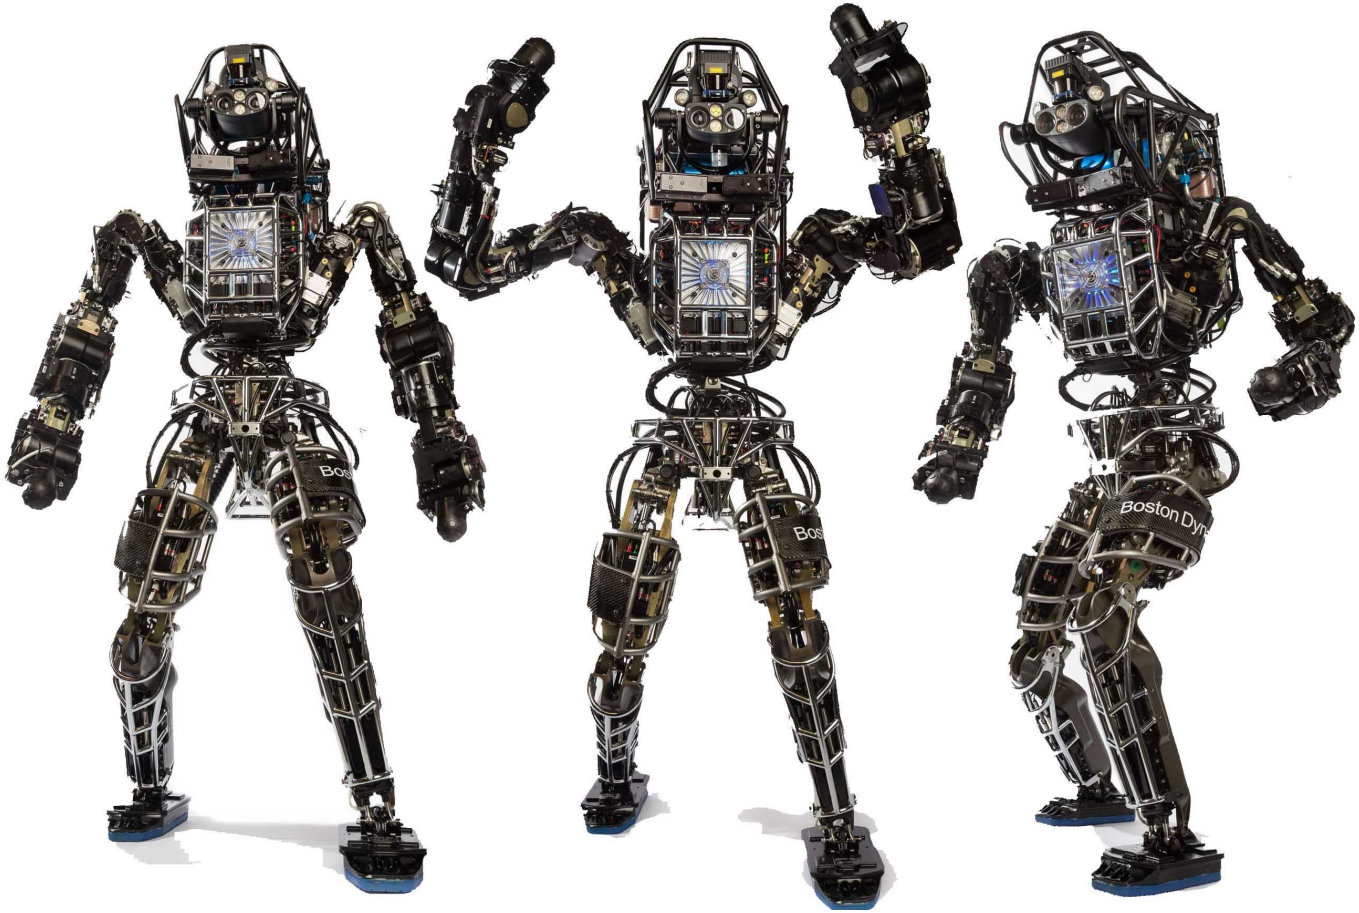
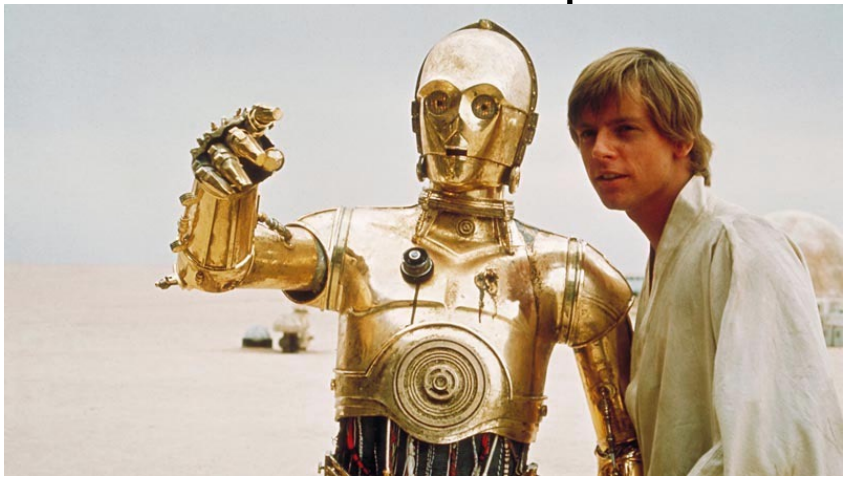
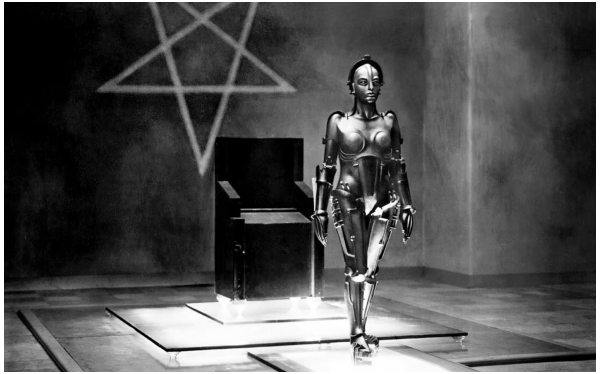

Boston Dyn

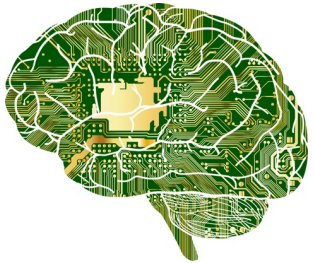
Where are the intelligent robots we were promised?



Maybe we took a **wrong turn** somewhere…



We originally imagined the mind as a kind of logical reasoning device coupled with a store of explicit data—a kind of combination logic machine and filing cabinet.



V.S.



We ignored the fact that minds evolved to make things happen. We ignored the fact that the biological mind is, first and foremost, an organ for controlling the biological body.



