Micro-data

the challenge of robot learning

Jean-Baptiste Mouret (Inria)


DARPA Robotics Challenge, 20
06/2015 UTC

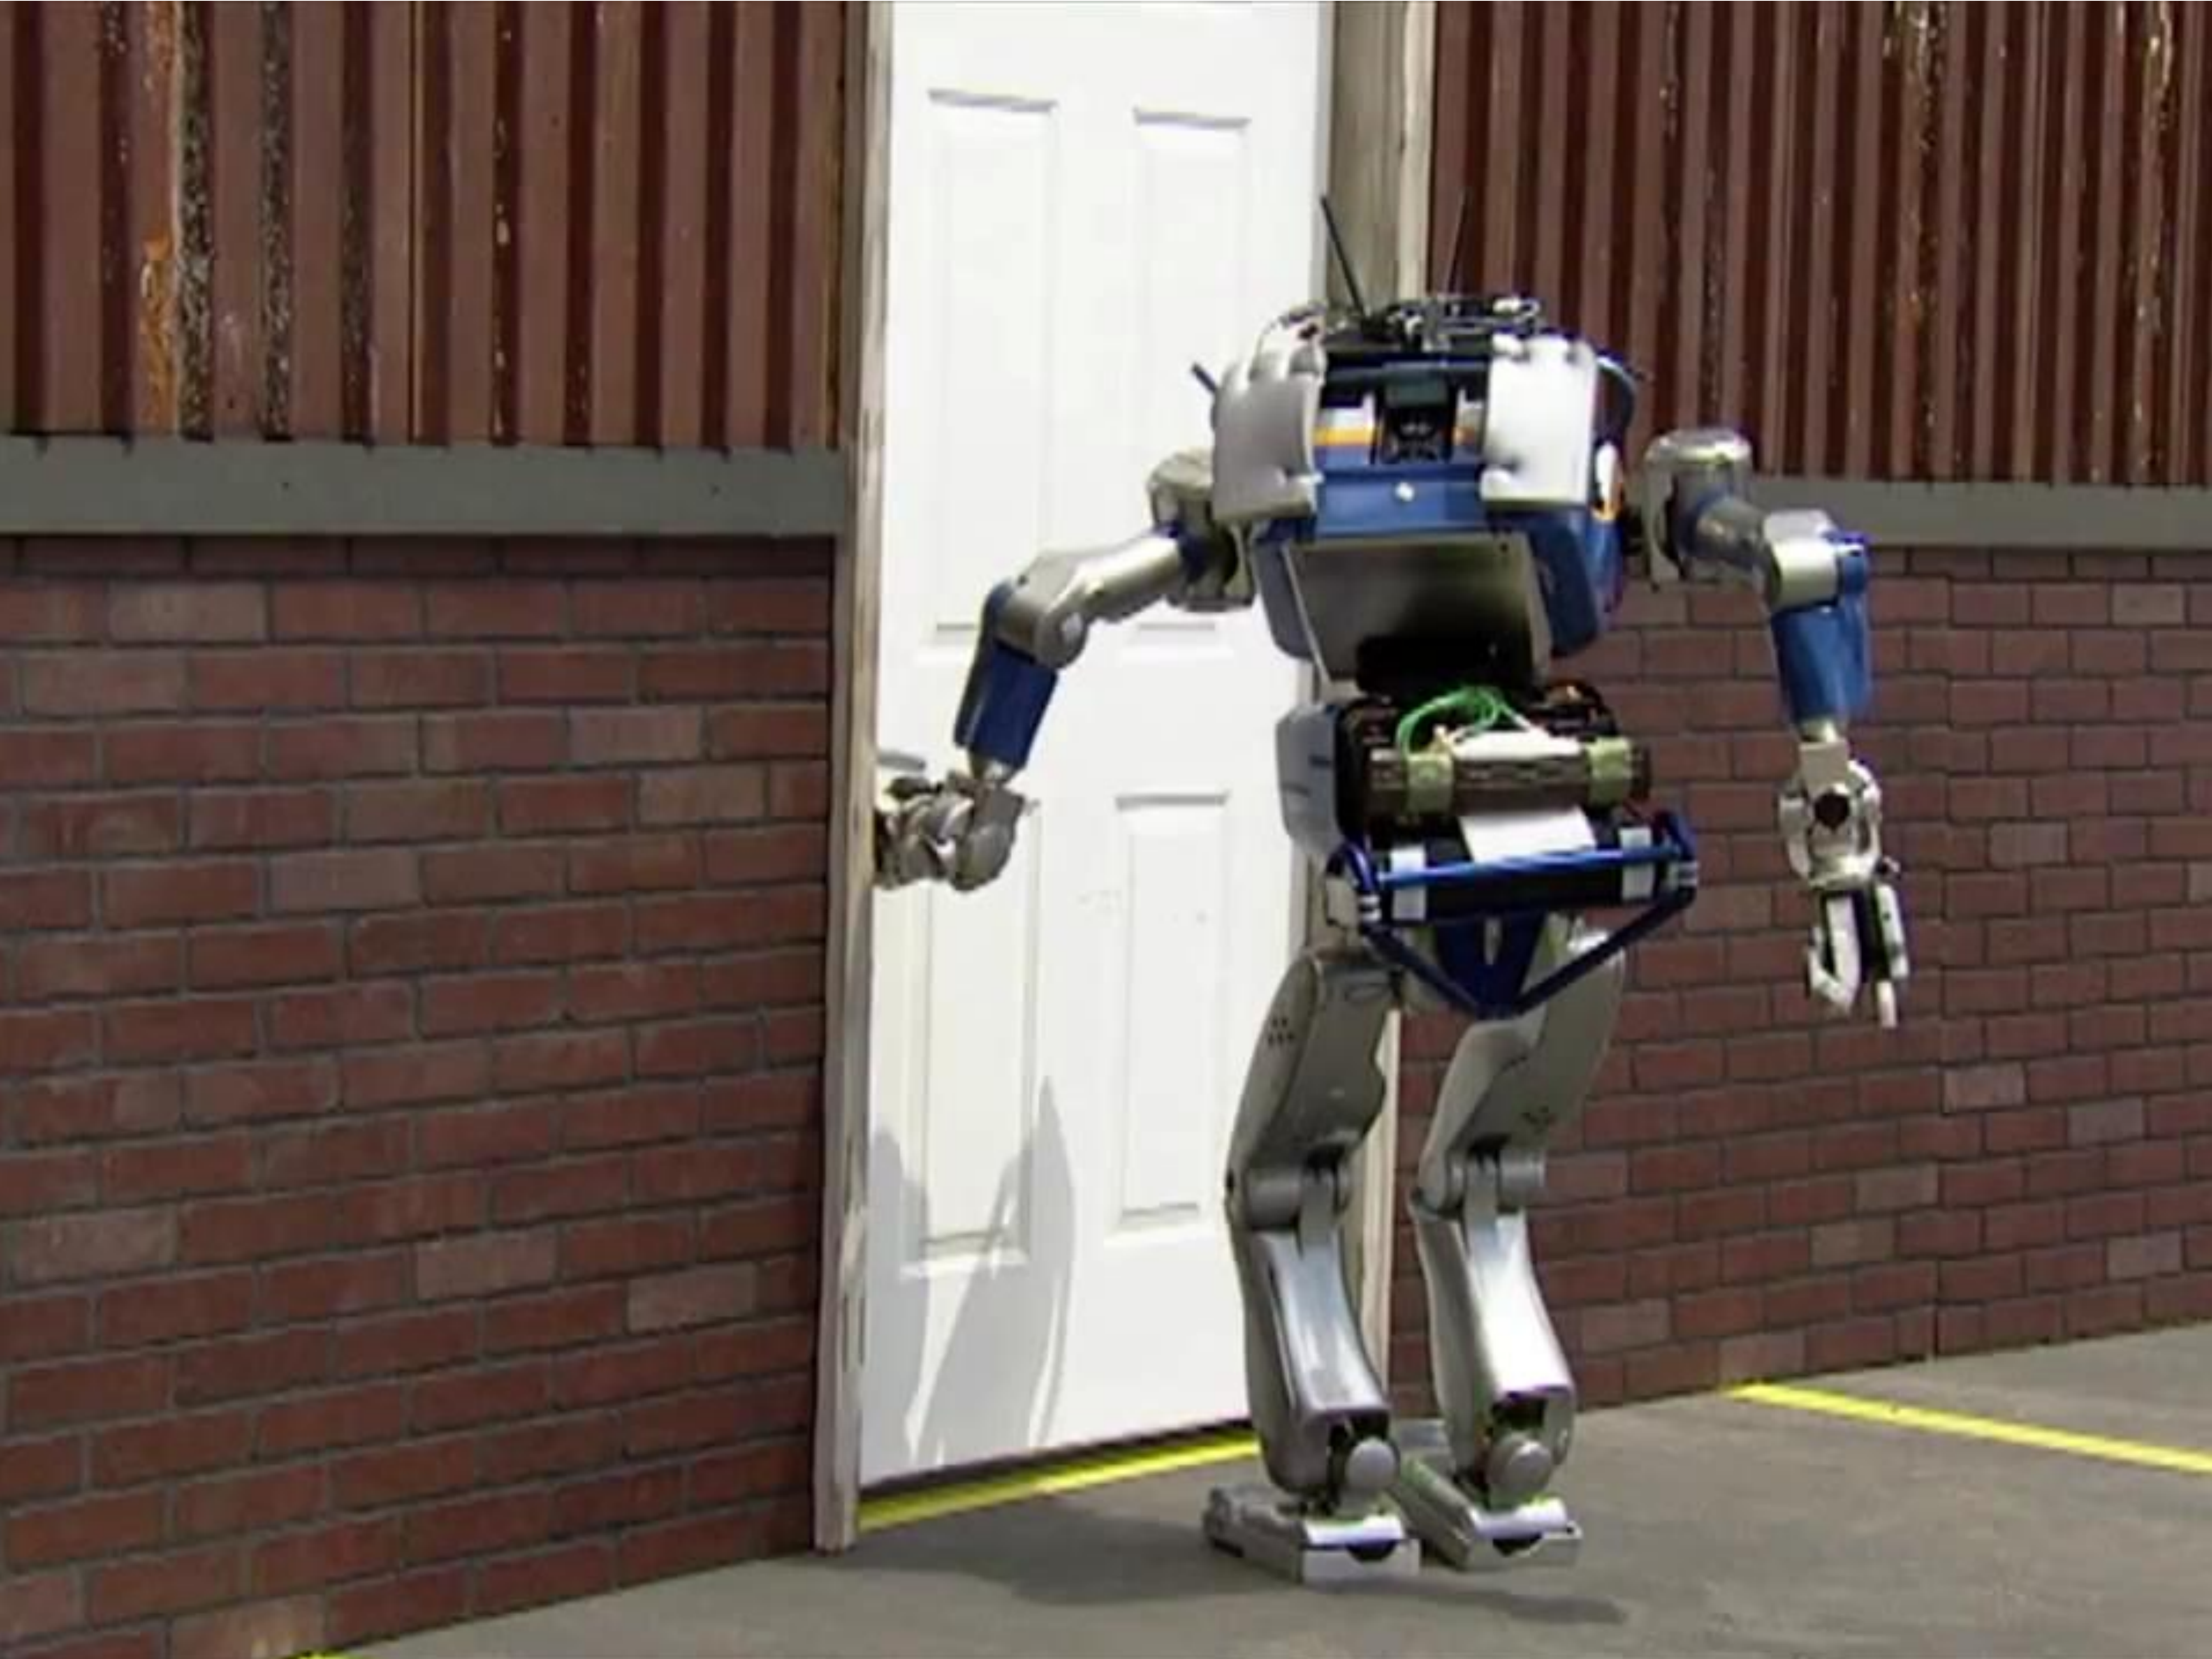

The issue with current robots is not that they fail…

… it is that they do not get back on their feet and try again

- ➡ Current robots do not deal well with unexpected situations (e.g., damage)
- ➡ Current robots do not learn from their mistakes
- ➡ Main reason: diagnosis / understanding problems is hard!

Micro-data: the other end of the spectrum

AlphaGo

4.9 million games
(self-play) / 40 days

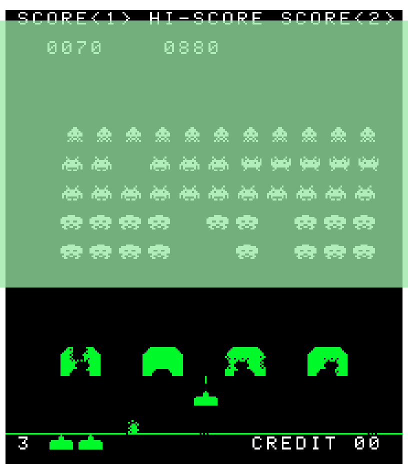

Atari games :
3[illegible] days

Robots need micro-data learning

Amount of data

