A Self-Modeling Quadruped Robot

- Work by Josh Bongard, Victor Zykov, and Hod Lipson, Cornell University (2006)
- “Starfish” robot autonomously learns about its own body
- Develops an internal self-model using a genetic algorithm
- Learns to walk based on self-model
- Can recover from damage

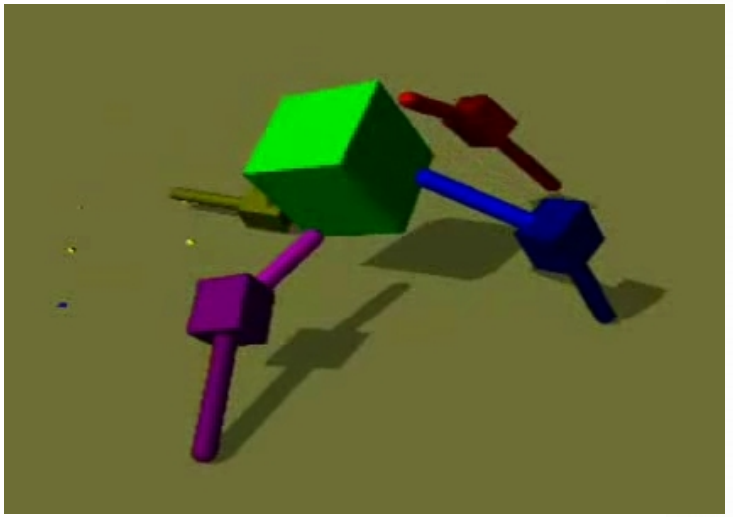

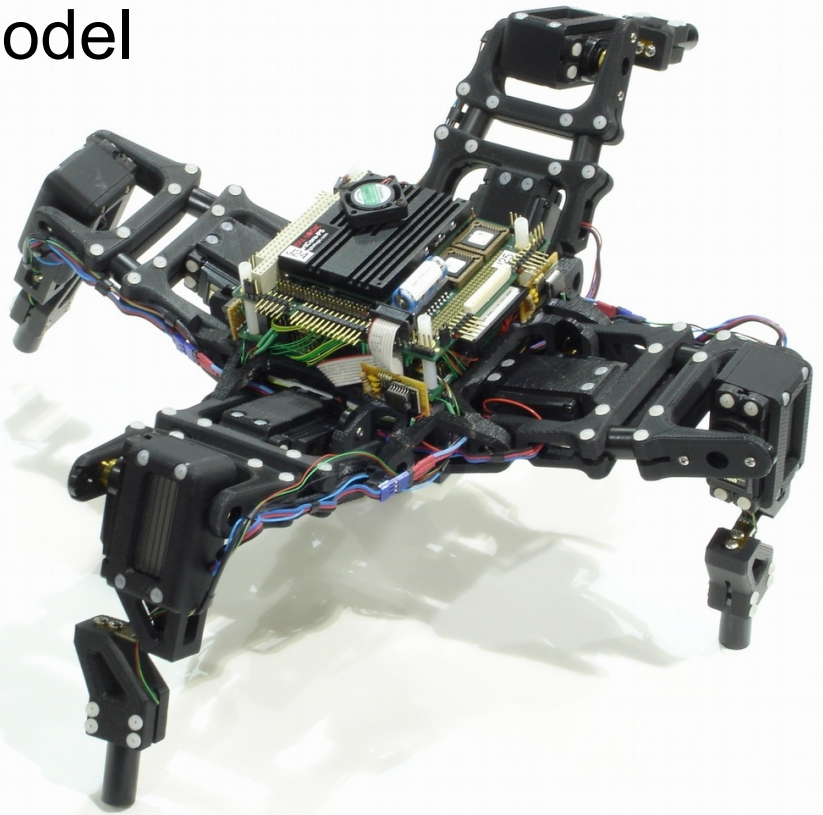

A Self-Modeling Quadruped Robot

- 9 body parts:
  - central body, 4 legs, 2 parts per leg
  - 2 body tilt sensors (left-right, forward-back)
  - 8 motorized joints
  - 8 joint angle sensors

