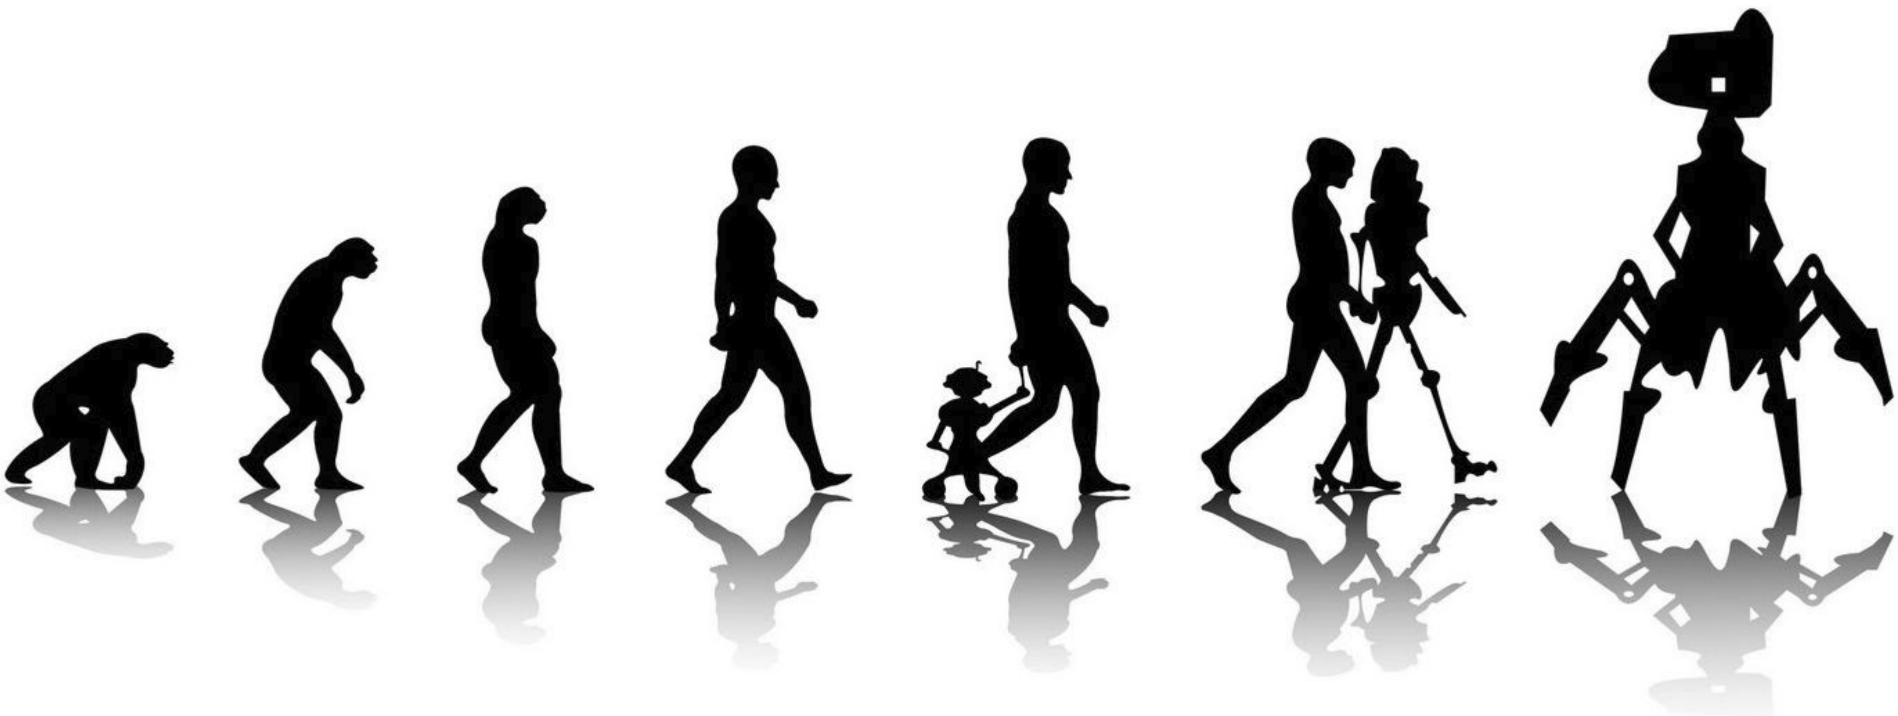

Modern Robotics: Evolutionary Robotics
COSC 4560 / COSC 5560

Professor Cheney
1/24/18

Announcements

TA Office Hours
(coding and assignment help):

Tuesday/Thursday 2-4pm

HW Assignment 1 posted!

University Calendar Revised:

- Jan. 31 – Last day to add/drop classes
  (one week from today!)
- Mar. 21 – Mid-semester grades
- Apr. 13 – Last day to withdraw
- May 11 – Our class' final (10:15-12:15)
- May 17 – Spring 2018 final grades

Meet and Greet

What's your name?

Where are you from?

Why are you here?
(What's your background/major?
What are you hoping to learn?)

One fun fact about yourself

Why Evolutionary Robotics?

Why Robotics?