« Big-Data » Simulated world

Deep learning?
Evolution?

Real world

« Micro-data »

1-20 trials
A few minutes

?

Mouret, JB. (2016) "Micro-Data Learning: The Other End of the Spectrum." ERCIM News.

Micro-data policy search

Parameters of the policy

Optimize: $$J(\boldsymbol{\theta}) = \mathbb{E}\left[\sum_{t=1}^{T} r(\mathbf{x}_t) \middle| \boldsymbol{\theta}\right]$$

Objective

Reward for state x_t

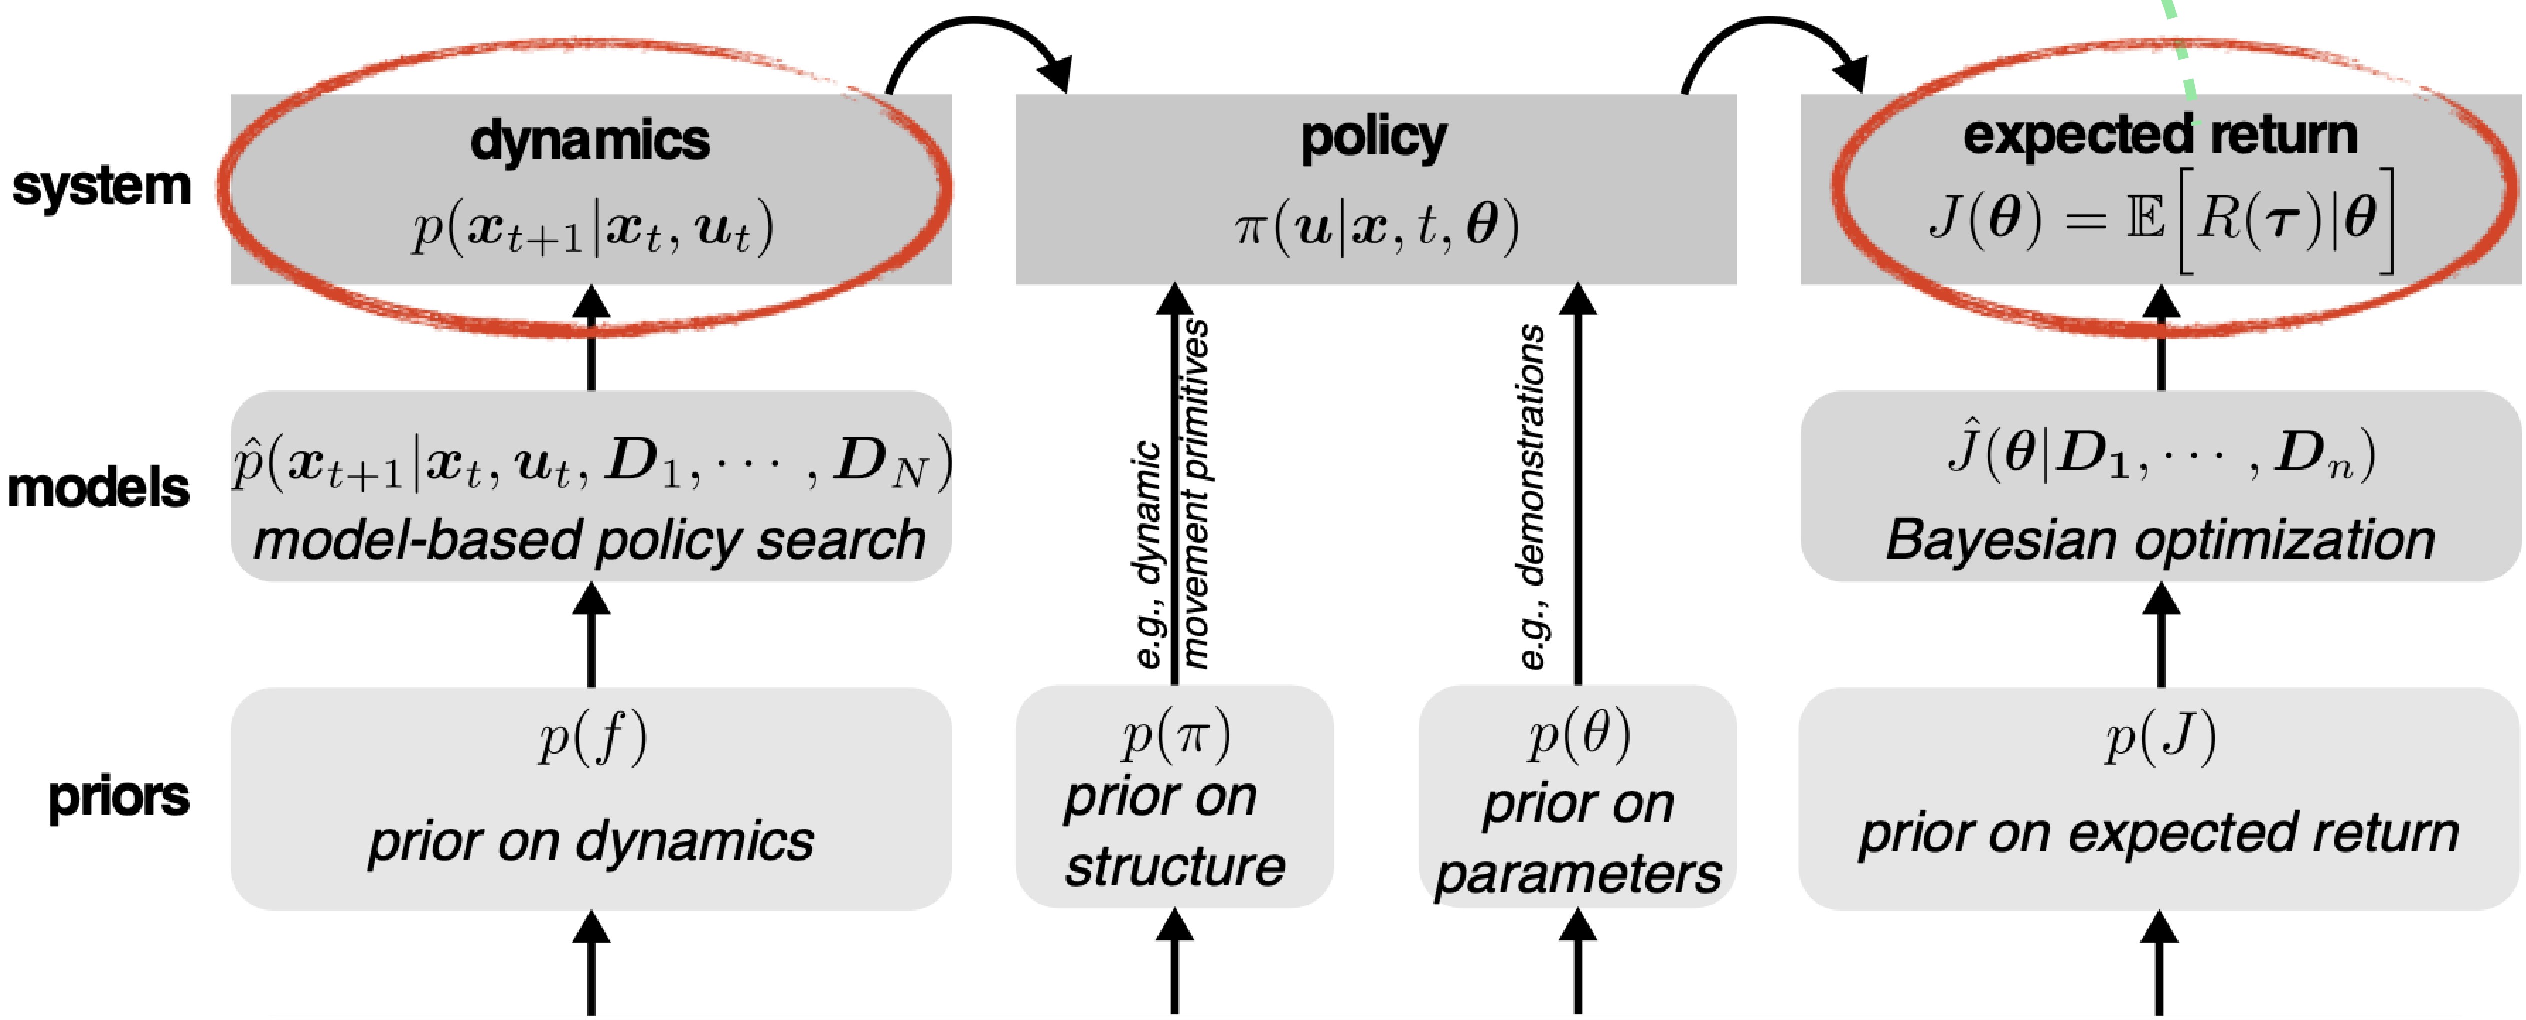

In nature: evolution & experience

K. Chatzilygeroudis, V. Vassiliades, F. Stulp, S. Calinon, J.-B. Mouret. (2018). A survey on policy search algorithms for learning robot controllers in a handful of trials. arXiv:1807.02303