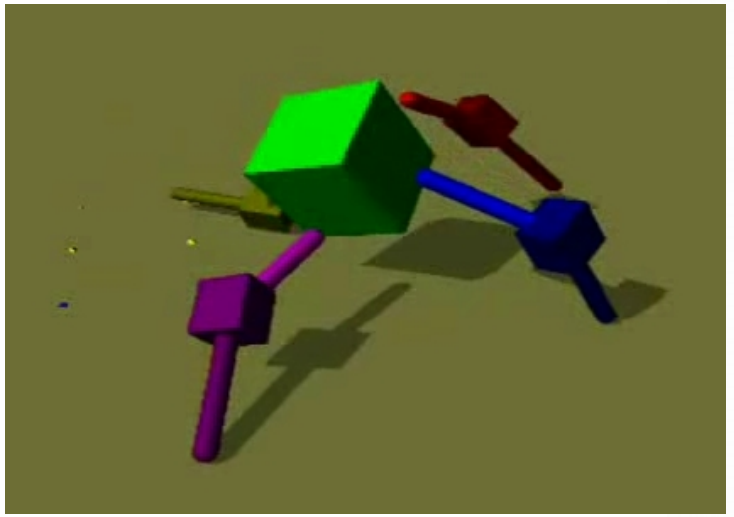

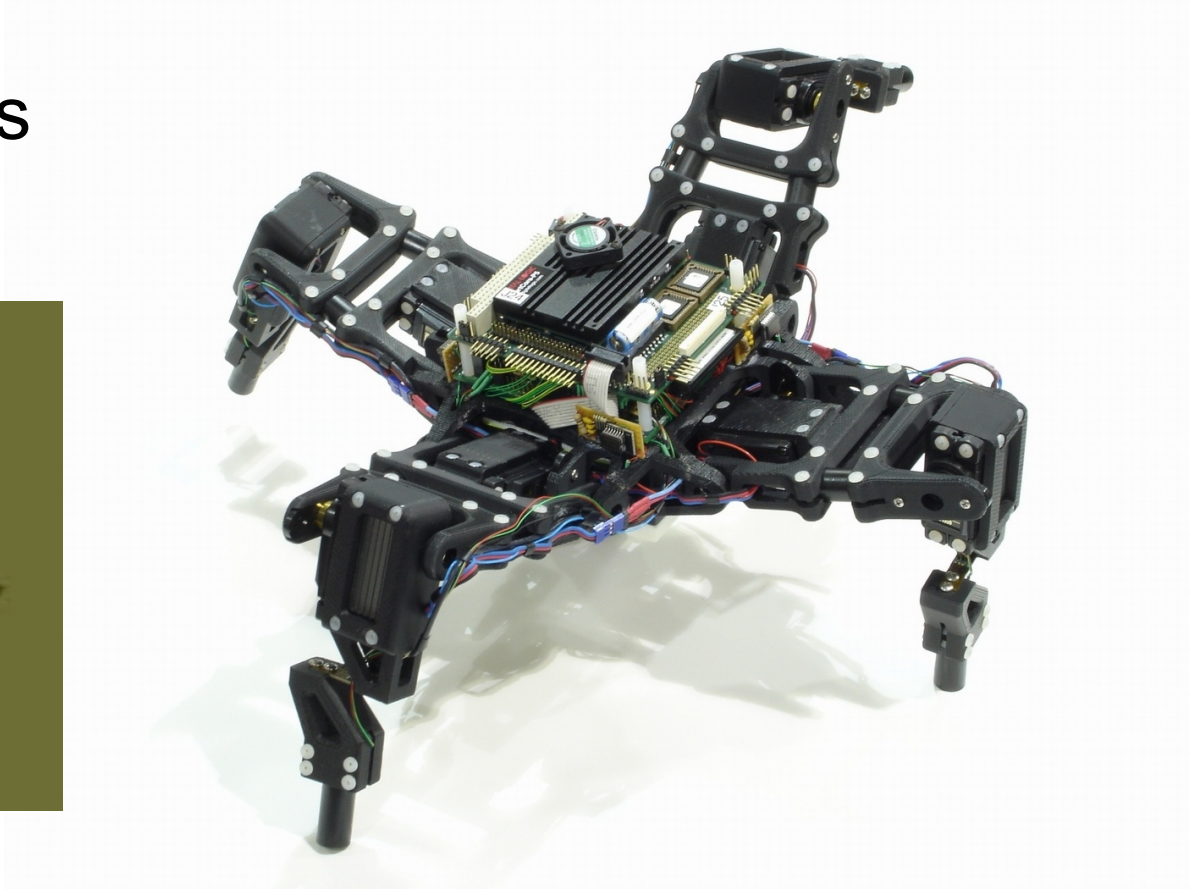

Self-Models

- Robot knows it has 9 body parts, but does NOT know which part is attached to which
- A self-model consists of **16 numerical parameters** that specify how leg parts are attached to each other
- Robot maintains 15 competing self-models internally