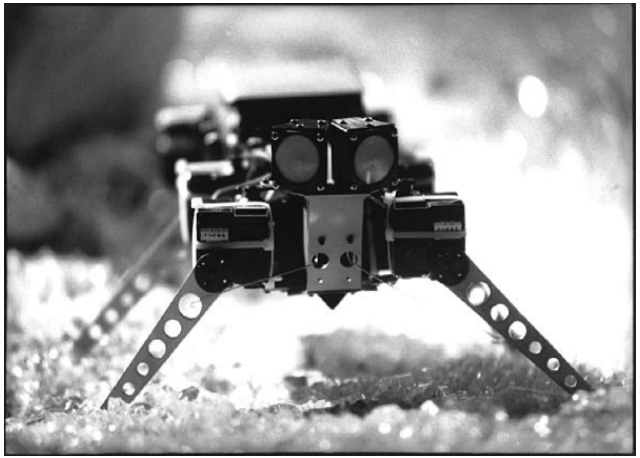
What is the block
in the road?

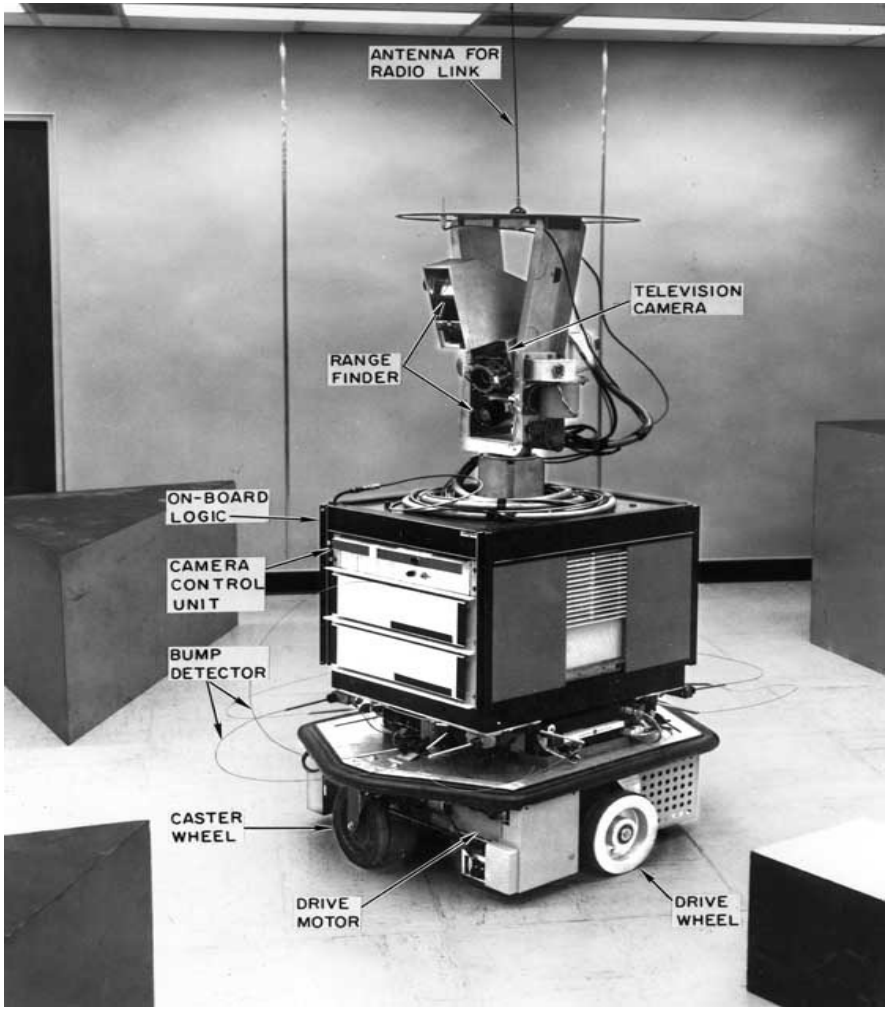
Some think it is knowledge—that the puzzles of mind will fall away once a nice big knowledge base is in place (like CYC).

Others think it's not knowledge but basic, real-time, real-world responsiveness.

Plan This Week:

- Introduce some key ideas from **robotics and embodied cognitive science.**
- Ask what difference **body, action, perception and environment** make to the **kinds of strategies** used by biological brains…





Shakey (1970) Stanford Research Institute.

One of the very first computer-controlled mobile robots

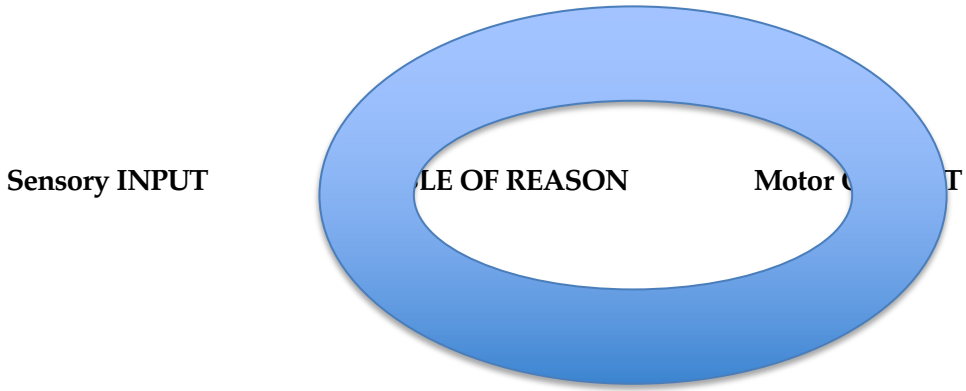
Shakey had a body, but was not interestingly 'embodied'!

For Shakey, the body and the environment were first and foremost problems to be solved.

- The environment was the problem arena.
- The sensors detected the lay-out in that arena.
- The reasoning system planned a solution.

The body and motor systems simply **put the solution into practice.**

Problem-solving using a **'bubble of reason'** (aka classical sandwich – see Susan Hurley (2001)) approach



Hurley (2001) Perception And Action: Alternative Views Synthese 129: 3-40.

INSIDE THE BUBBLE

World Models (can 'throw away the world)
Representations
Inference, Computation
Reason, Thought

SENSE- **THINK**- ACT

Traditionally, the sciences of mind tended to look mostly **inside the 'think' bubble….**