Strategy 1: Learning the dynamical model

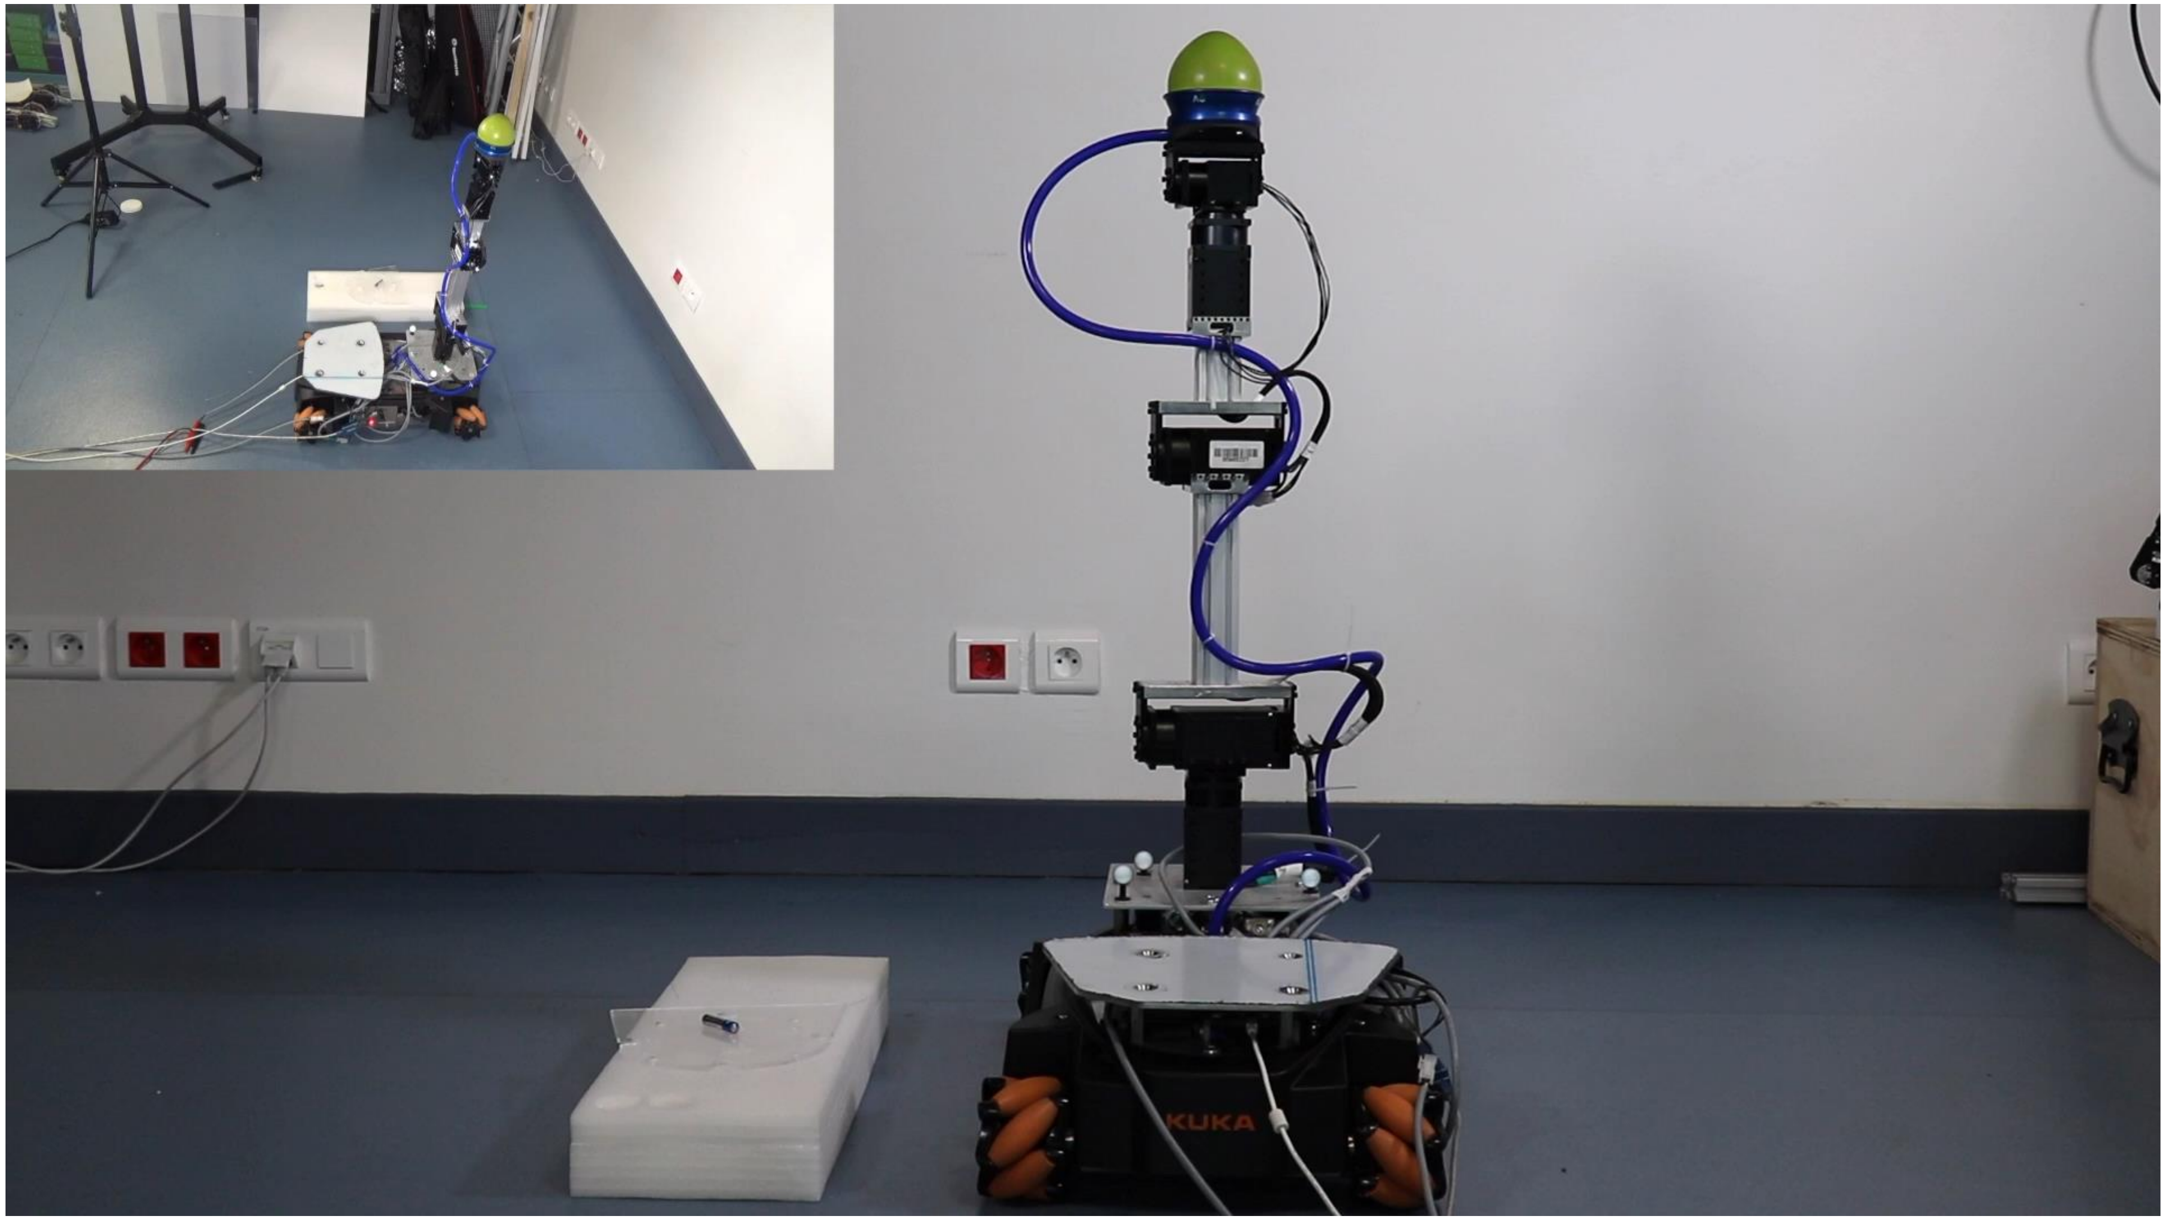

$$\mathbf{x}_{t+1} = \mathbf{x}_t + f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}$$

next state — state — dynamics — noise

Probabilistic model → **Optimal policy (in the model)** → **Probabilistic model**

Chatzilygeroudis K, Rama R, Kaushik R, Goepp D, Vassiliades V, Mouret JB. (2017) Black-Box Data-efficient Policy Search for Robotics. Proc. of IEEE IROS.

Deisenroth, M. P., Fox, D., & Rasmussen, C. E. (2015). Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2), 408-423.