Robotics Markets Are Growing Even Faster Than Expected

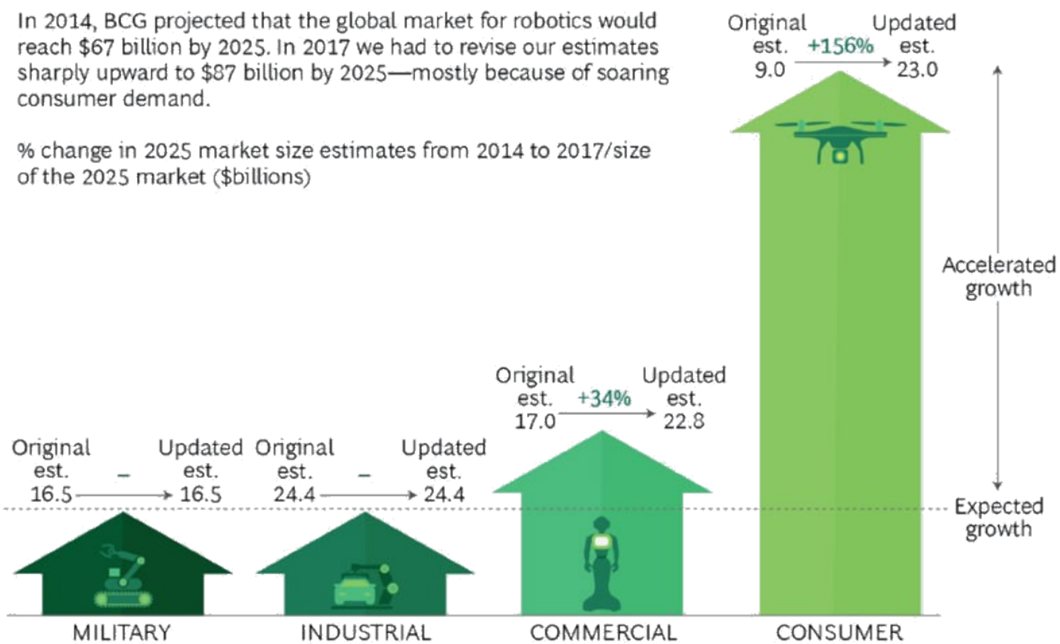

Source: BCG analysis

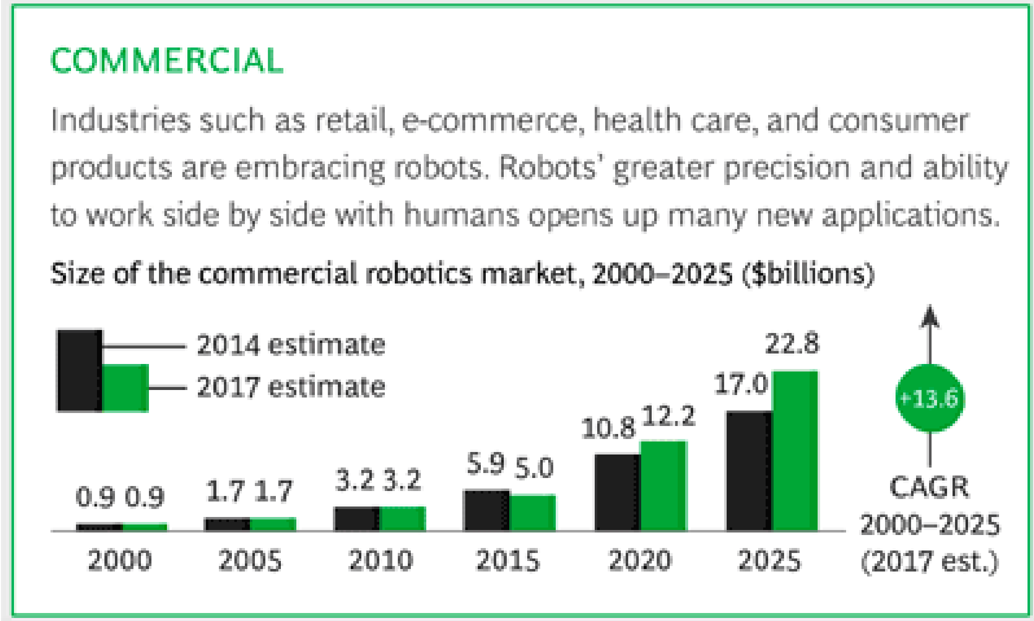

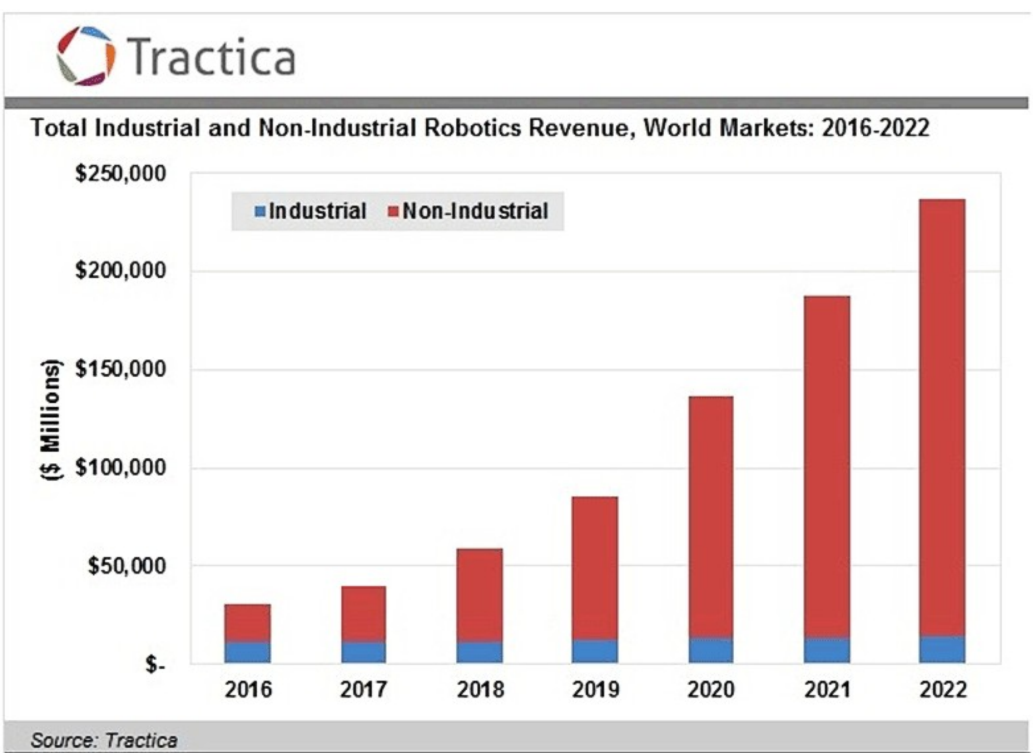

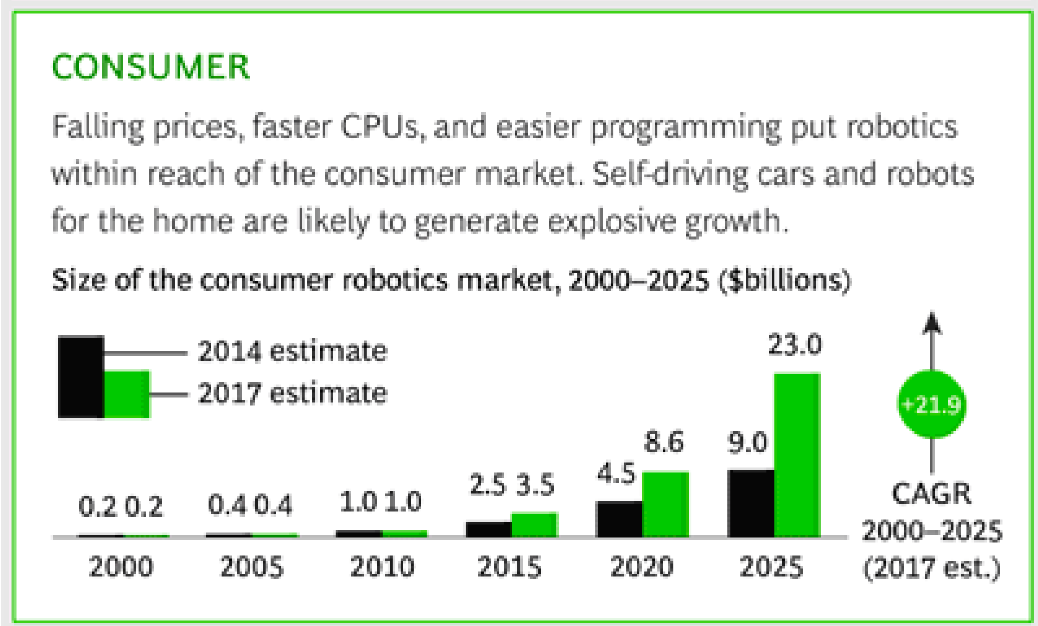

Source: Boston Dynamics Group

TECH • ROBOTICS

The Multi-Billion Dollar Robotics Market Is About to Boom

A Fanuc R-2000iC robot spot-welding the body of a truck. Photo: Courtesy of Fanuc

By JONATHAN VANIAN February 24, 2016

Here's another sign that the robotics industry is poised to see big gains in the future.

On Wednesday, a report by International Data Corporation said worldwide spending on robotics and related services will hit $135.4 billion in 2019. The research firm said that global robotics spending in 2015 was $71 billion, and is set to grow at a compound annual growth rate of 17%.

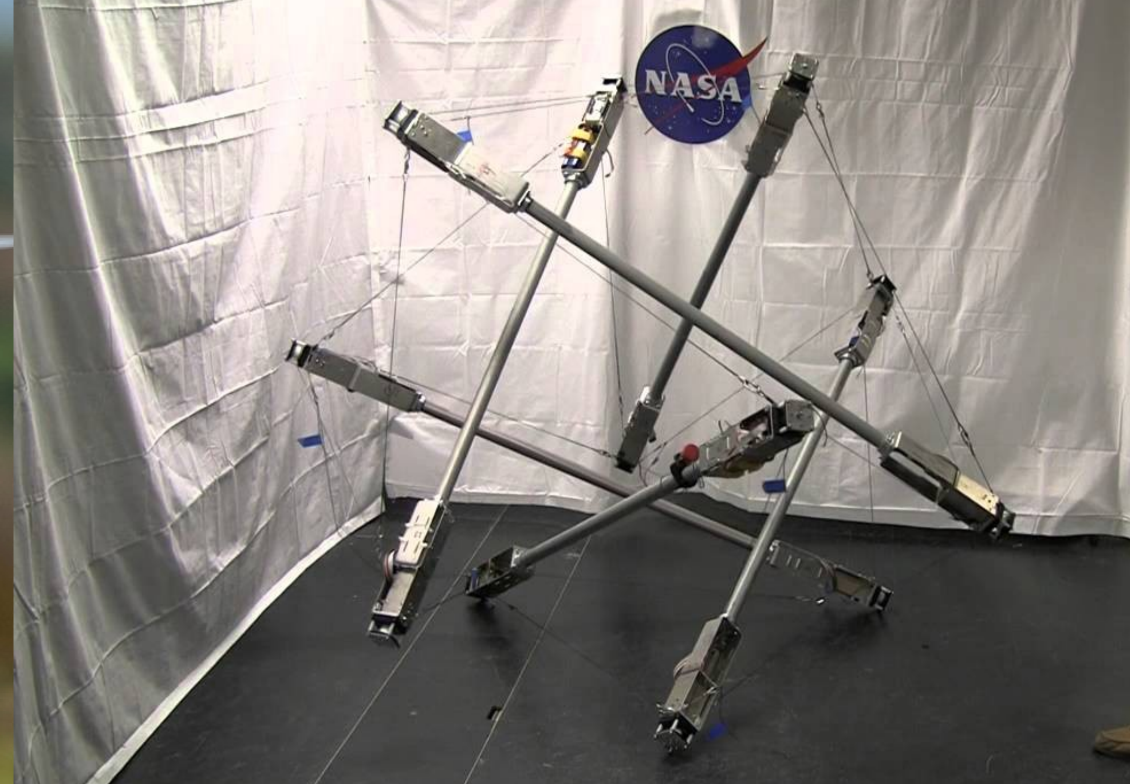
NASA


UBER
ADVANCED
TECHNOLOGIES
CENTER
UBERATC.COM/

Why Evolutionary Robotics?