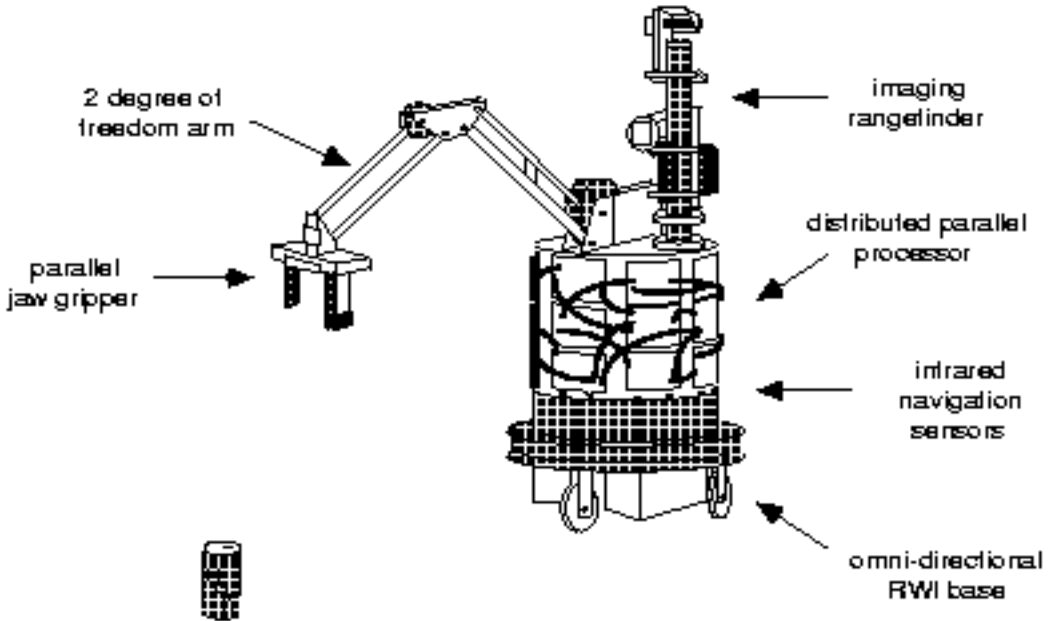
In the late 1980's, Rodney Brooks' mobile robots ('mobots'), such as Herbert, replaced the **single big bubble of reason** with lots of **small bubbles** – lots of **semi-independent routines running directly from sensory intake to motor action**.





Herbert Task: collect soft drink cans left around a crowded laboratory.

Herbert = a collection of **coarse sensors** and simple, relatively independent, **behavioural routines**.



Herbert's Routines:

- Basic obstacle avoidance via a ring of **ultrasonic sound sensors**
- General locomotion (randomly directed) - **interrupted** if Herbert's simple visual system detected a roughly table-like outline.
- At this point a **new routine kicks in and the table surface is swept using a laser.**
- If the outline of a can is detected, the **whole robot rotates until the can is centered** in its field of vision.
- This physical action simplifies the **pick-up procedure** by creating a standard action-frame in which the **robot arm, equipped with simple touch sensors, gently skims the table surface dead ahead.**
- Once a can is encountered, it is grasped, collected and the robot moves on.

This is known as a **‘subsumption architecture’**, because the layers each constitute a complete behavior-producing system and interact only in simple ways such as by one layer **subsuming** (turning off) the activity of another (see Brooks (1991)).

Herbert succeeds – in its simple little task - without using any conventional planning techniques and without creating and updating any **detailed inner model of the environment.**

Herbert's 'world' is composed of undifferentiated obstacles and rough table-like and can-like outlines.

Within this world the robot **also exploits its own bodily actions** (rotating the 'torso' to center the can in its field of view) so as to greatly **simplify the computational problems** involved in eventually reaching for the can.

Herbert succeeds – in its simple little task - without using any conventional planning techniques and without creating and updating any **detailed inner model of the environment.**

Herbert's 'world' is composed of undifferentiated obstacles and rough table-like and can-like outlines.

Within this world the robot **also exploits its own bodily actions** (rotating the 'torso' to center the can in its field of view) so as to greatly **simplify the computational problems** involved in eventually reaching for the can.

Herbert is thus a super-simple example of a system that succeeds using **minimal internal representational resources**

Aside: In thinking about this stuff, Brooks was inspired to be thinking about how evolution might have started with very simple routines then built more on top, layer-by-layer, without undoing (too much of) what went before.

Herbert is a bunch of **simple tricks, orchestrated by the environment** itself.

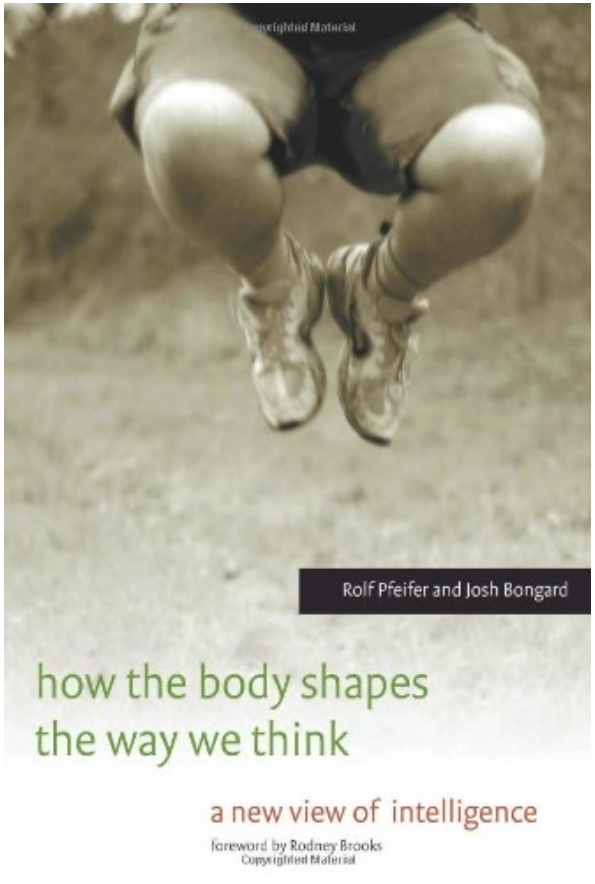
Is this a good model of basic adaptive intelligence?


ANY
QUESTIONS?

Real-world robots have come a long way since the early days of Brooks' '**mobots**'.

But the newer developments still owe much to these bare foundations – to the idea of being **frugal in the use of internal resources** while making the most of whatever help the body, world, and action can provide.





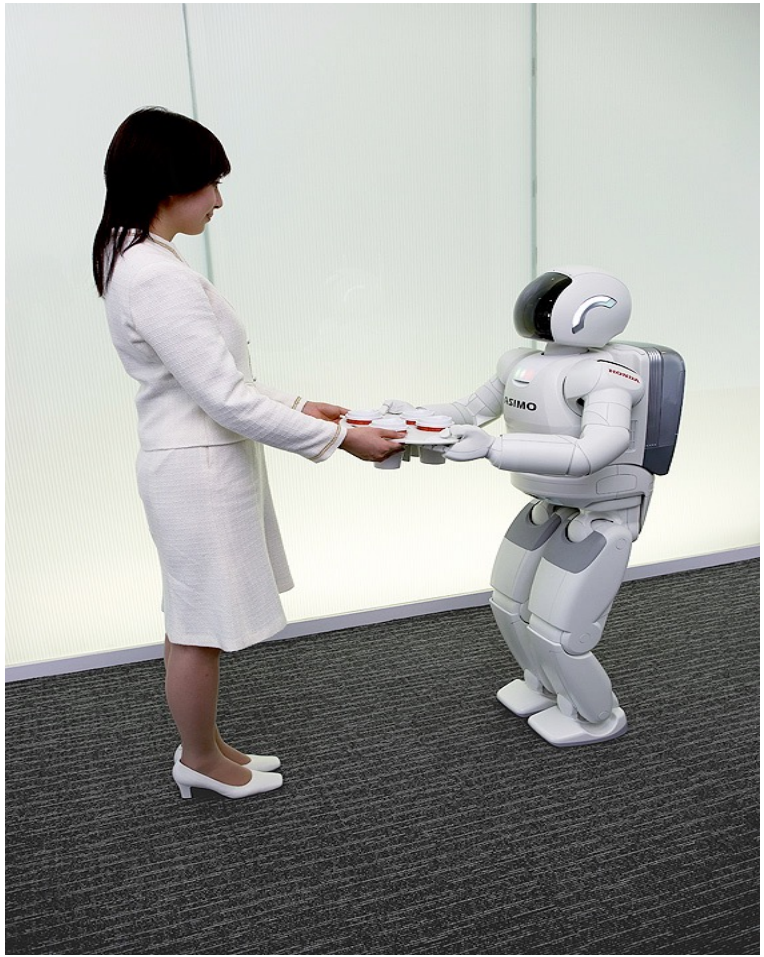
Robotics has also taught us important lessons about how to make the most of your own basic bodily form and structure, using inbuilt ‘bio-mechanics’ to simplify the tasks of the neural controller.