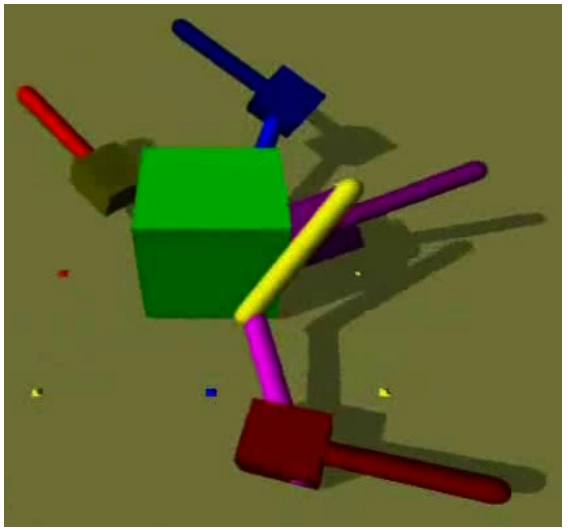

A poor self-model

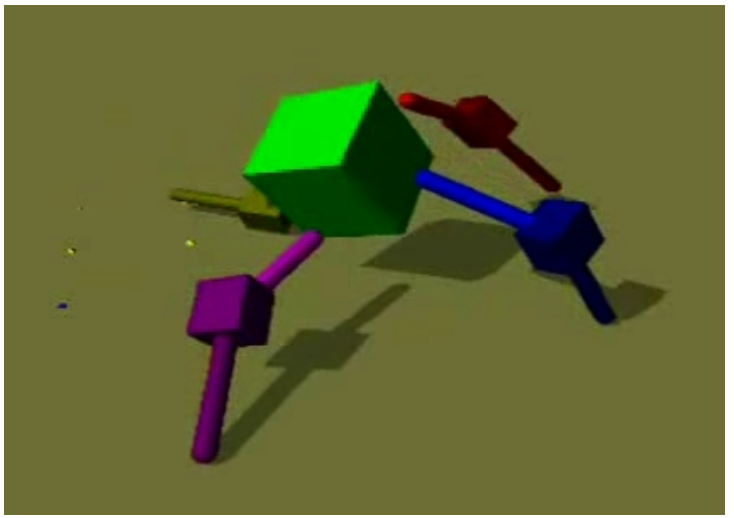

A good self-model

Stage I: Learning About Itself

(a) Robot performs a randomly chosen motor action

(b) 15 internal self-models are randomly created

(c) Self-models evolve for 200 cycles. Fitness: ability to predict observed results of robot's actions so far

(d) A new action is chosen that **maximizes disagreement** among the models about that action's predicted effect

(e) Robot performs the action and observes the result

Stage I: Learning About Itself

(a) Robot performs a randomly chosen motor action

(b) 15 internal self-models are randomly created

(c) Self-models evolve for 200 cycles. Fitness: ability to predict observed results of robot's actions so far


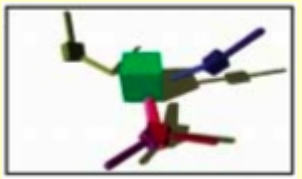
. . .


(d) A new action is chosen that **maximizes disagreement** among the models about that action's predicted effect