Designing robots is hard!

Honda Asimo, 2006

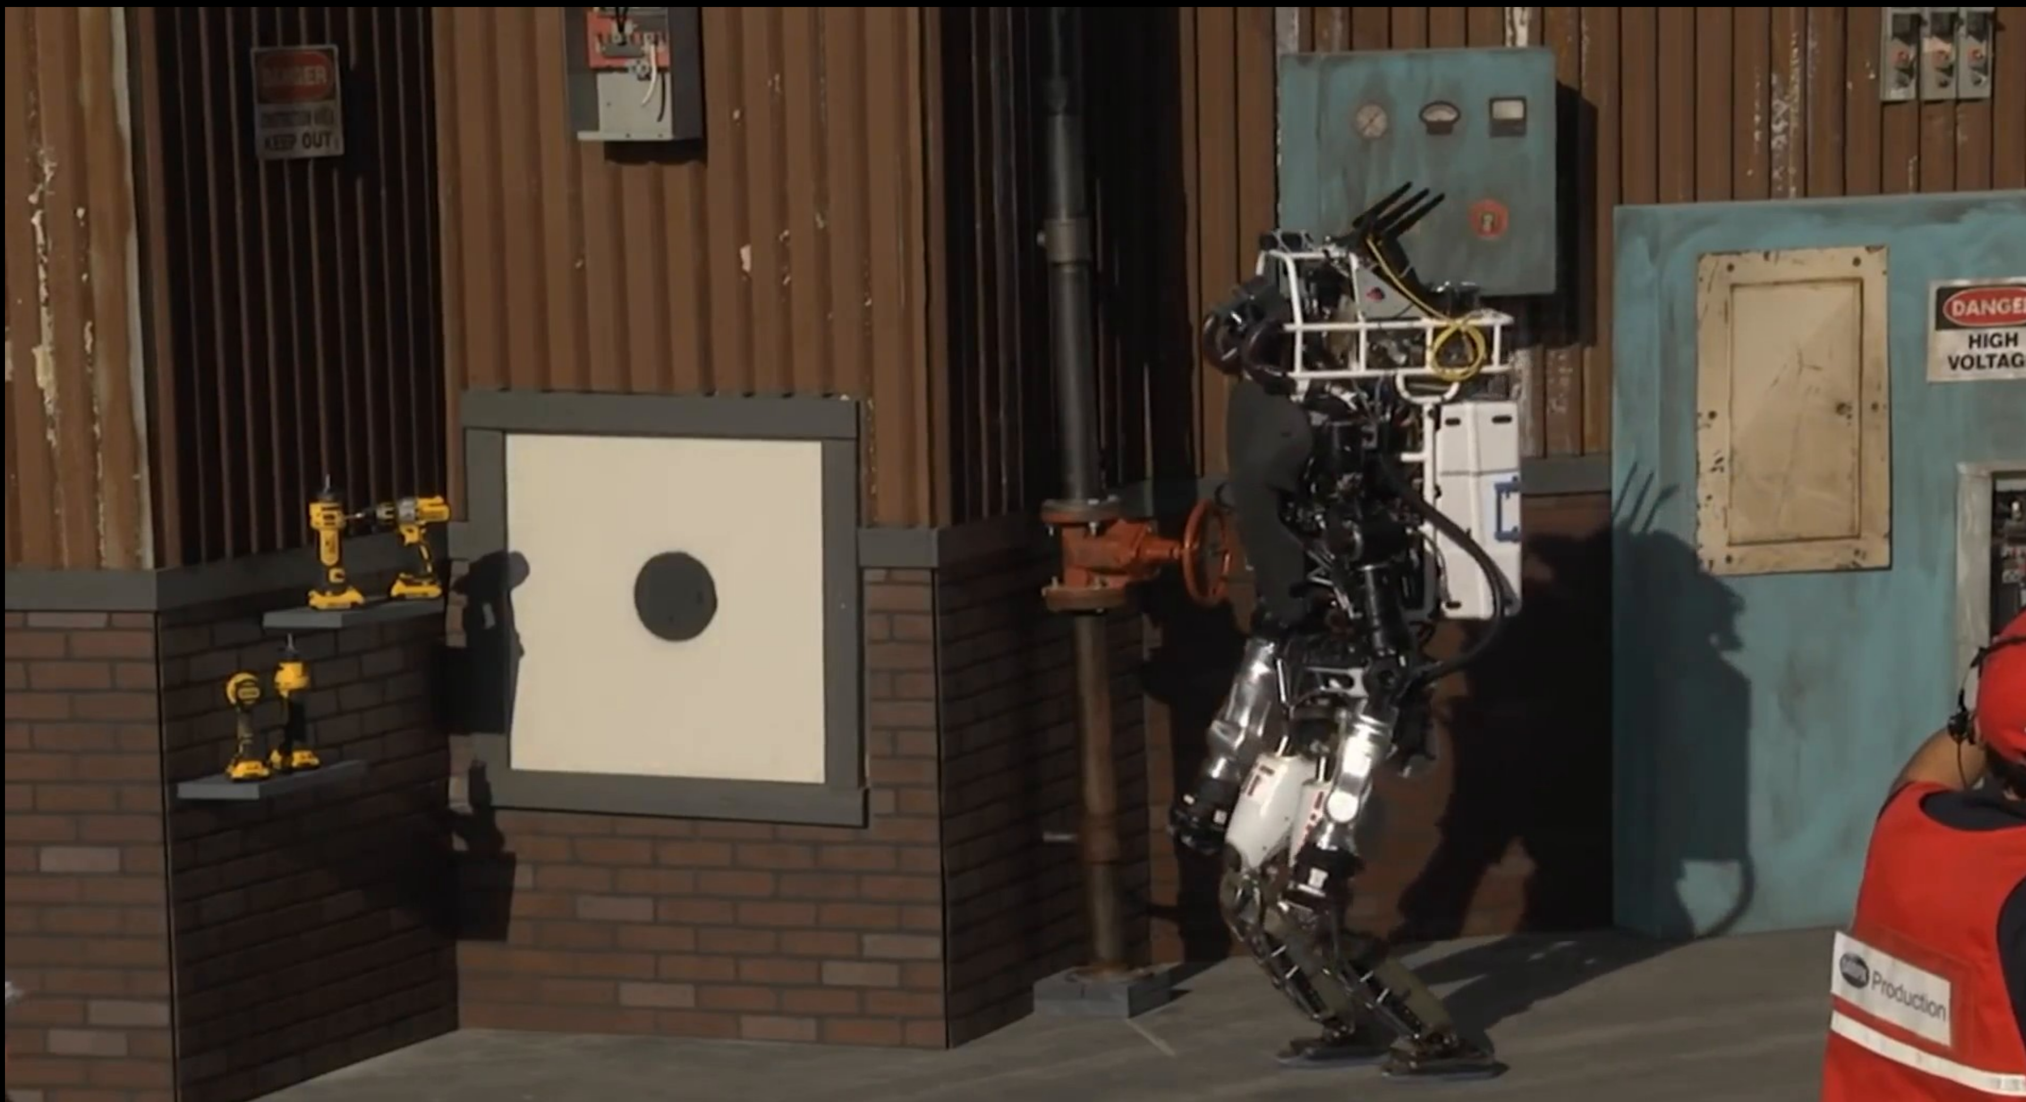

DARPA Robotics Challenge, 2015

DARPA Robotics Challenge, 2015

Robots have not historically been designed for autonomous behavior or use in dynamic and unstructured environments