Consider the task of **controlling two-legged locomotion**….



Honda’s ASIMO





Honda's Asimo

A highly advanced humanoid robot displaying a full 26 degrees of freedom.

But ASIMO does not (or not yet!) exploit its own body to the full.

Consider **energy efficiency….**

Energy efficiency can be measured via **'specific cost of transport' = the amount of energy required to carry a unit weight a unit distance.**

SCORES (the lower the better: taken from Collins and Ruina (2005)

Asimo: 3.2
Typical Human: 0.2

Compare: **Passive Dynamic Walkers (PDW's)**

No actuation (power) except gravity, and no use of calculated joint angle control at any time!

Yet surprisingly, PDW's are capable (when set on a gentle incline) of very stable, human-looking walking.



A third generation Passive Dynamic Walker (S. Collins)

Here, **the body is part of the solution.**



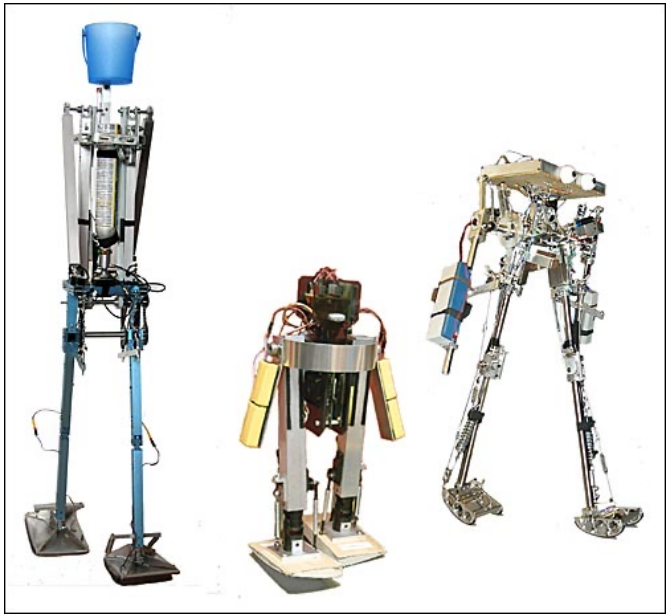
Q/ How to build on this kind of fluency in a powered walking device?

A/ Rely on a control regime that systematically **pushes, damps and tweaks** a system in which Passive Dynamic effects continue to play a very major role = **control as a gentle nudge to a complex system..**

In this way, a low energy source, a simple control system, and the body (and gravity!) can 'collaborate' to solve the walking problem.

= the power of **co-evolving morphology (shape) and control.**

A flurry of new robots have been built to make the most of this new design and control paradigm for walking robots



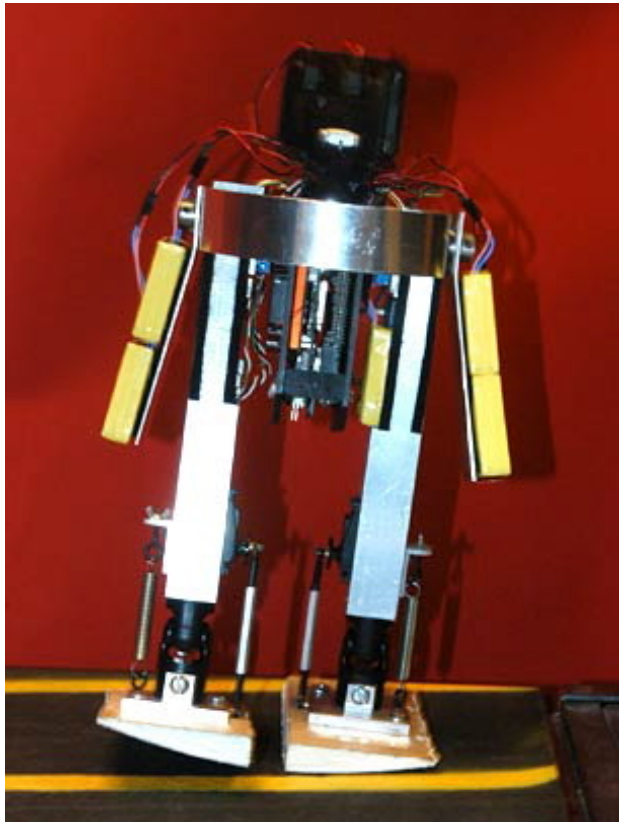


"Efficient Bipedal Robots Based on Passive-Dynamic Walkers" S. Collins, A. Ruina, R. Tedrake, M. Wisse *Science* 307, 1082-1085 (2005