(e) Robot performs the action and observes the result

Steps (c)-(e) are repeated for 16 different exploratory actions

Stage I: Results

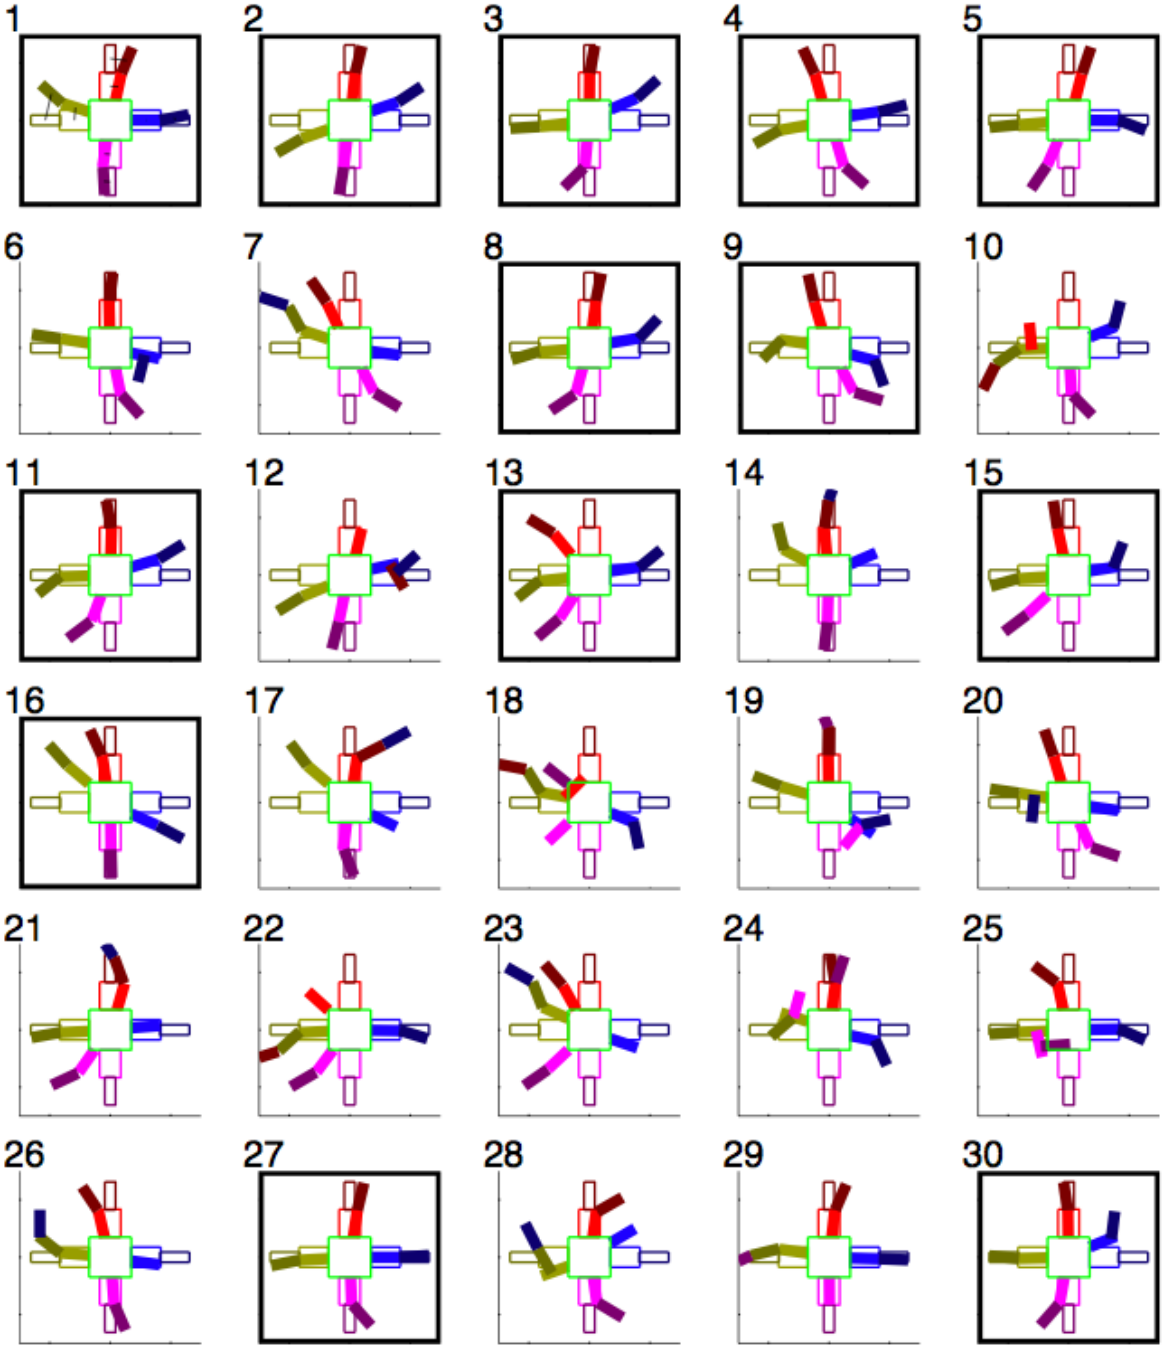

Stage II: Learning To Walk

- Start with the most accurate **self-model** generated in Stage I
- Use a **genetic algorithm** to evolve a population of behaviors
  - a behavior is a set of numbers specifying a temporal sequence of joint angles
  - fitness of a behavior: how far does it cause the self-model to move forward?
- Run the **best evolved behavior** on the physical robot