NO STEP
NO STEP
8X-060-04

But robot controller and body-plans that are hand-designed by engineers (“top-down”) can't possibly account for every single instance and specify the ideal behavior in that scenario

Behaviors in biology are “emergent” (“bottom-up”) properties that are built by experience and trial-and-error, rather than specified *a priori*


400

Hand-engineered solution often rely on complex multi-step computation pipelines

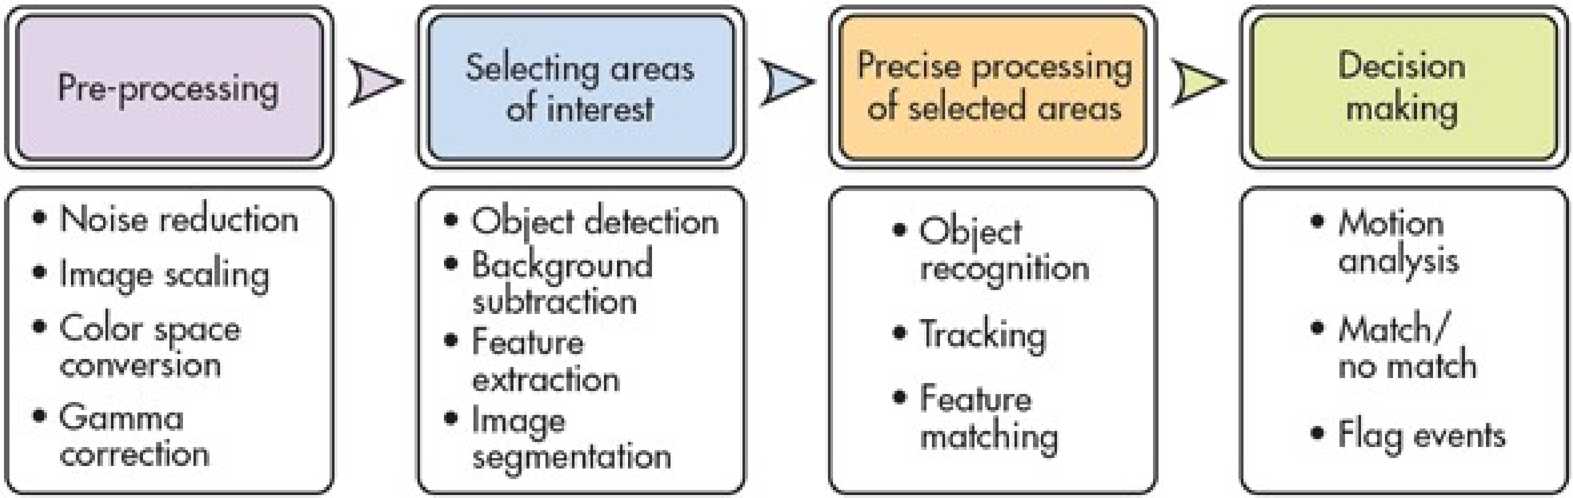

Biology find much simpler “hacks”

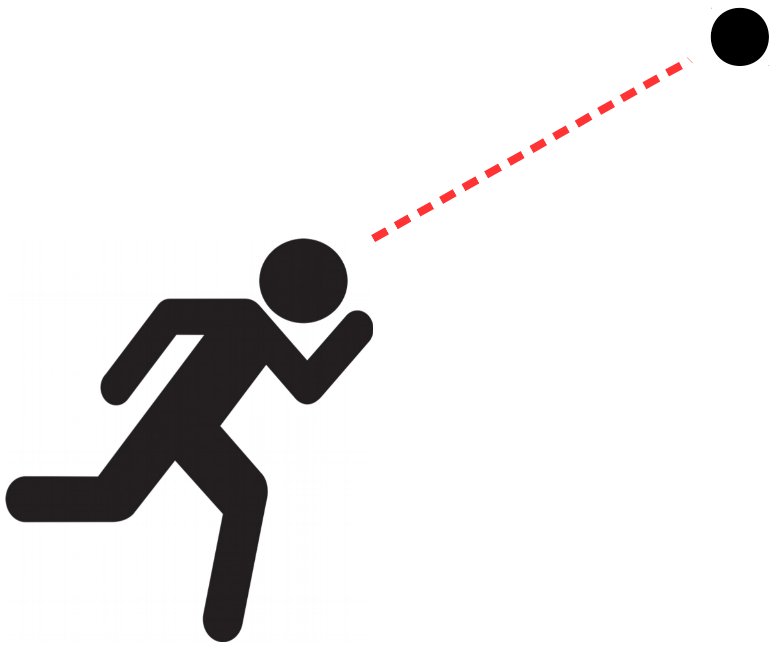

McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there?. Journal of experimental psychology: human perception and performance, 22(3), 531.

Thus, the best approaches are bio-inspired!