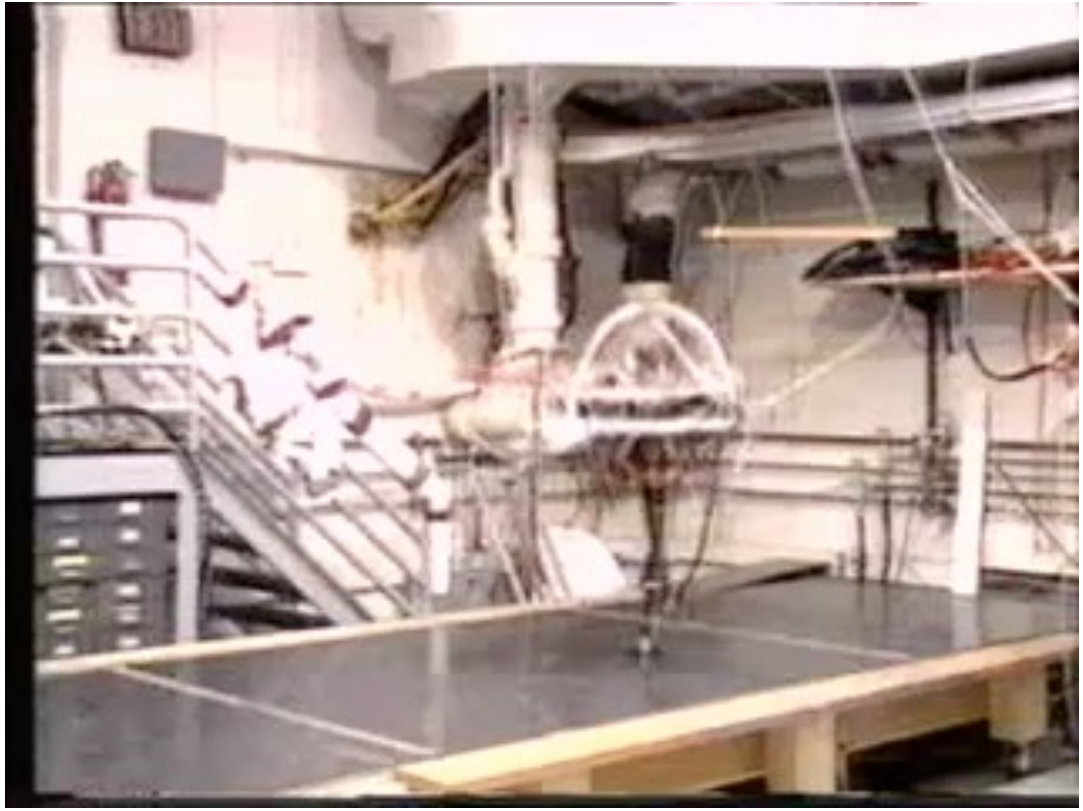
‘Robotoddler’ (Russ Tedrake, MIT, 2005)

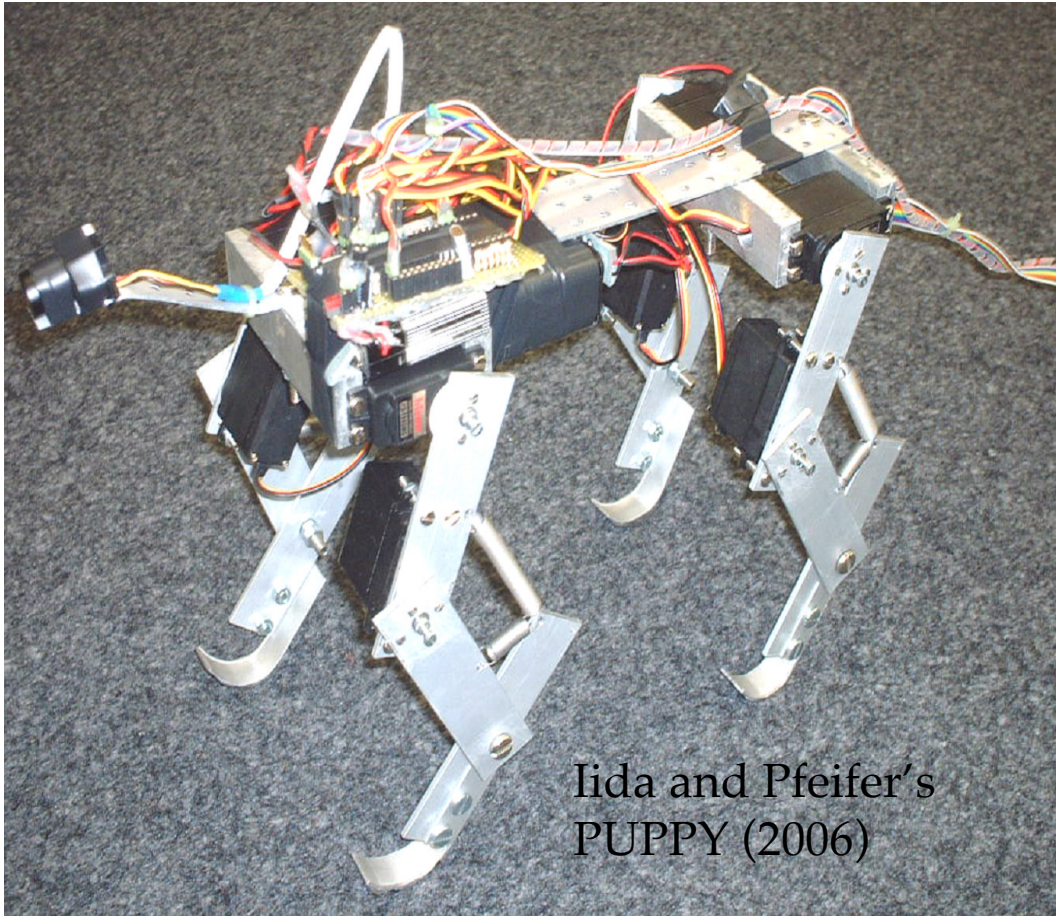
Learns a **control policy that exploits the passive dynamics of the body.**

Change speeds
Forward and backwards
On different terrains



Or…. the bouncing hopping robot, a very early (MIT leglab) example of combining control with rich bodily dynamics




Iida and Pfeifer's
PUPPY (2006)



Puppy's **aluminum legs and feet** play an important adaptive role: they induce, as the video shows, small amounts of **slippage** on most surfaces.

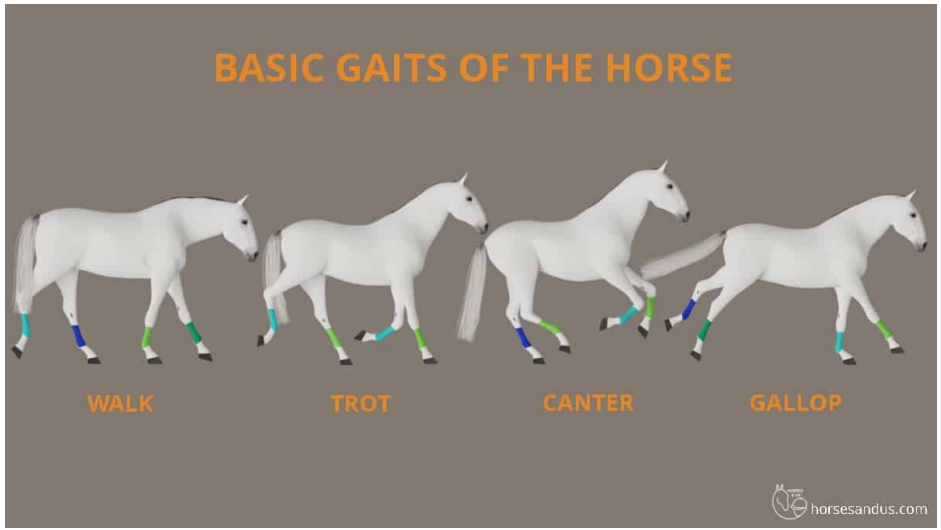
Reducing the slippage by adding rubber pads to the feet actually caused the robot to begin to fall over!

The subtle slippage was actually playing a stabilizing role, effectively **enabling the robot to rapidly search for a stable way to proceed**

(see Pfeifer and Bongard (2007) pp 96-100, 125-128 for discussion).