Stage II: Results

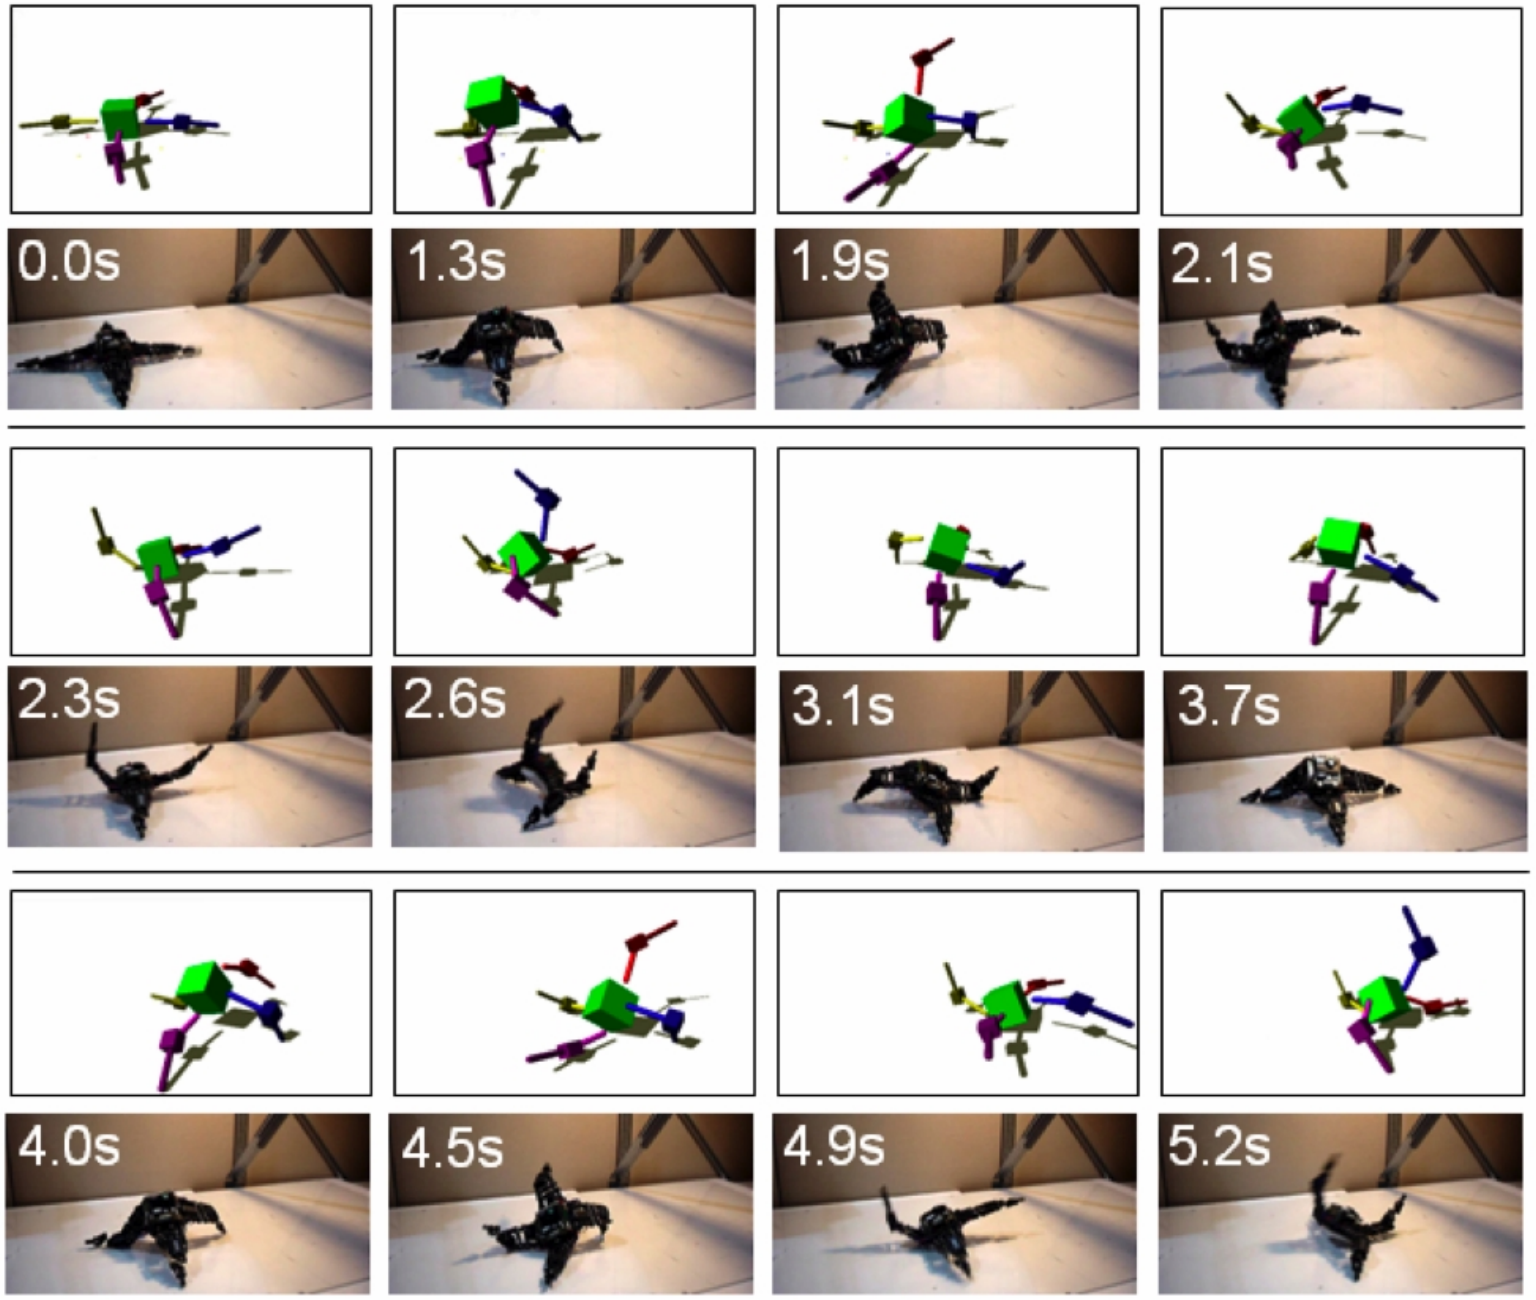

Damage Recovery

- If the robot becomes damaged, self-modeling and experimentation recommence with the best model so far
- The GA varies the relative sizes of leg parts within models until a new, more accurate self-model emerges
- The new, improved self-model is used to synthesize a new walking behavior