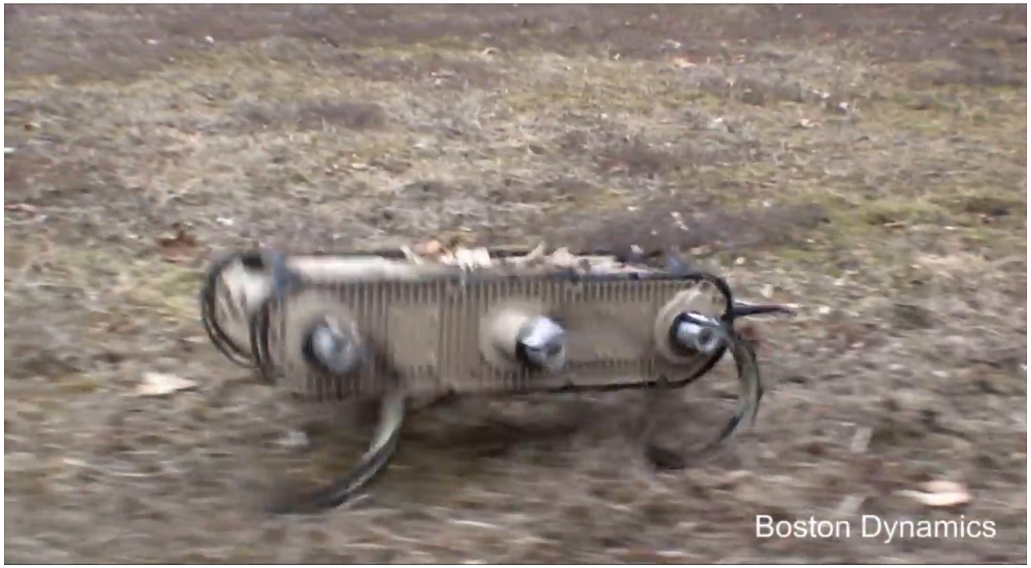
Boston Dynamics

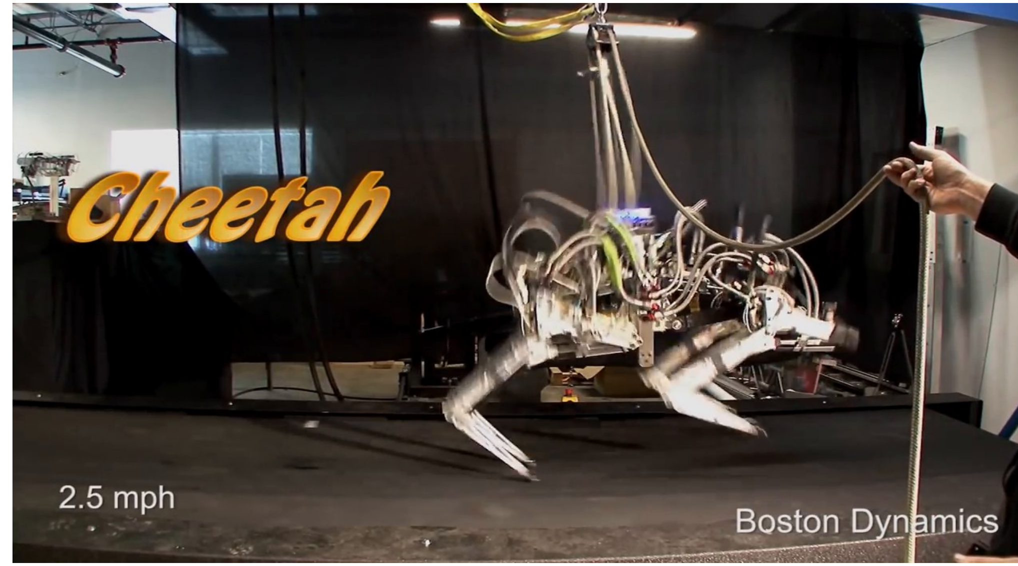
Cheetah
2.5 mph
Boston Dynamics

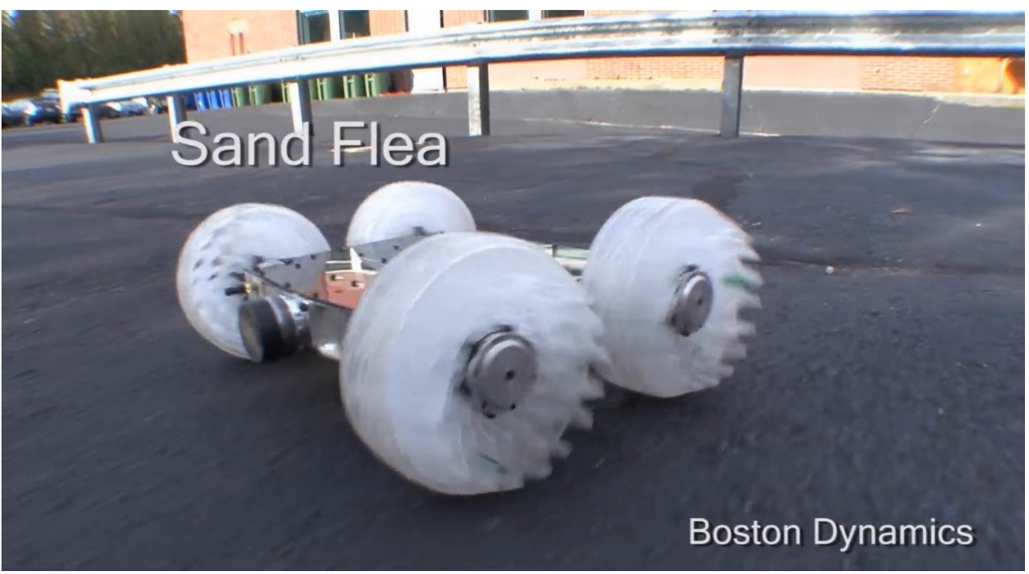
Sand Flea
Boston Dynamics

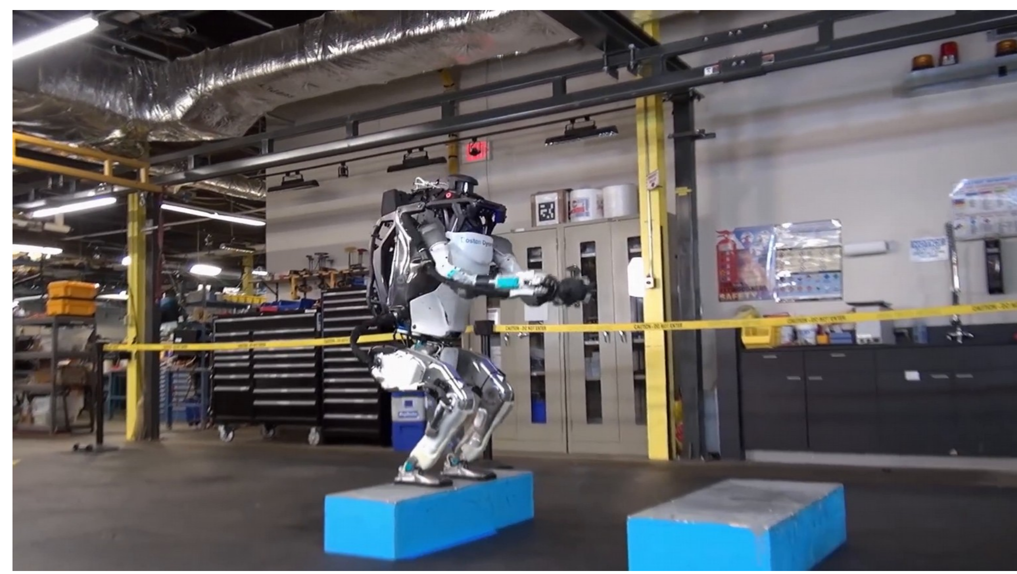

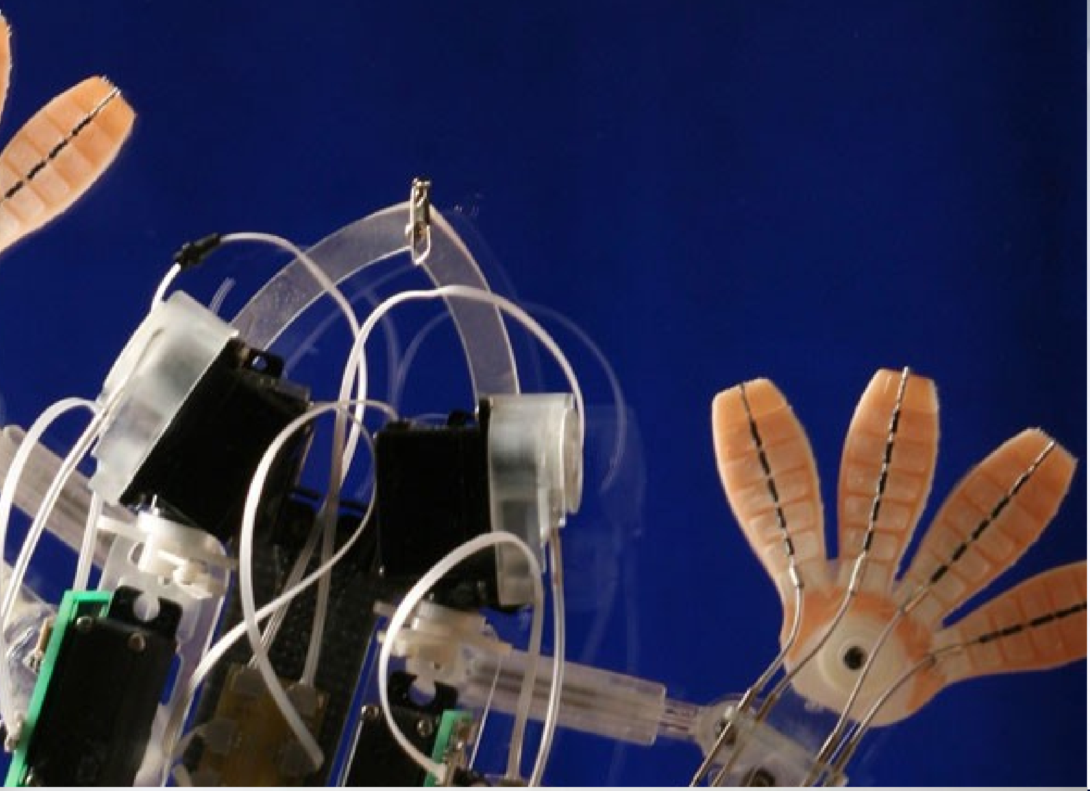
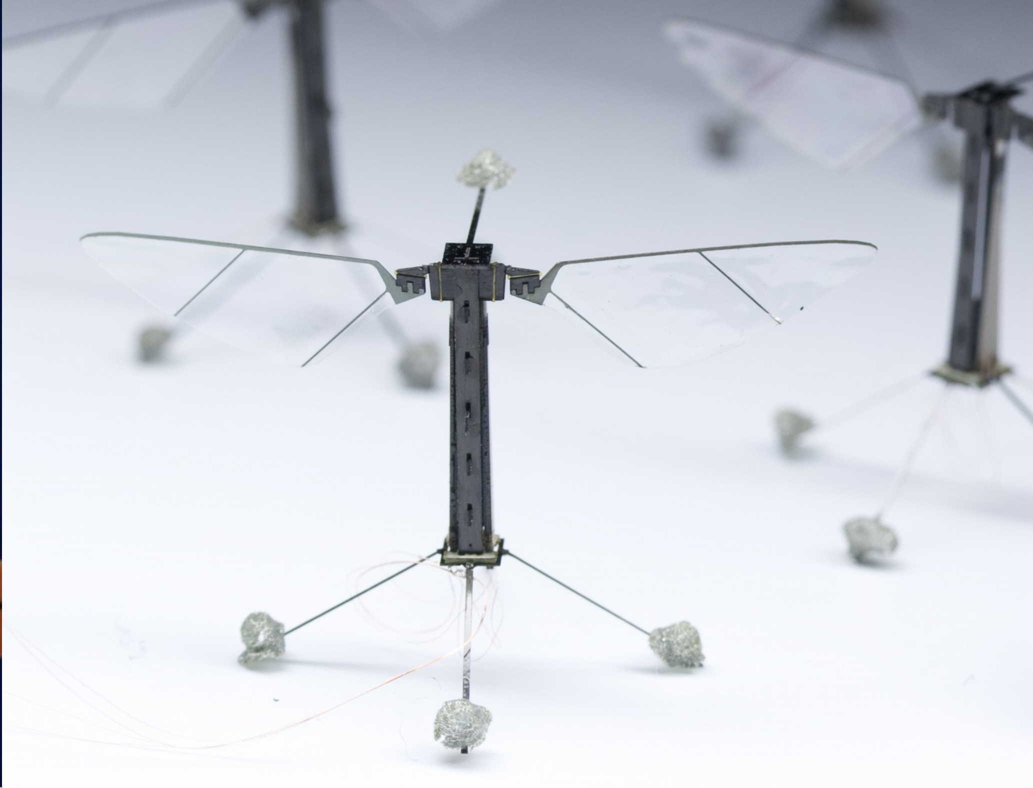
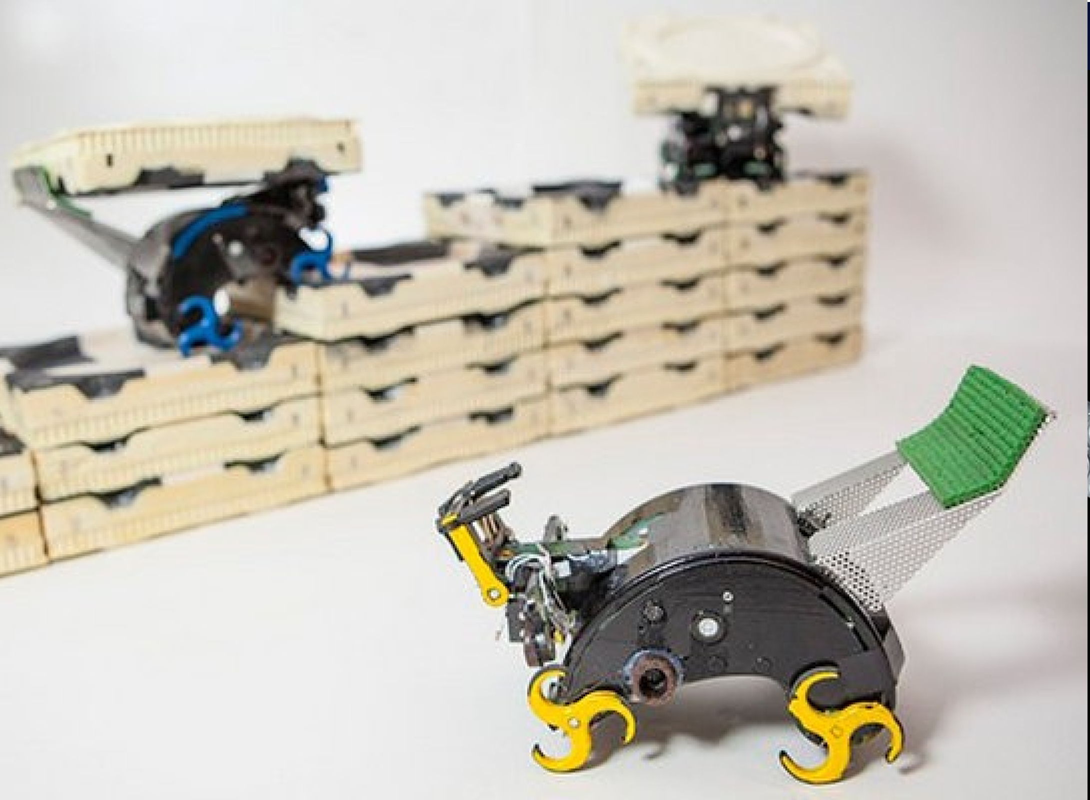


But designing a new robot for each task, environmental niche, or desired features set is time consuming and expensive

(and again assumes a task and behavior *a priori*)

Boston Dynamics Gets $10 Million from DARPA for New Stealthy, Bulletproof LS3

By Evan Ackerman
Posted 23 Sep 2013 | 14:42 GMT

You'd think that Boston Dynamics would be all kinds of busy building (and supporting) a small army of Atlas robots for the DARPA Robotics Challenge. But, it looks like they've somehow managed to find the time to continue working on all of