An important observation:

Complex bodily forms (e.g. combinations of mechanically linked legs, muscles, and tendons) determine **landscapes for control** (Pfeifer and Bongard, 2007)



An important observation:

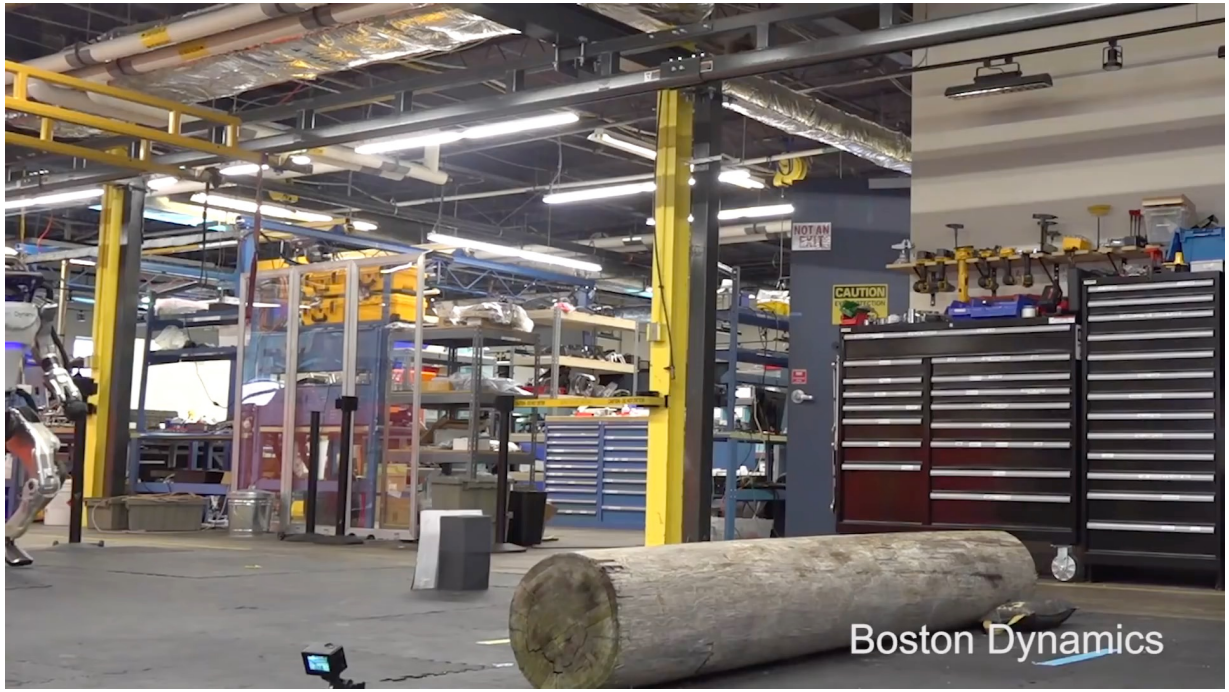
Complex bodily forms (e.g. combinations of mechanically linked legs, muscles, and tendons) determine **landscapes for control** (Pfeifer and Bongard, 2007)

This is because the bodily biomechanics themselves nudge the system towards a number of **stable (minimal energy expenditure) gaits** such as jumping, running, walking, trotting, and skipping, which the system repeatedly falls into **given its motor speeds and the local terrain.**

These biomechanically constituted ‘attractor states’ can be simply ‘assumed’ by **high-level control strategies** since the system will tend to fall back into these states whenever it is brought anywhere close to one of them.

In this way we get by with **nudges and tweaks to a system with rich intrinsic dynamics,** rather than more **complex commands** in a system with **simple intrinsic dynamics** (such as Asimo)


Boston Dynamics


NOT AN EXIT
CAUTION
Boston Dynamics

These are all cases of what Pfeifer and Bongard (2007, p.100) call ‘morphological computation’

“bodies can perform functions that would otherwise have to be performed by brains”

The First Moral of the Embodied Mind:

Control and processing leak into the body…

Bodily shape (morphology) and bodily bio-mechanics can simplify the neural commands required to bring about complex behaviours, effectively delegating aspects of control and processing to **the body itself.**

This will be important later when we come back to mind.

This makes good evolutionary sense, since **bodily forms, action repertoires, and the neural control system have co-evolved.**

It's not as if nature first evolved (say) the body of a crab and THEN had to find a neural control system to engage that body!!!

More?



Dr. Mark Miller

MONASH University UNIVERSITY OF TORONTO CHAIN

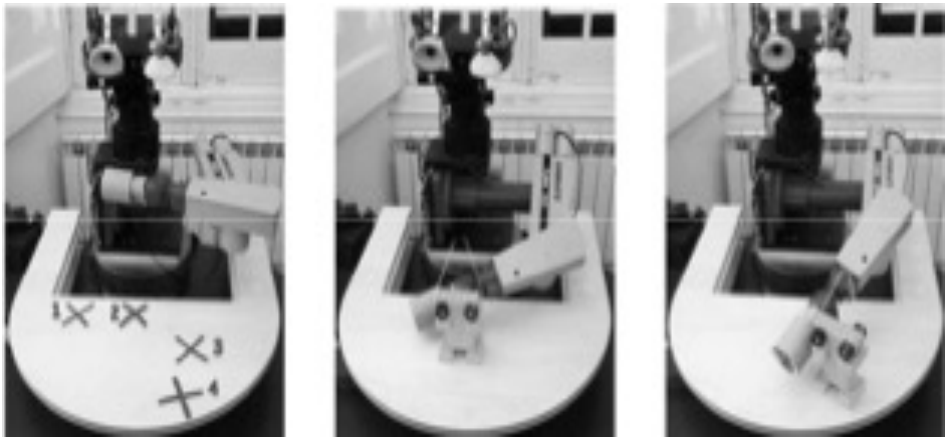
Next: **Action**

It is not just the gross form and bio-mechanics, that may be doing unexpected work, but also the **skilled use of the body in action.**

Key contribution of action concerns the power of what Lungarella and Sporns call the **active self-structuring of a data flow**

Lungarella, M., Sporns, O. (2005) Information self-structuring: key principle for learning and development. *Proceedings 2005 IEEE Intern. Conf. Development and Learning*, pp. 25-30.

A simple example: **BABYBOT** (Metta and Fitzpatrick) learns about object boundaries by poking and shoving.



Metta, G. and Fitzpatrick, P. (2003). Early integration of vision and manipulation. *Adaptive Behavior*, 11(2):109—128.

So BABYBOT's own motor action has now induced an **informative sensorimotor correlation**: the sudden spread of motion **identifies the boundaries of an object.**



This isn't just a trick for babies…

We are rampant **self-structured of our own informational flows**.

In higher cognition we constantly **sketch, scribble, gesture and talk to ourselves**, creating structured flows of visual and auditory stimulation that can enhance and guide our own cognitive activity

Experimental work by Goldin-Meadow suggest that **spontaneous gesture is doing important cognitive work**, and is not merely a side-effect or aid to communication.