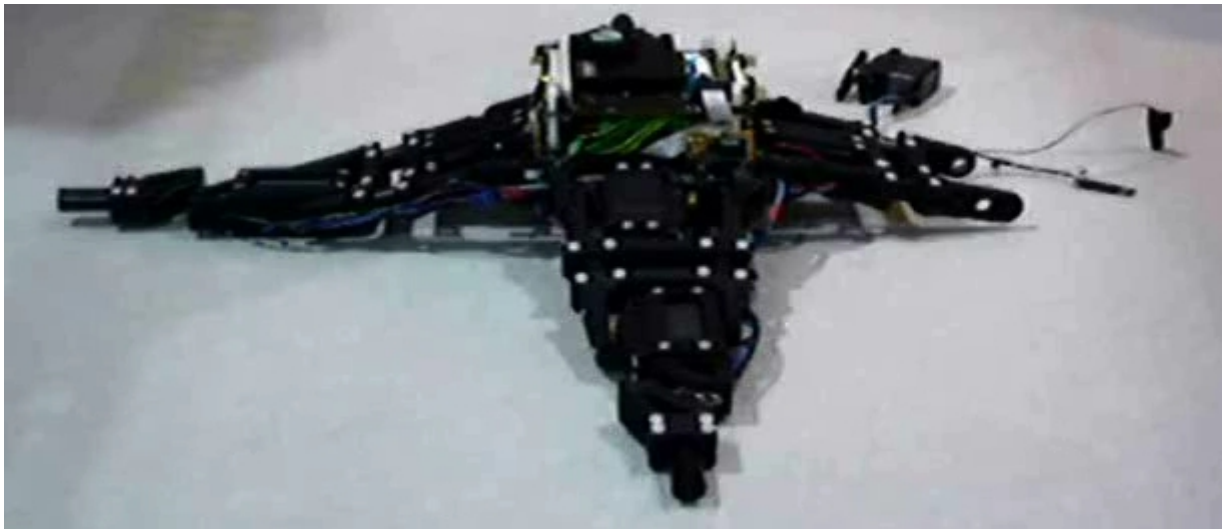

Damage Recovery: Results