Policy Search for damage recovery & adaptation

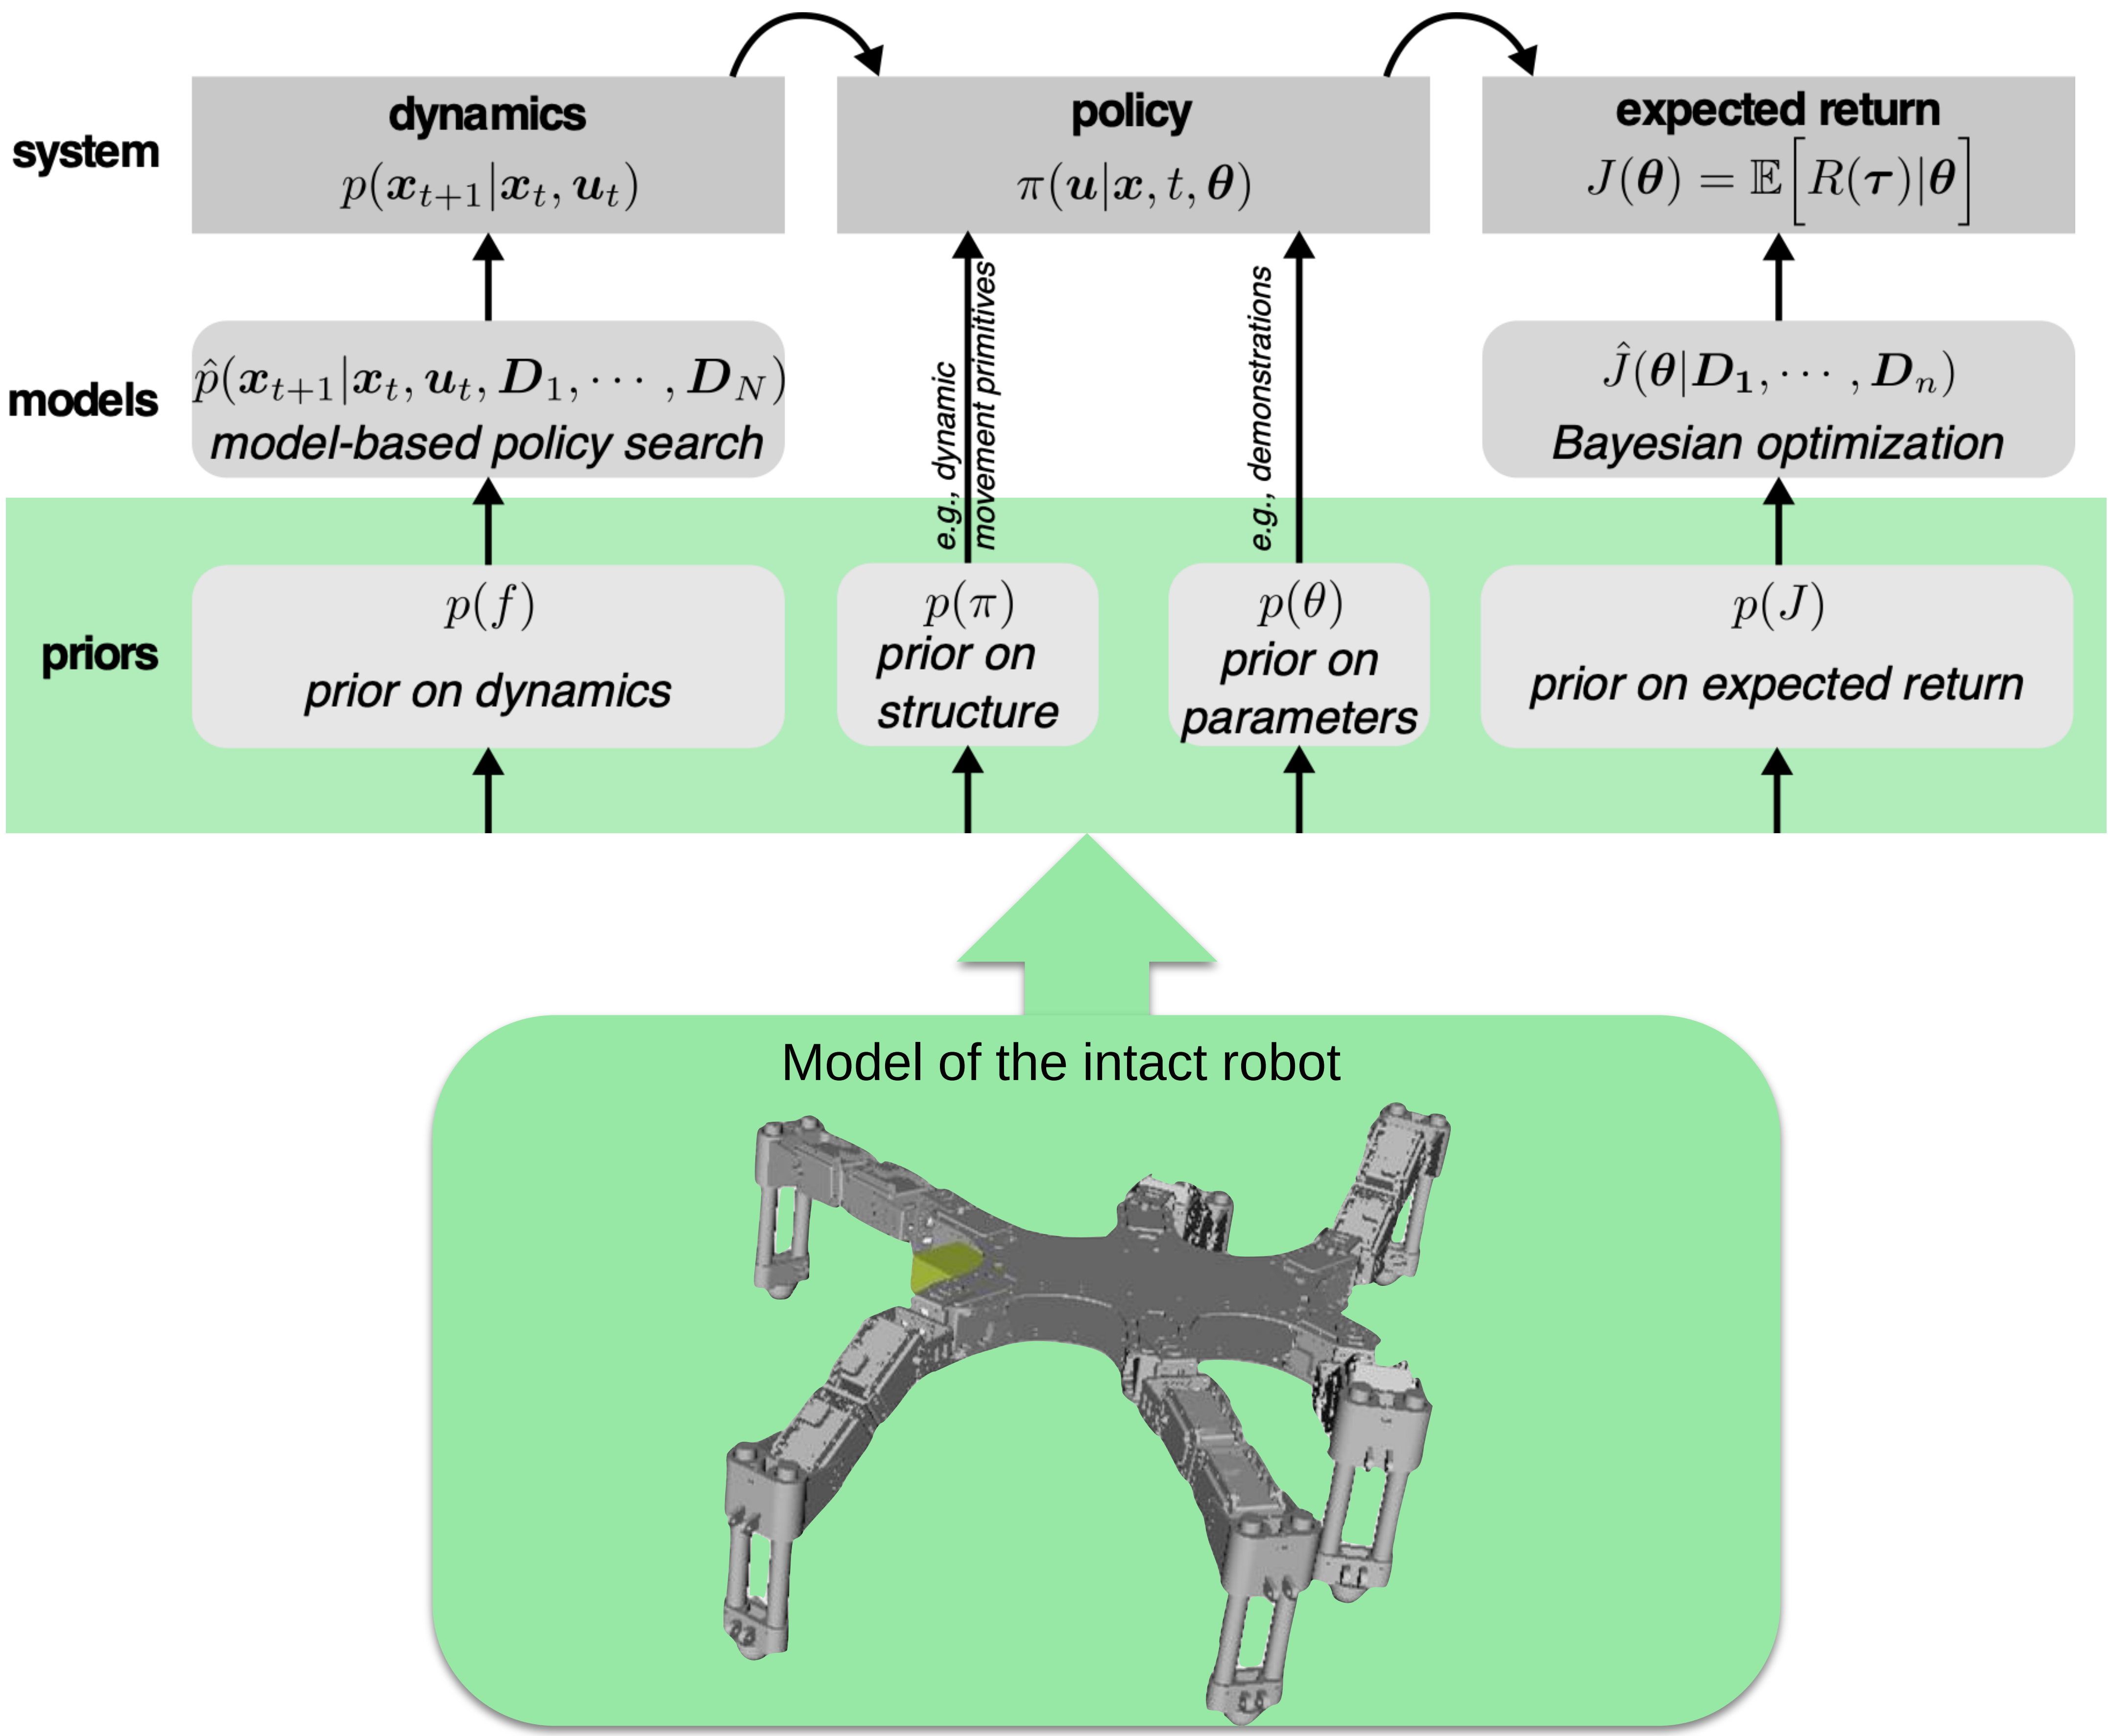

Black-DROPS + priors + identification

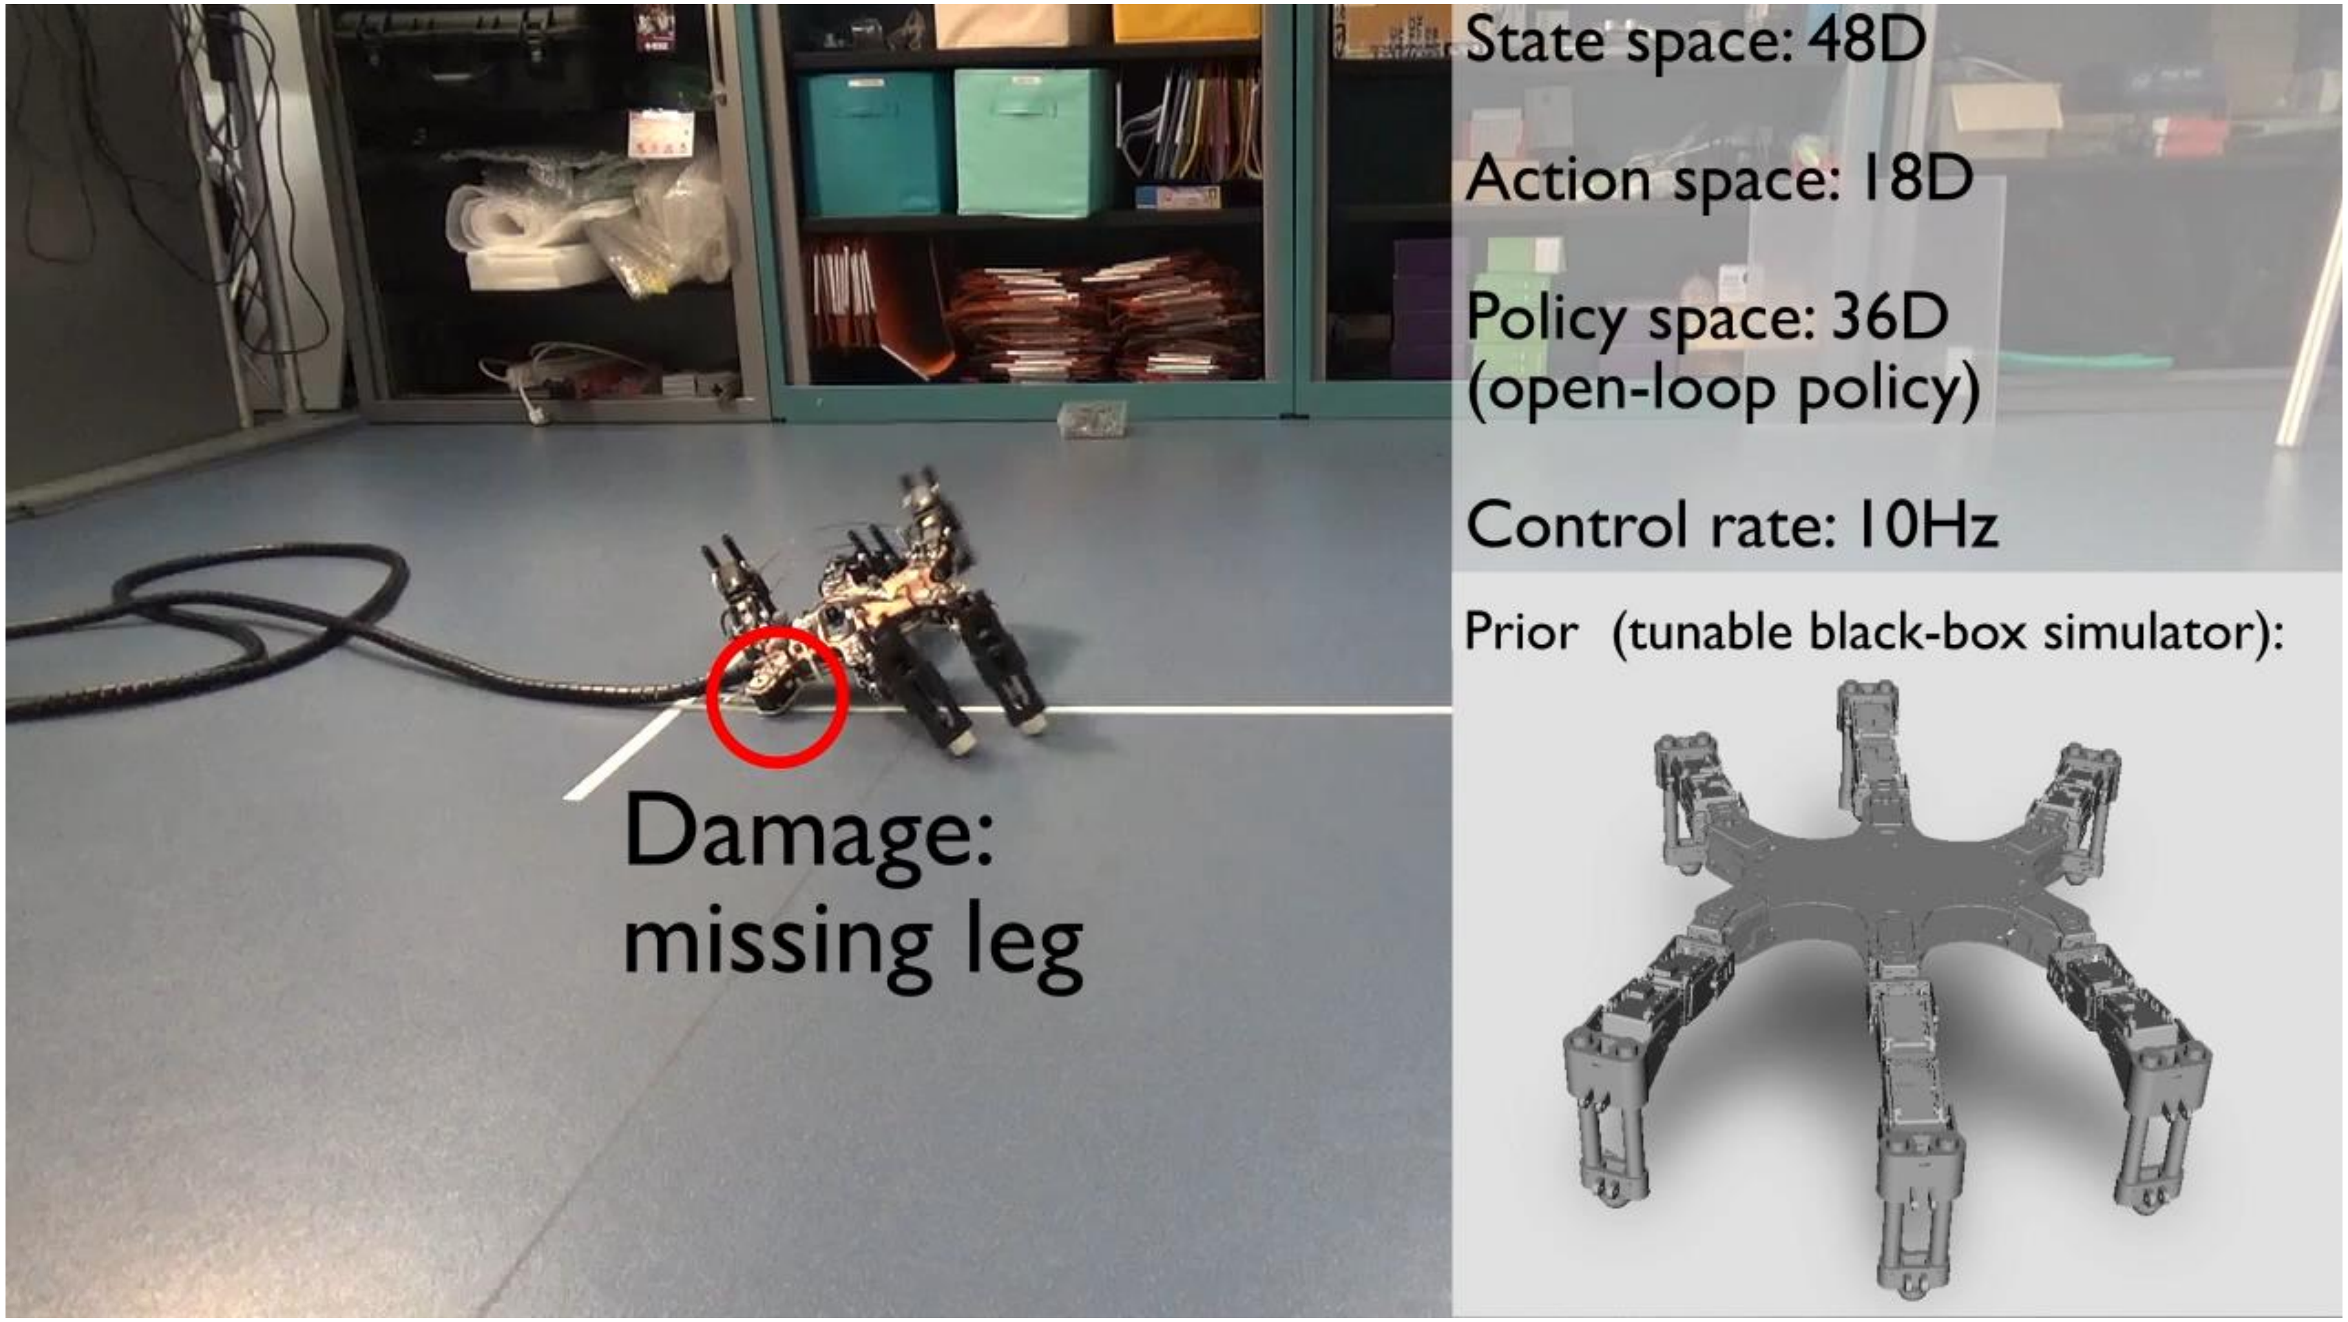

$$\mathbf{x}_{t+1} = \mathbf{x}_t + M(\mathbf{x}_t, \mathbf{u}_t, \phi_M) + f(\mathbf{x}_t, \mathbf{u}_t, \phi_K) + \mathbf{w}$$

next state — Simulator / model — Other dynamics

probabilistic model → policy

Chatzilygeroudis K, Mouret J.-B. (2018) Using Parameterized Black-Box Priors to Scale Up Model-Based Policy Search for Robotics. Proc. of ICRA.