Some more clues…

- We do it when we talk on the phone
- We do it when we talk to ourselves
- We do it in the dark, no one can see
- Gesturing increases when genuinely reasoning about a problem, rather than just describing a known solution.

Interestingly,

Speakers blind from birth, who have never spoken to a visible listener, and never seen others moving their hands as they speak, also gesture when they speak

Moreover, they do so even when they are speaking to others they know to be blind

More empirical research showing that if you suppress gestures you suppress cognitive task

A possible model:

In gesture, as when we write to talk, we materialize our own thoughts, we bring something concrete into being

That thing can then systematically reflect, but also alter and effect, our own ongoing thinking and reasoning

What Dennet (1991) calls the power of **cognitive self-stimulation**

Speech, gesture and neural activity continuously inform and are informed by each other, together constituting a single integrated cognitive system" (Iversen & Thelen 1999)

"A dynamic mutuality [involving words, gesture and neural activity] such that activity in any one component can potentially entrain activity in any other" (2001 p.37)

"the gesture, the actual motion of the gesture itself, is a dimension of thinking" (Mc Neil 2005 p.98)


MASTER MIND
ABACUS

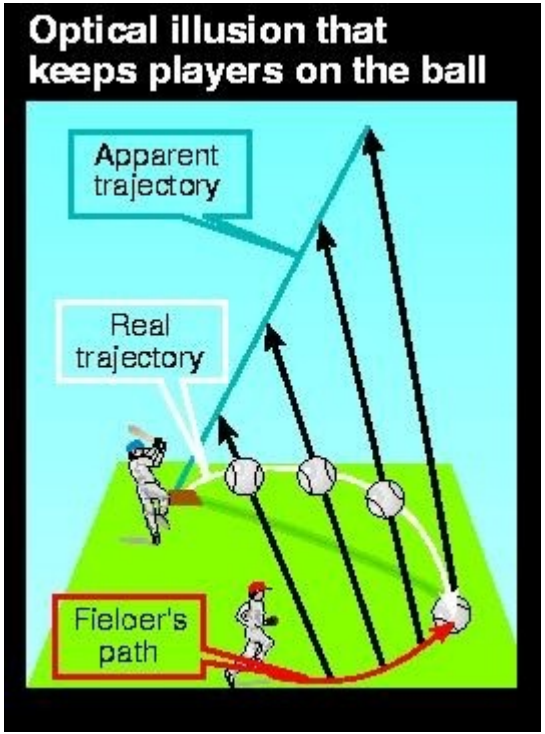
A general picture:

At many levels, **self-generated motor activity** acts as a "**complement** to neural information-processing", creating new, whole, **highly complex, systems** in which

"'information structuring' by motor activity and 'information processing' by the neural system are continuously linked to each other through sensorimotor loops" Lungarella and Sporns (2005).

Next: The Role of **Sensing**…

Active embodied agents can also use **sensing** in ways that are rather different to the way imagined in the classical bubble of reason/ Sense-Think-Act cycle picture



Consider running to catch a fly ball in baseball…

Run so you cancel the **vertical acceleration of the ball's optical projection** – see the neat VR experiments reported in Fink, Foo, and Warren (2009)

McBeath, M. K., Shaffer, D. M., & Kaiser, M. K. (1995). How baseball outfielders determine where to run to catch fly balls. *Science*, 268, 569–573.

Fink, P. W., Foo, P. S., & Warren, W. H. (2009). Catching fly balls in virtual reality: A critical test of the outfielder problem. *Journal of Vision*, 9(13):14, 1–8



Similar for how dogs catch frisbees (Shaffer et al (2004))

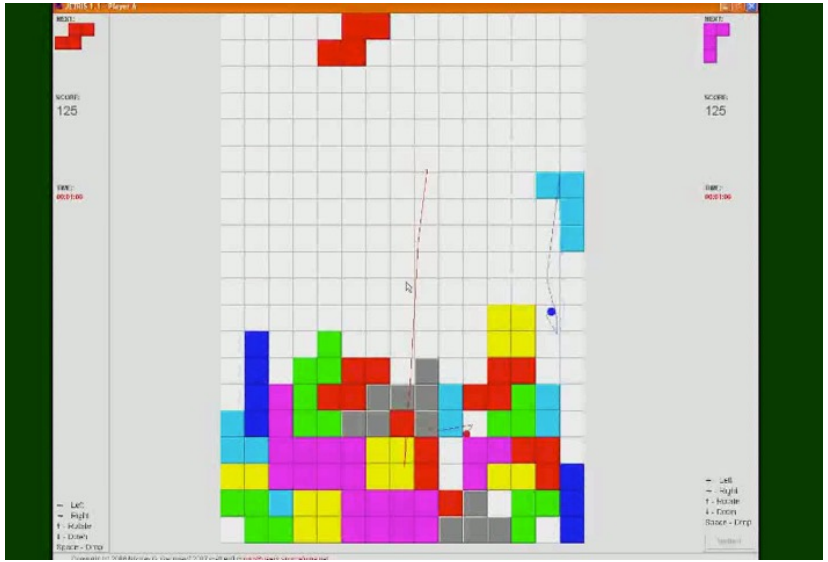
Sensing is not used, in this case, to get sufficient information inside the brain to allow it to solve the whole problem in one go

(Not: sense-**ponder**-run to catch ball)

Instead, **sensing is geared to acting, and acting is geared to sensing, in a rolling cycle**

You run so that the **visual input stays within certain bounds**, and that solves the problem..

Frequent saccades enable us to circumvent the need to build enduring and detailed models of our visual surroundings.



"We can use the world as its own best model and visit and rev

See Kirsh and Maglio's (1994)

The Point:

Sensing need not always be used **re-constructively** i.e. to prime a rich world-model that enables the agent to reason over the model **instead of** the world.

Instead, active systems can use sensing in more direct ways, as the fly-ball case suggests.

A lot of expert sports skill probably depends on these kinds of uses of sensing.

Second Moral of the Embodied Mind: **Cognition leaks into whole action-perception loops.**

The presence of a self-controlled, acting, sensing body allows an agent to **sculpt her own sensory input streams** in ways that promote learning, reasoning, and efficient problem-solving.

The engine of some forms of cognizing then isn't the naked brain - but rather whole complex system including brain and the sensing-acting-moving body.