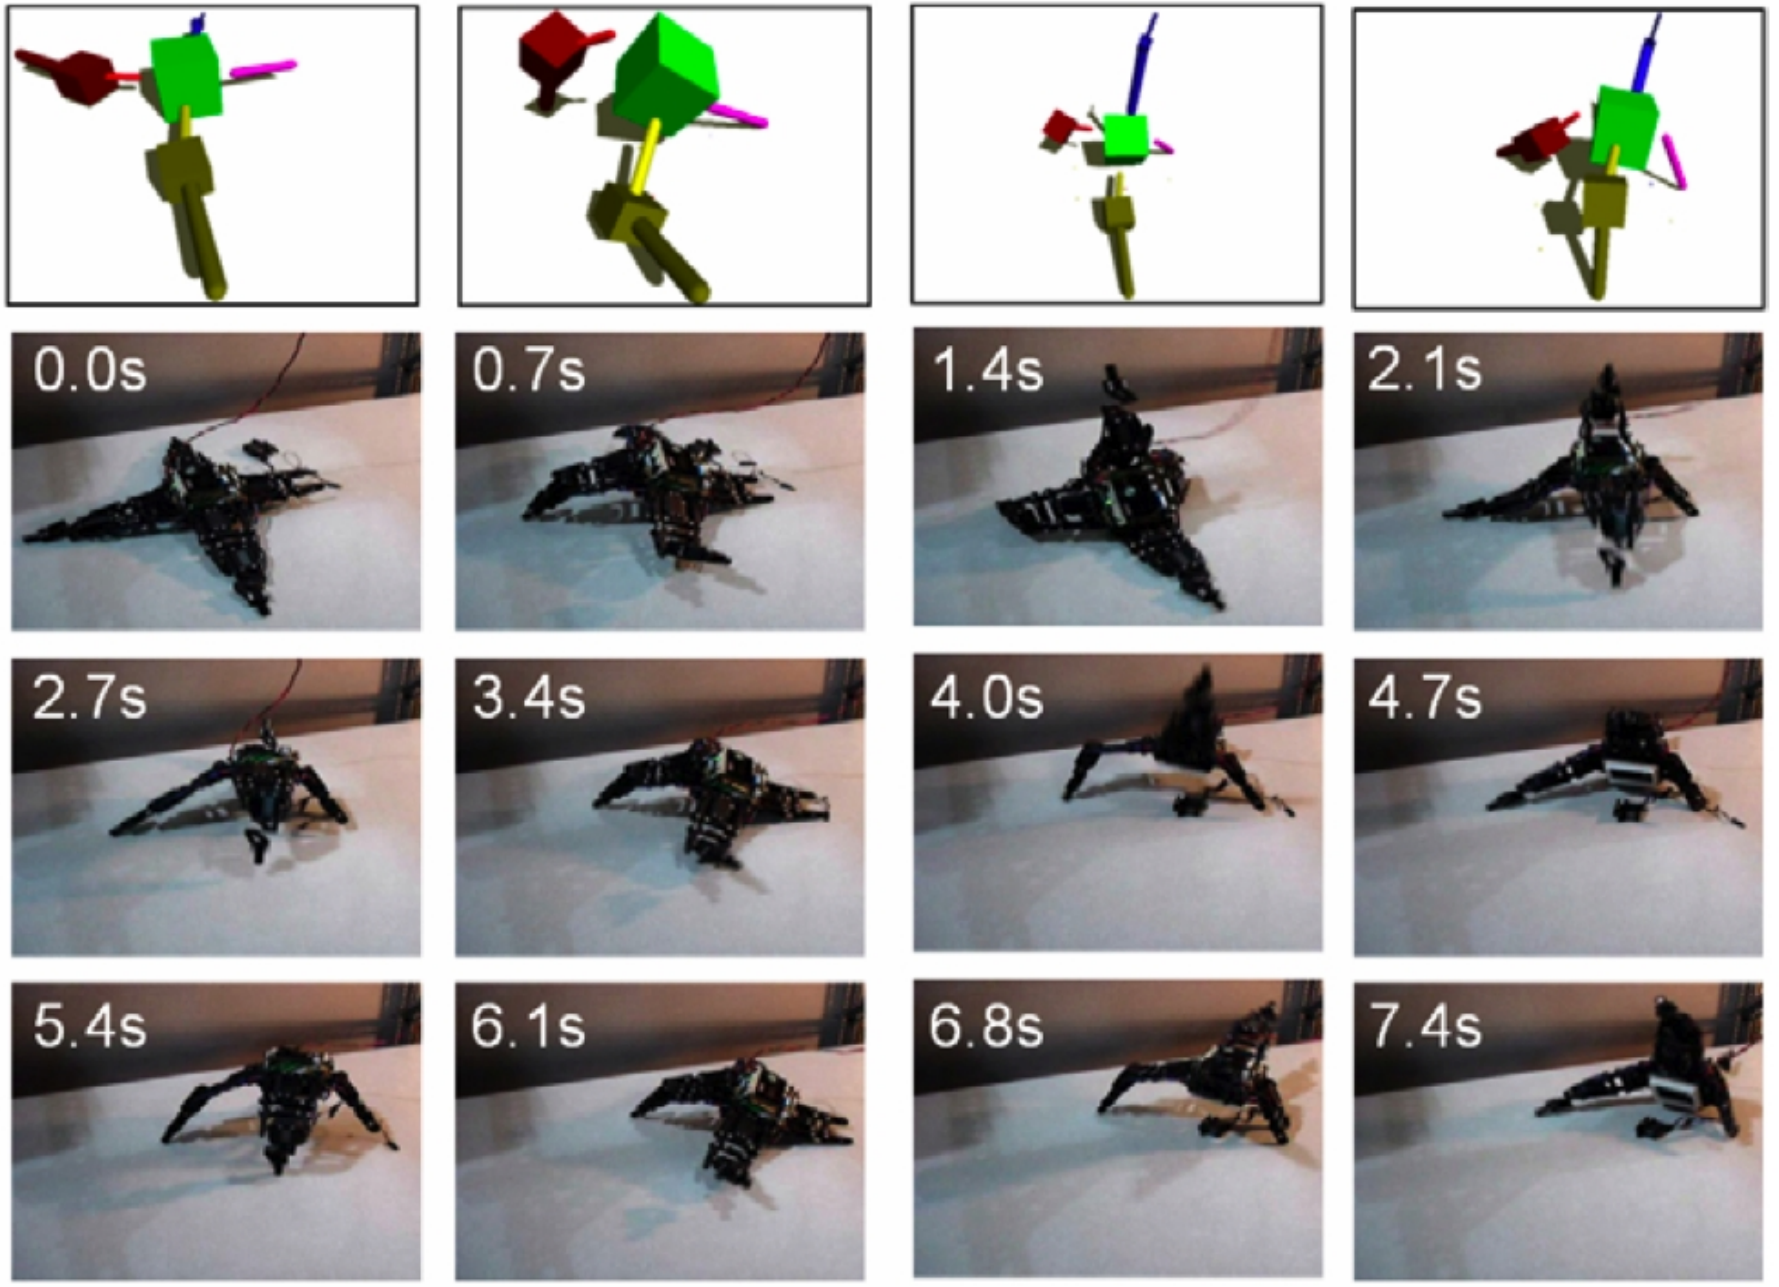

Feelin' Good!