their other systems as well, like BigDog's big brother, LS3. Last week, DARPA committed to investing an extra $10 million towards a more robust and (eventually) deployable robot.

DARPA is very specific about what they want to see in the next-gen LS3, namely:

So instead of taking inspiration from
specific biological creatures,

let's take inspiration from biology's higher-level
optimization principles and methodologies

since all biological creatures have these in common,
we only need to do it once!

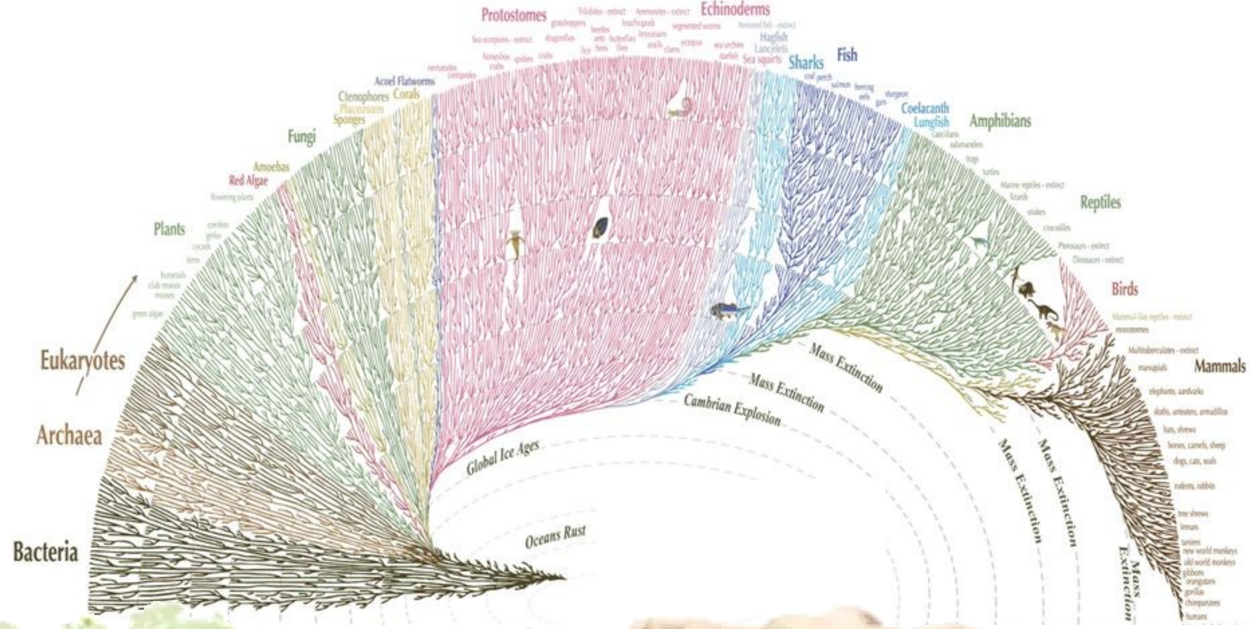

Protostomes
Echinoderms
Hagfish
Lancelets
Sea squirts
Sharks
Fish
Coelacanth
Lungfish
Amphibians
Reptiles
Birds
Mammals
Acoel Flatworms
Ctenophores
Corals
Placozoans
Sponges
Fungi
Amoebas
Red Algae
Plants
Eukaryotes
Archaea
Bacteria
Global Ice Ages
Oceans Rust
Cambrian Explosion
Mass Extinction
Mass Extinction
Mass Extinction
Mass Extinction
Mass Extinction