Spreading into the **environment**…

The Bluefin Tuna Puzzle:

The aquatic performance of the Bluefin Tuna far exceeds its basic physical capacities.

e.g. they should not be Abel to go as fast, turn as sharp, or excel with such force…

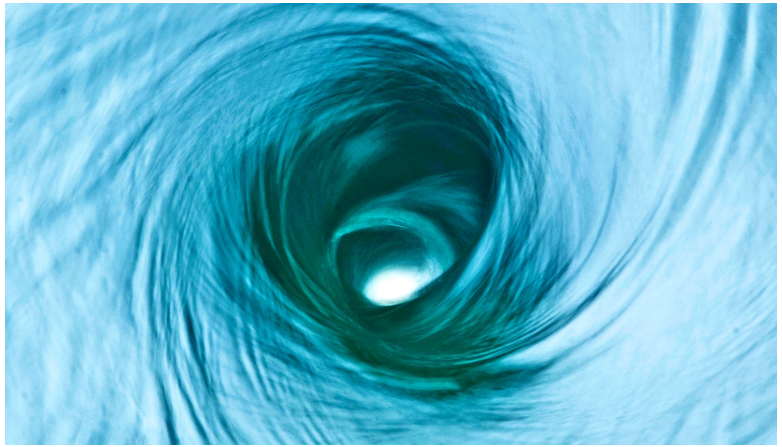
Too weak by a factor of 7

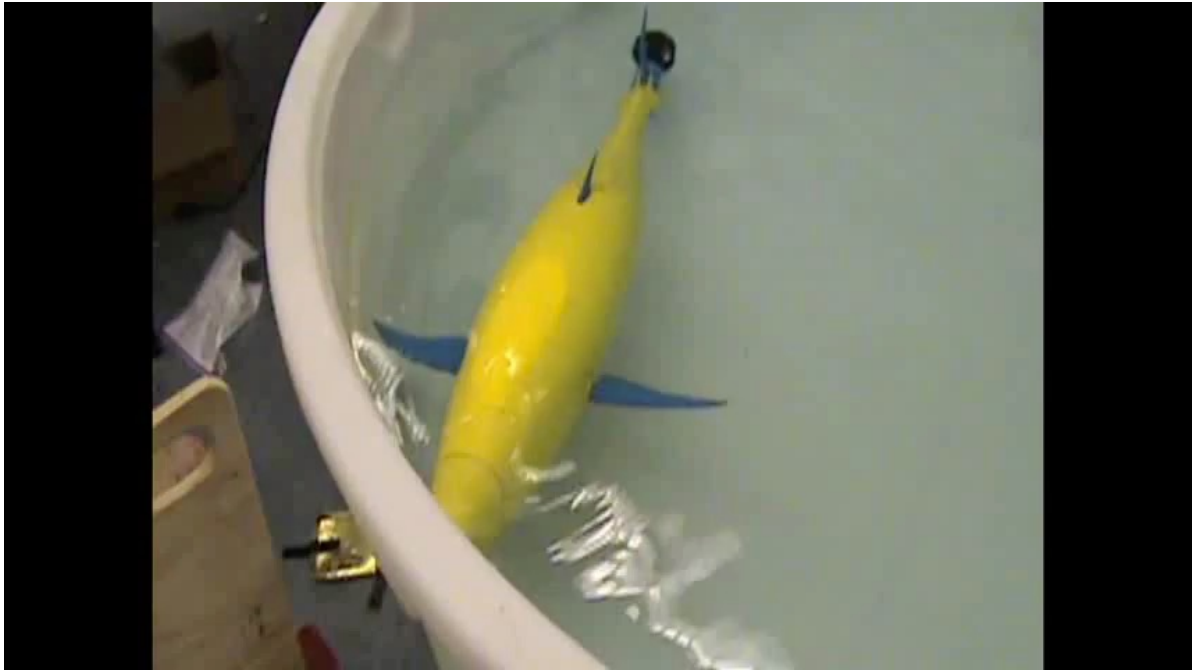


How does this work:

Tuna makes the most of the world around them, using naturally occurring eddies and cortices (e.g. where water hits a rock)

They use their own tail-flaps to create local vortices and pressure gradients that they then exploit for power.





Now we're creeping up on **higher cognition**…

We also **structure our worlds** to simplify our thinking.

Consider now our own artifact-enhanced cognitive processing

Classic Example: **Long Multiplication**

Most of us can do simple multiplications by learning a few rote patterns e.g

7 x 7 is 49

But what about

7466796 X 74965

Familiar but potent trick is to **combine** simple neural pattern completion strategies and the **use and manipulation of external symbols and media** e.g. pen and paper – reducing the complex problem to a set of simpler pattern-completing operations.

e.g. physically shuffling the **scrabble** tiles to prompt word recall.



Famous exchange between **Richard Feynman** (the Nobel laureate physicist) and the historian **Charles Weiner**

> " Weiner once remarked casually that [a batch of notes and sketches] represented "a record of [Feynman's] day-to-day work," and Feynman reacted sharply.
>
> "I actually did the work on the paper," he said.
>
> "Well," Weiner said, "the work was done in your head, but the record of it is still here."
>
> "No, it's not a *record*, not really. It's *working*. You have to work on paper and this is the paper. Okay?" "

As quoted in *Genius* (Gleick's biography of Feynman)

Stronger Claim: The Embodied/Extended Mind

Its not that all the real thinking happens inside the head and loops out into the body, and external media is just for offloading or storage

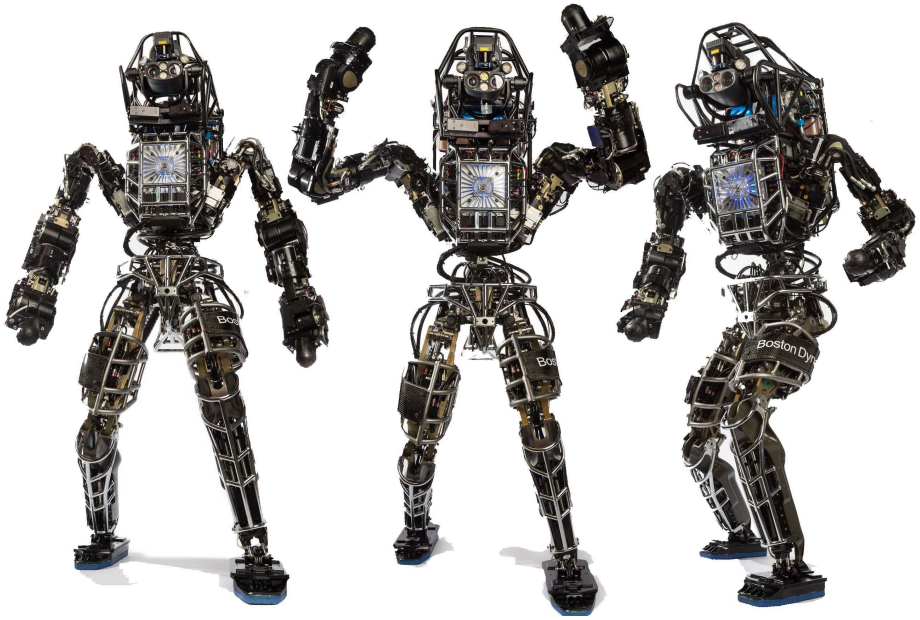
But all of these things **become part of thinking**

Questions to ponder:

Can high-level reason do without **richly reconstructive inner world-models**?

What **role** (if any) does embodiment play in cases of higher-level reasoning?

In what ways might the presence of **spoken language** and other cognitive technologies transform the tasks involved in higher-level reasoning?



The End

Thanks!



Dr. Mark Miller

MONASH University
UNIVERSITY OF TORONTO
CHAIN