The algorithms are conceptually simple

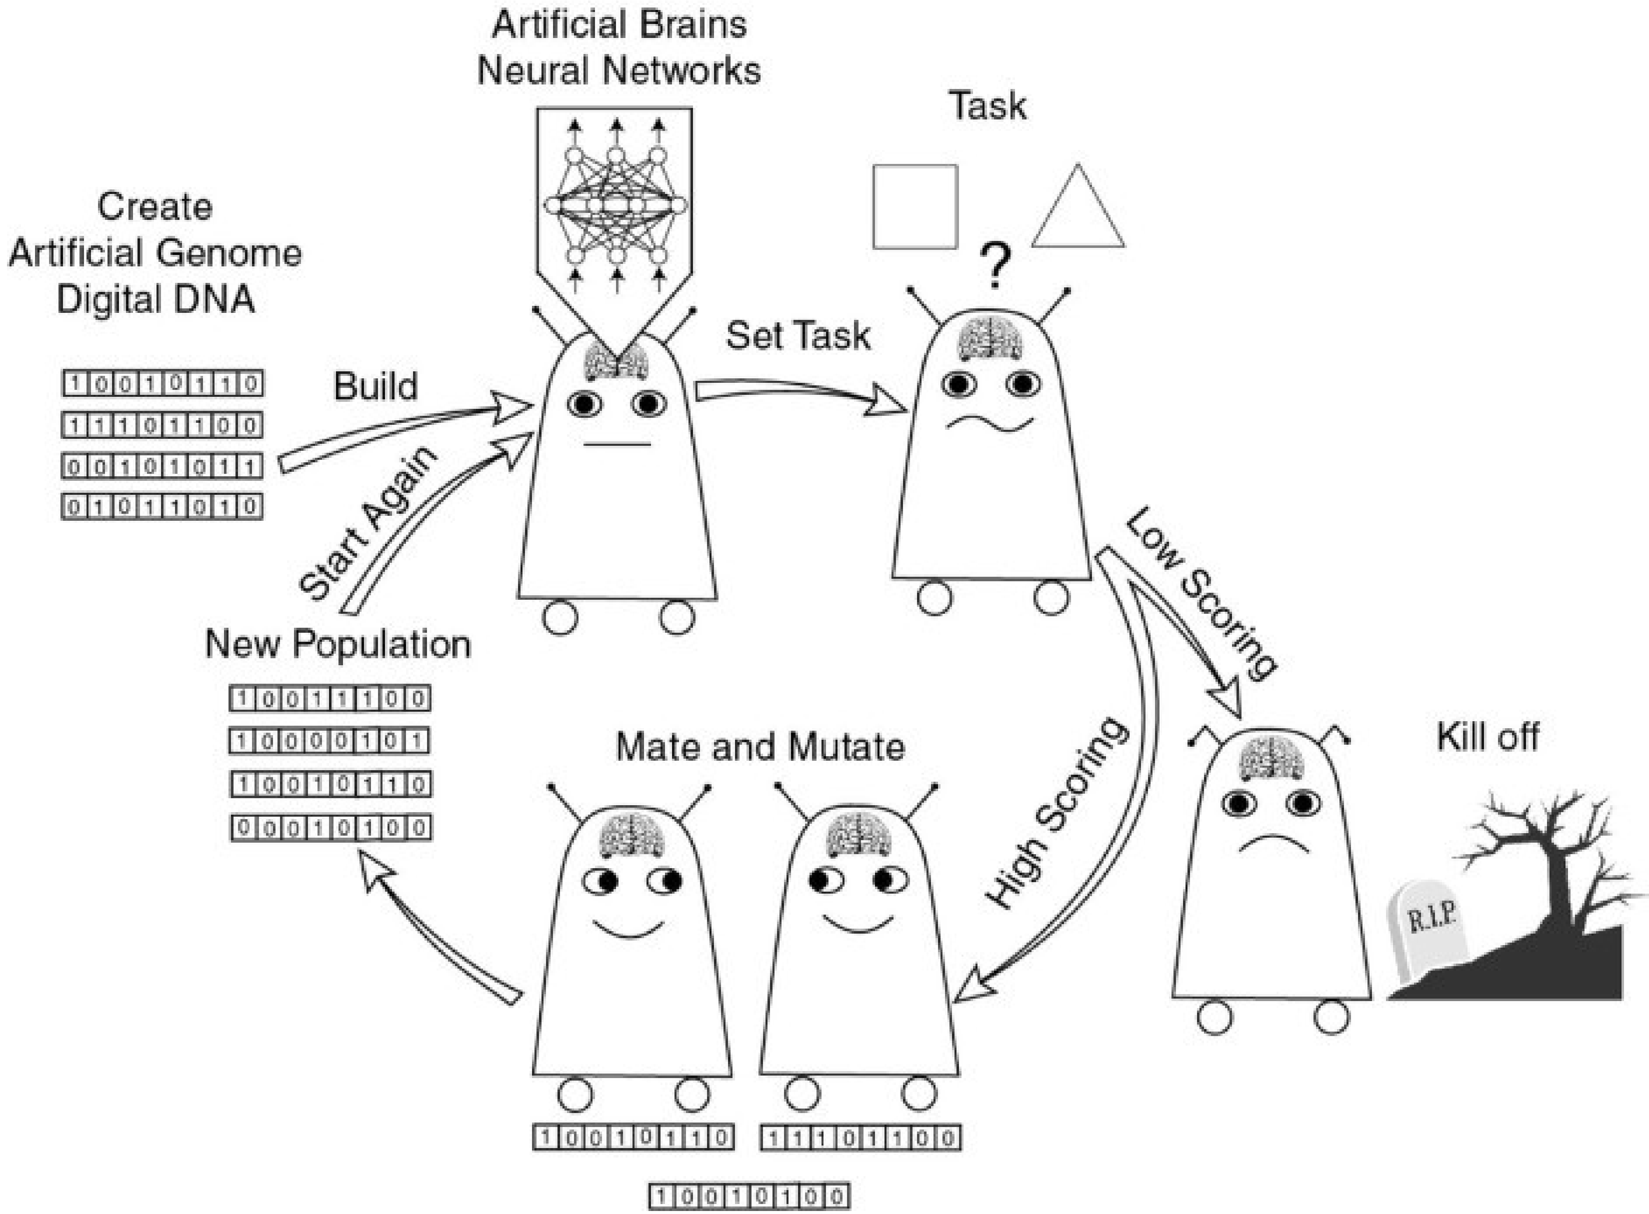
Artificial Brains
Neural Networks
Task
?
Create
Artificial Genome
Digital DNA
10010110
11101100
00101011
01011010
Build
Set Task
Start Again
Low Scoring
New Population
10011100
10000101
10010110
00010100
Mate and Mutate
High Scoring
Kill off
R.I.P.
10010110
11101100
10010100

As population-based optimizers,
they don't require closed-form expression
for the cost function for the
behavior/robot being optimized

(nor a supervised example of ideal behavior)
“creative machines”

Softmax Regression

Cost Function

We now describe the cost function that we'll use for softmax regression. In the equation below, $1\{\cdot\}$ is the **indicator function,** so that $1\{\text{a true statement}\} = 1$, and $1\{\text{a false statement}\} = 0$. For example, $1\{2+2=4\}$ evaluates to 1; whereas $1\{1+1=5\}$ evaluates to 0. Our cost function will be:

$$J(\theta) = -\frac{1}{m}\left[\sum_{i=1}^{m}\sum_{j=1}^{k} 1\left\{y^{(i)} = j\right\}\log\frac{e^{\theta_j^T x^{(i)}}}{\sum_{l=1}^{k} e^{\theta_l^T x^{(i)}}}\right]$$

Notice that this generalizes the logistic regression cost function, which could also have been written:

$$\begin{aligned} J(\theta) &= -\frac{1}{m}\left[\sum_{i=1}^{m}(1-y^{(i)})\log(1-h_\theta(x^{(i)})) + y^{(i)}\log h_\theta(x^{(i)})\right] \\ &= -\frac{1}{m}\left[\sum_{i=1}^{m}\sum_{j=0}^{1} 1\left\{y^{(i)} = j\right\}\log p(y^{(i)} = j | x^{(i)}; \theta)\right] \end{aligned}$$

The softmax cost function is similar, except that we now sum over the k different possible values of the class label. Note also that in softmax regression, we have that $p(y^{(i)} = j | x^{(i)}; \theta) = \frac{e^{\theta_j^T x^{(i)}}}{\sum_{l=1}^{k} e^{\theta_l^T x^{(i)}}}$.

There is no known closed-form way to solve for the minimum of $J(\theta)$, and thus as usual we'll resort to an iterative optimization algorithm such as gradient descent or L-BFGS. Taking derivatives, one can show that the gradient is:

$$\nabla_{\theta_j} J(\theta) = -\frac{1}{m}\sum_{i=1}^{m}\left[x^{(i)}\left(1\{y^{(i)} = j\} - p(y^{(i)} = j | x^{(i)}; \theta)\right)\right]$$

They can be used as a testbed for controlled, repeatable experiments that are difficult or impossible to perform in a biological setting

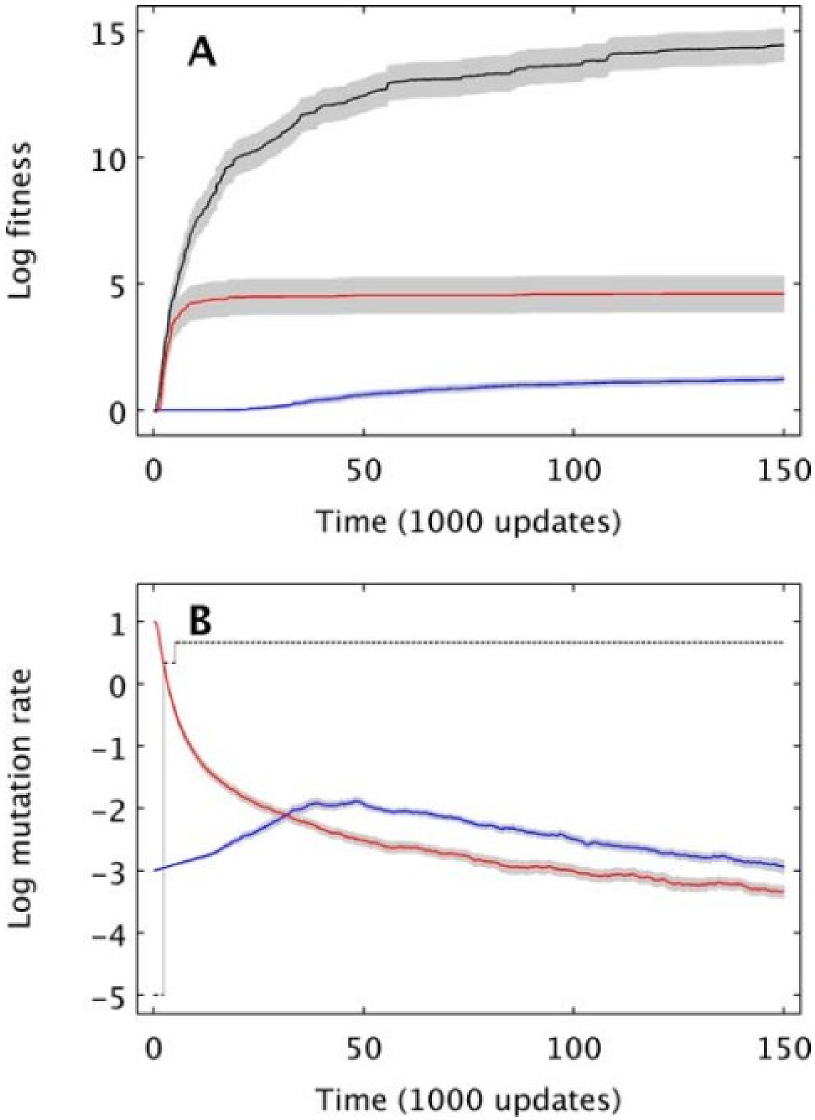

Figure 2. Evolutionary trajectories for fitness and mutation rate on a complex fitness landscape. (A) Evolution of average log-fitness±1 s.e.m. for treatments with the mutation rate fixed at $U_{opt} = 4.641$ (black) and for treatments with variable mutation rates starting at either 10 (red) or 10^{-3} (blue). (B) Evolution of average log genomic mutation rate±1 s.e.m. for treatments with variable mutation rates starting at either 10 (red) or 10^{-3} (blue). The black line indicates the mutation rate that had produced the highest average fitness for that time point.

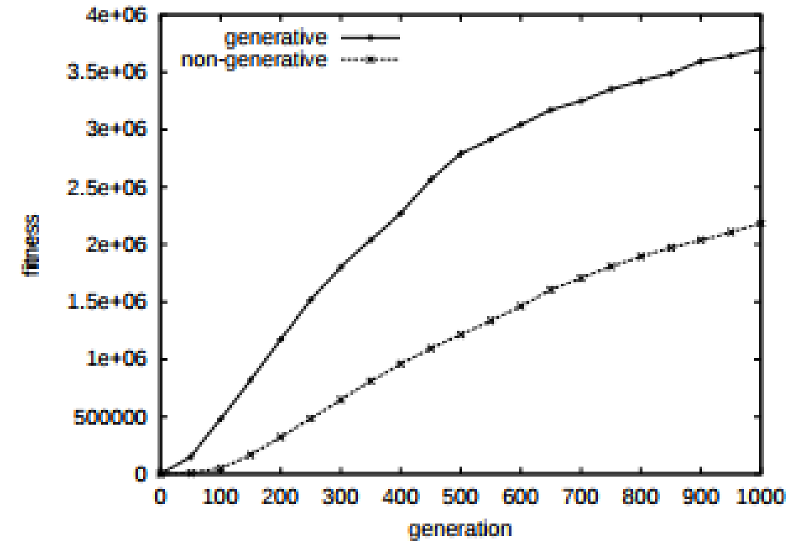

Figure 2: Performance comparison between the non-generative encoding and the P0L-system generative encoding.

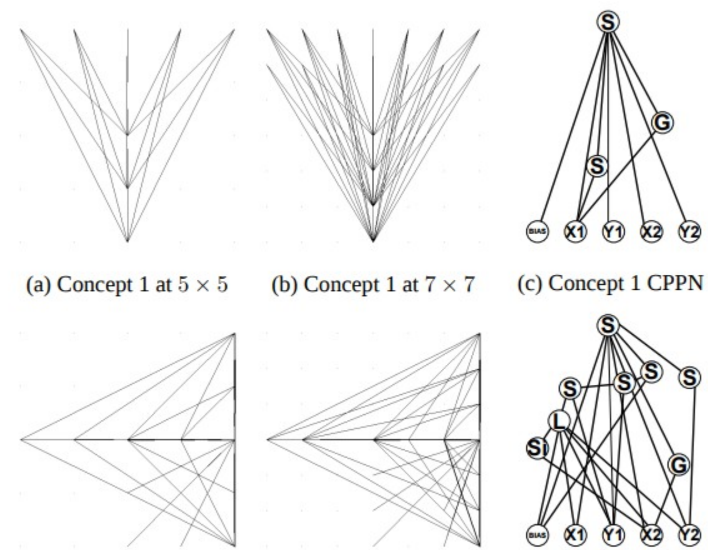

(a) Concept 1 at 5 × 5 (b) Concept 1 at 7 × 7 (c) Concept 1 CPPN

(d) Concept 2 at 5 × 5 (e) Concept 2 at 7 × 7 (f) Concept 2 CPPN

Figure 8: **Equivalent Connectivity Concepts at Different Substrate Resolutions.** Two connectivity concepts are depicted that were evolved through interactive evolution. The CPPN that generates the first concept at 5 × 5 (a) *and* 7 × 7 (b) is shown in (c). The CPPN in (f) similarly generates the second concept at *both resolutions* (d) and (e). This illustration demonstrates that CPPNs represent a mathematical concept rather than a single structure. Thus, the same CPPN can produce patterns with the same underlying concept at different substrate resolutions (i.e. different node densities). *CPPN activation functions in this paper are denoted by G for Gaussian, S for sigmoid, Si for sine, A for absolute value, and L for linear.*

ARTIFICIAL LIFE 13

13th INTERNATIONAL CONFERENCE ON THE SIMULATION & SYNTHESIS OF LIVING SYSTEMS

The Biology of Possible Life

- life-as-we-know-it vs. life-as-it-could-be
  - Biology: study of *life-as-we-know-it,* based on carbon-chain chemistry, the only kind of life available for study.
  - Is it possible to derive general theories from single examples?
  - Life, as a dynamic physical process, could "haunt" other physical material. What matters is the organization of such material.

Frank and Ernest