Strategy 2: modelling the return + prior

The MAP-Elites algorithm generates the search space (prior)

- in simulation, with an intact robot
- many evaluations [simulation]
- "take the needles out of the haystack"
- provide an expected performance for the "needles"

Prior-based Bayesian optimization does the online learning

- search only among good solutions ("needles")
- fast trial-and-error (thanks to pre-computed "needles")
- few evaluations [real robot]

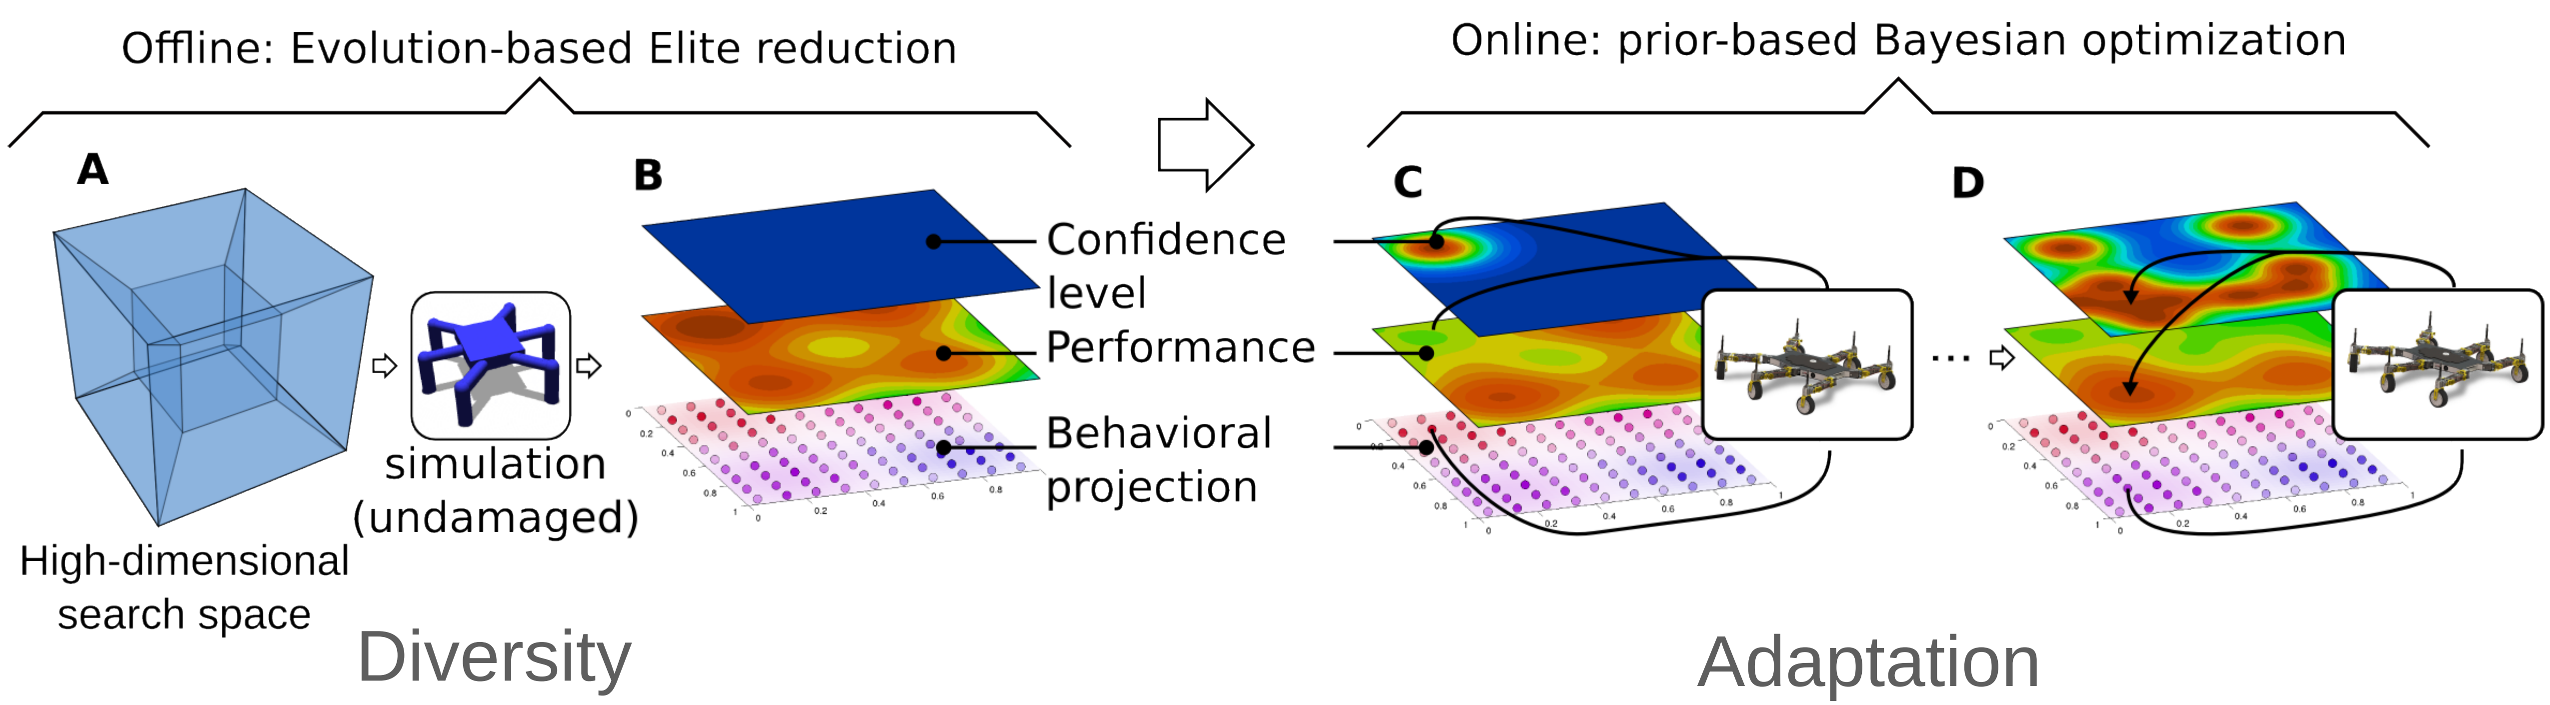

Strategy 2: computing a prior

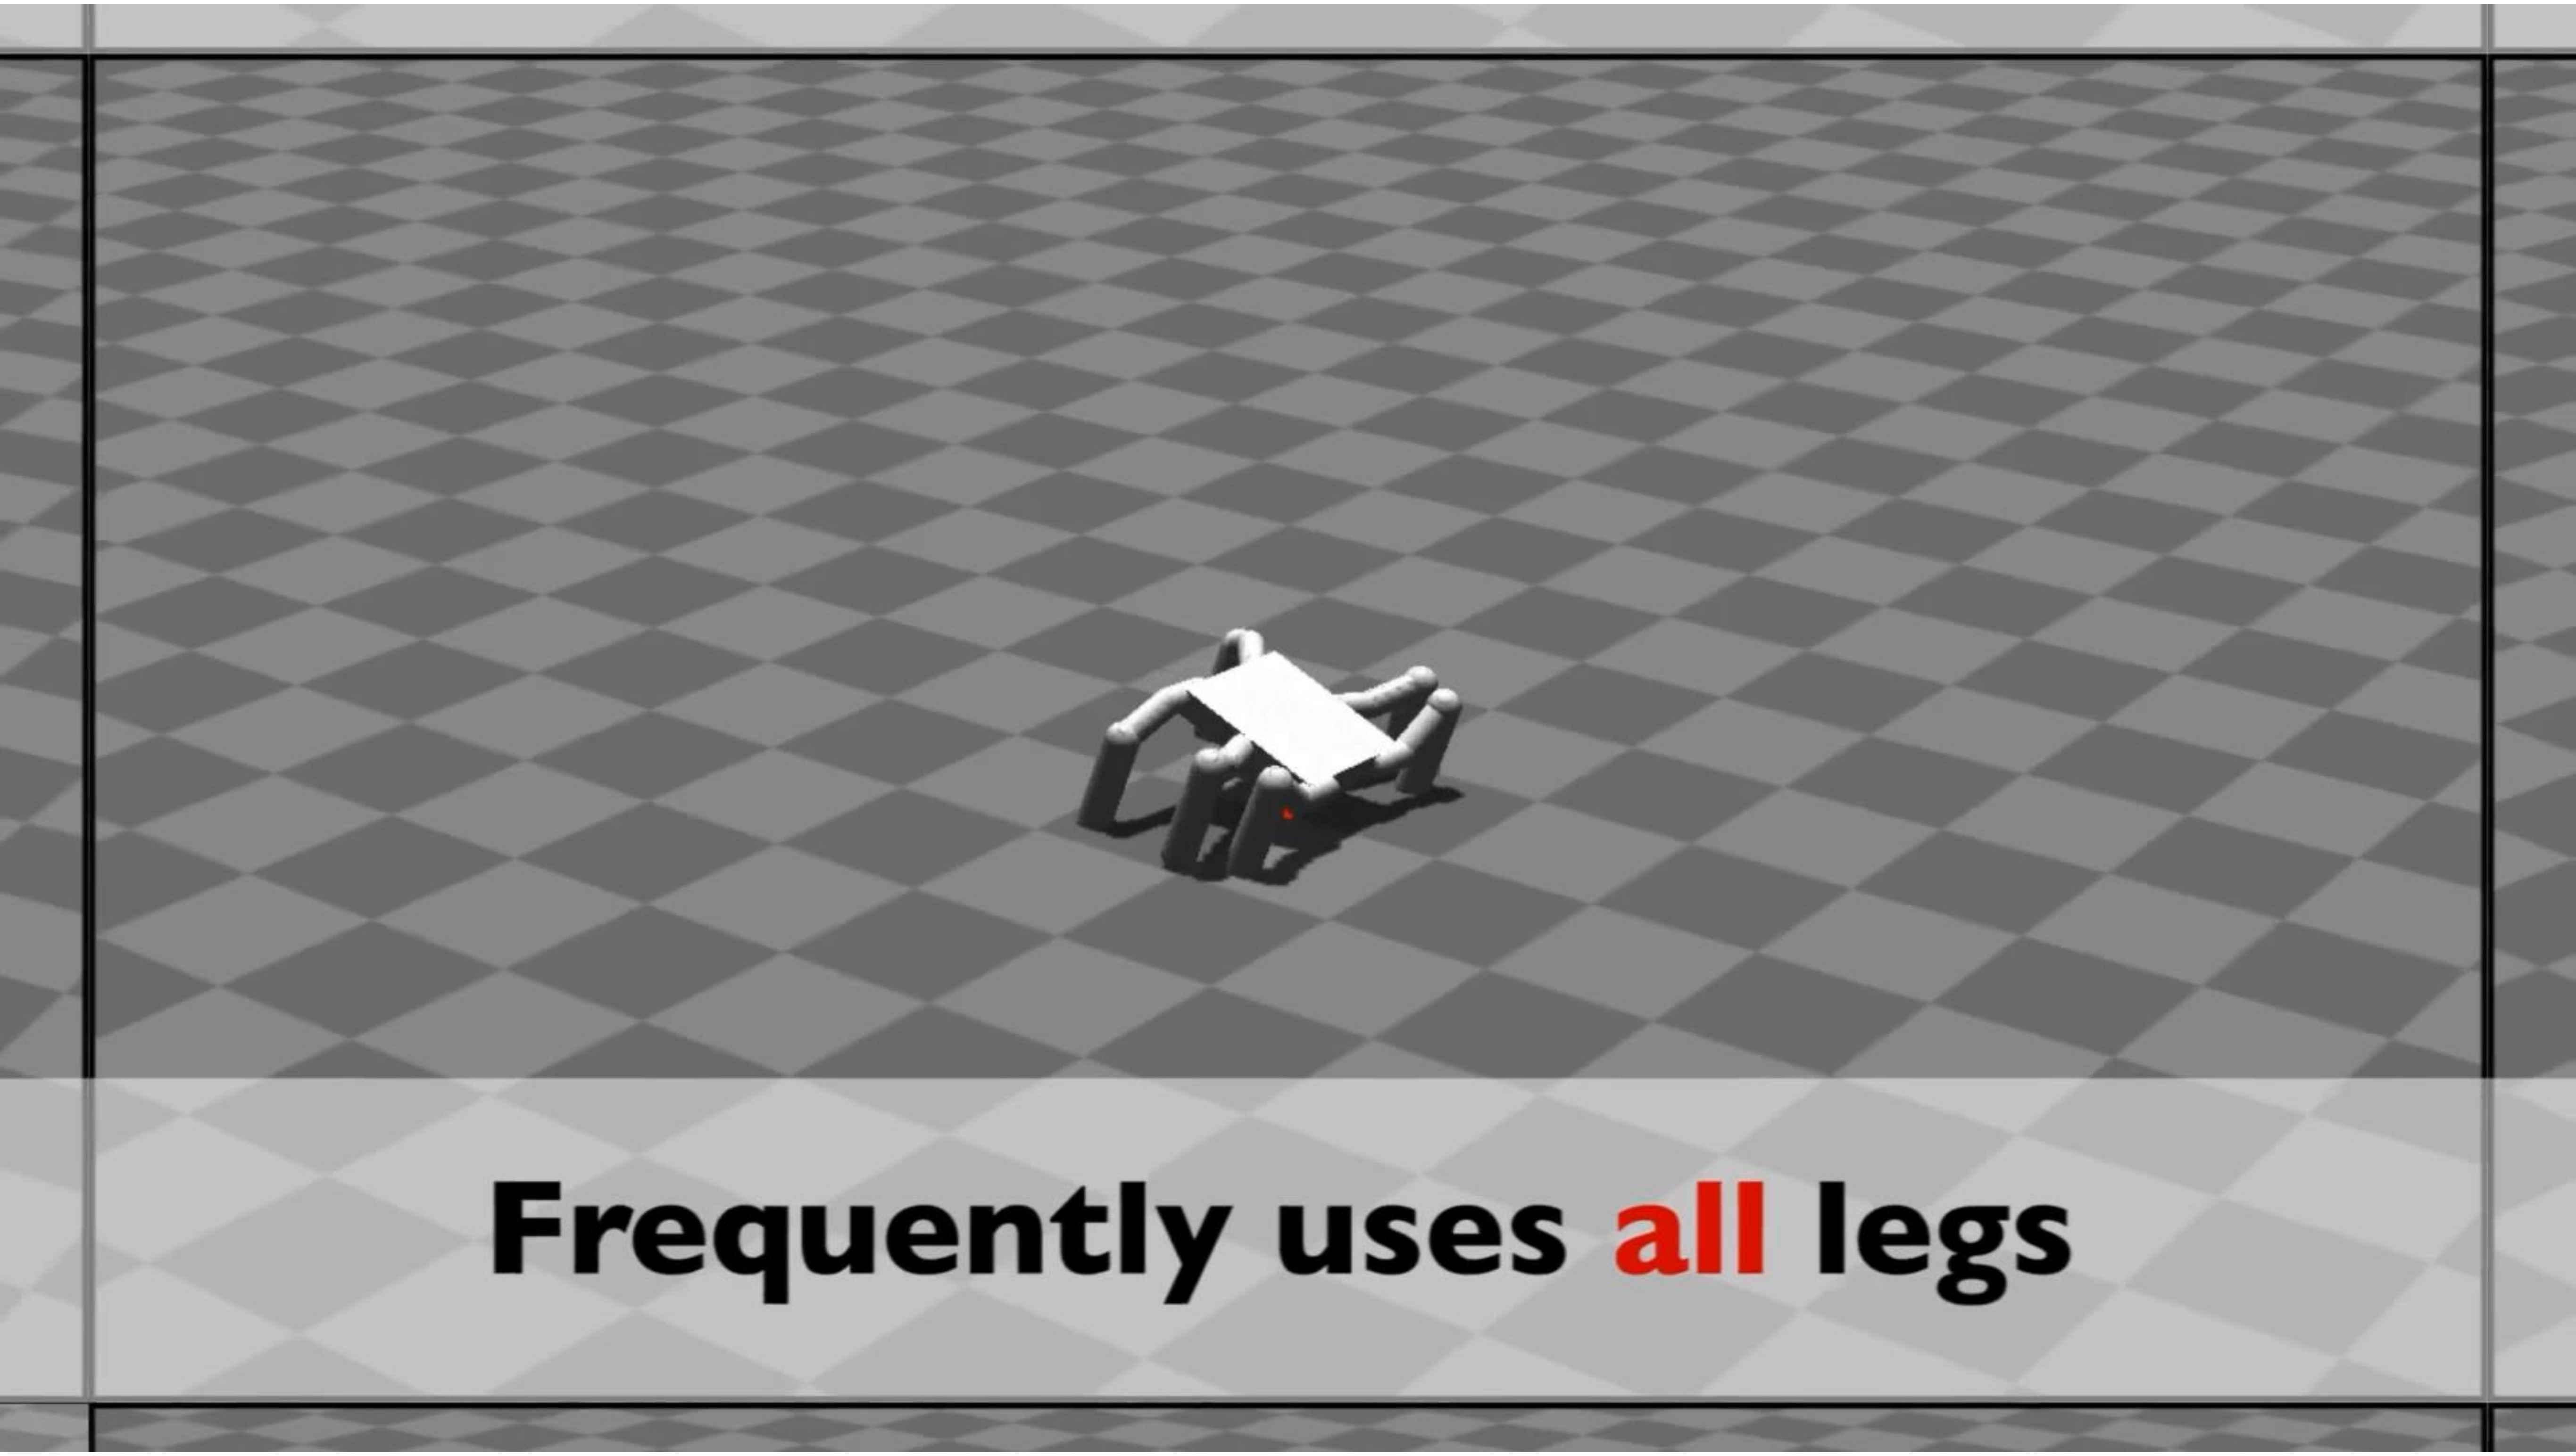

Challenge: damaged 6-legged robot

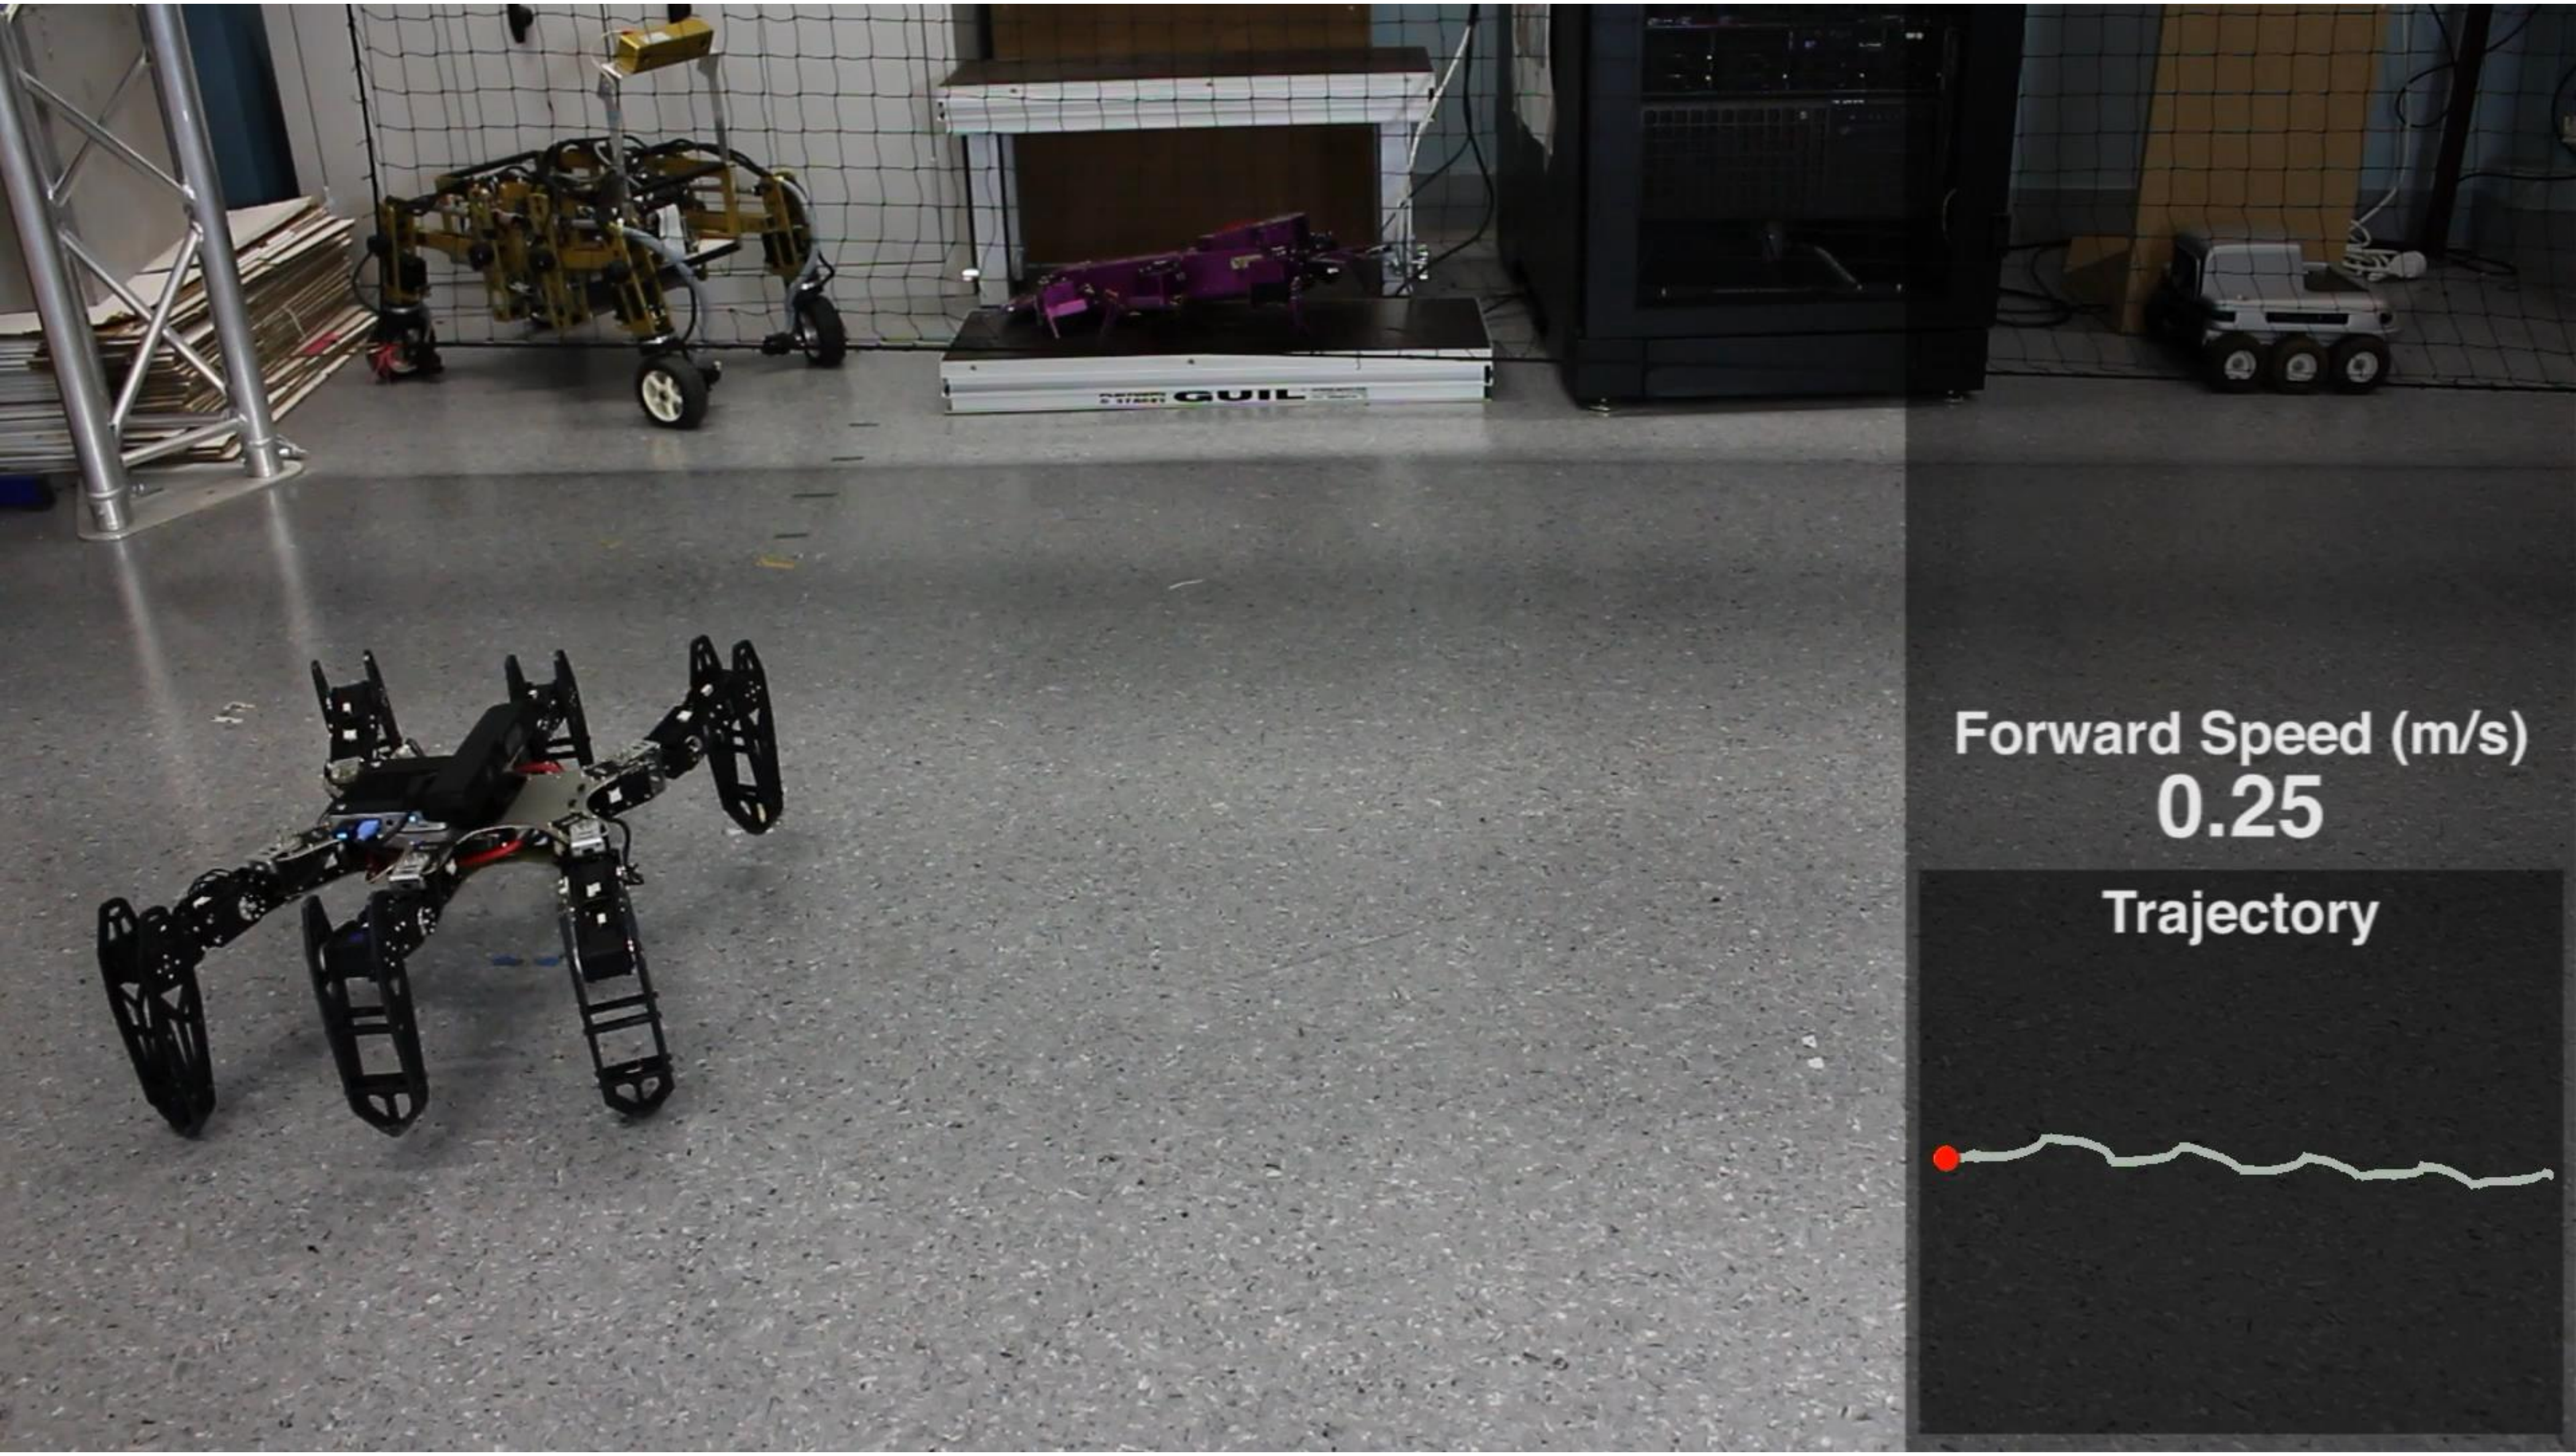

Strategy 2: prior + modelling the expected return

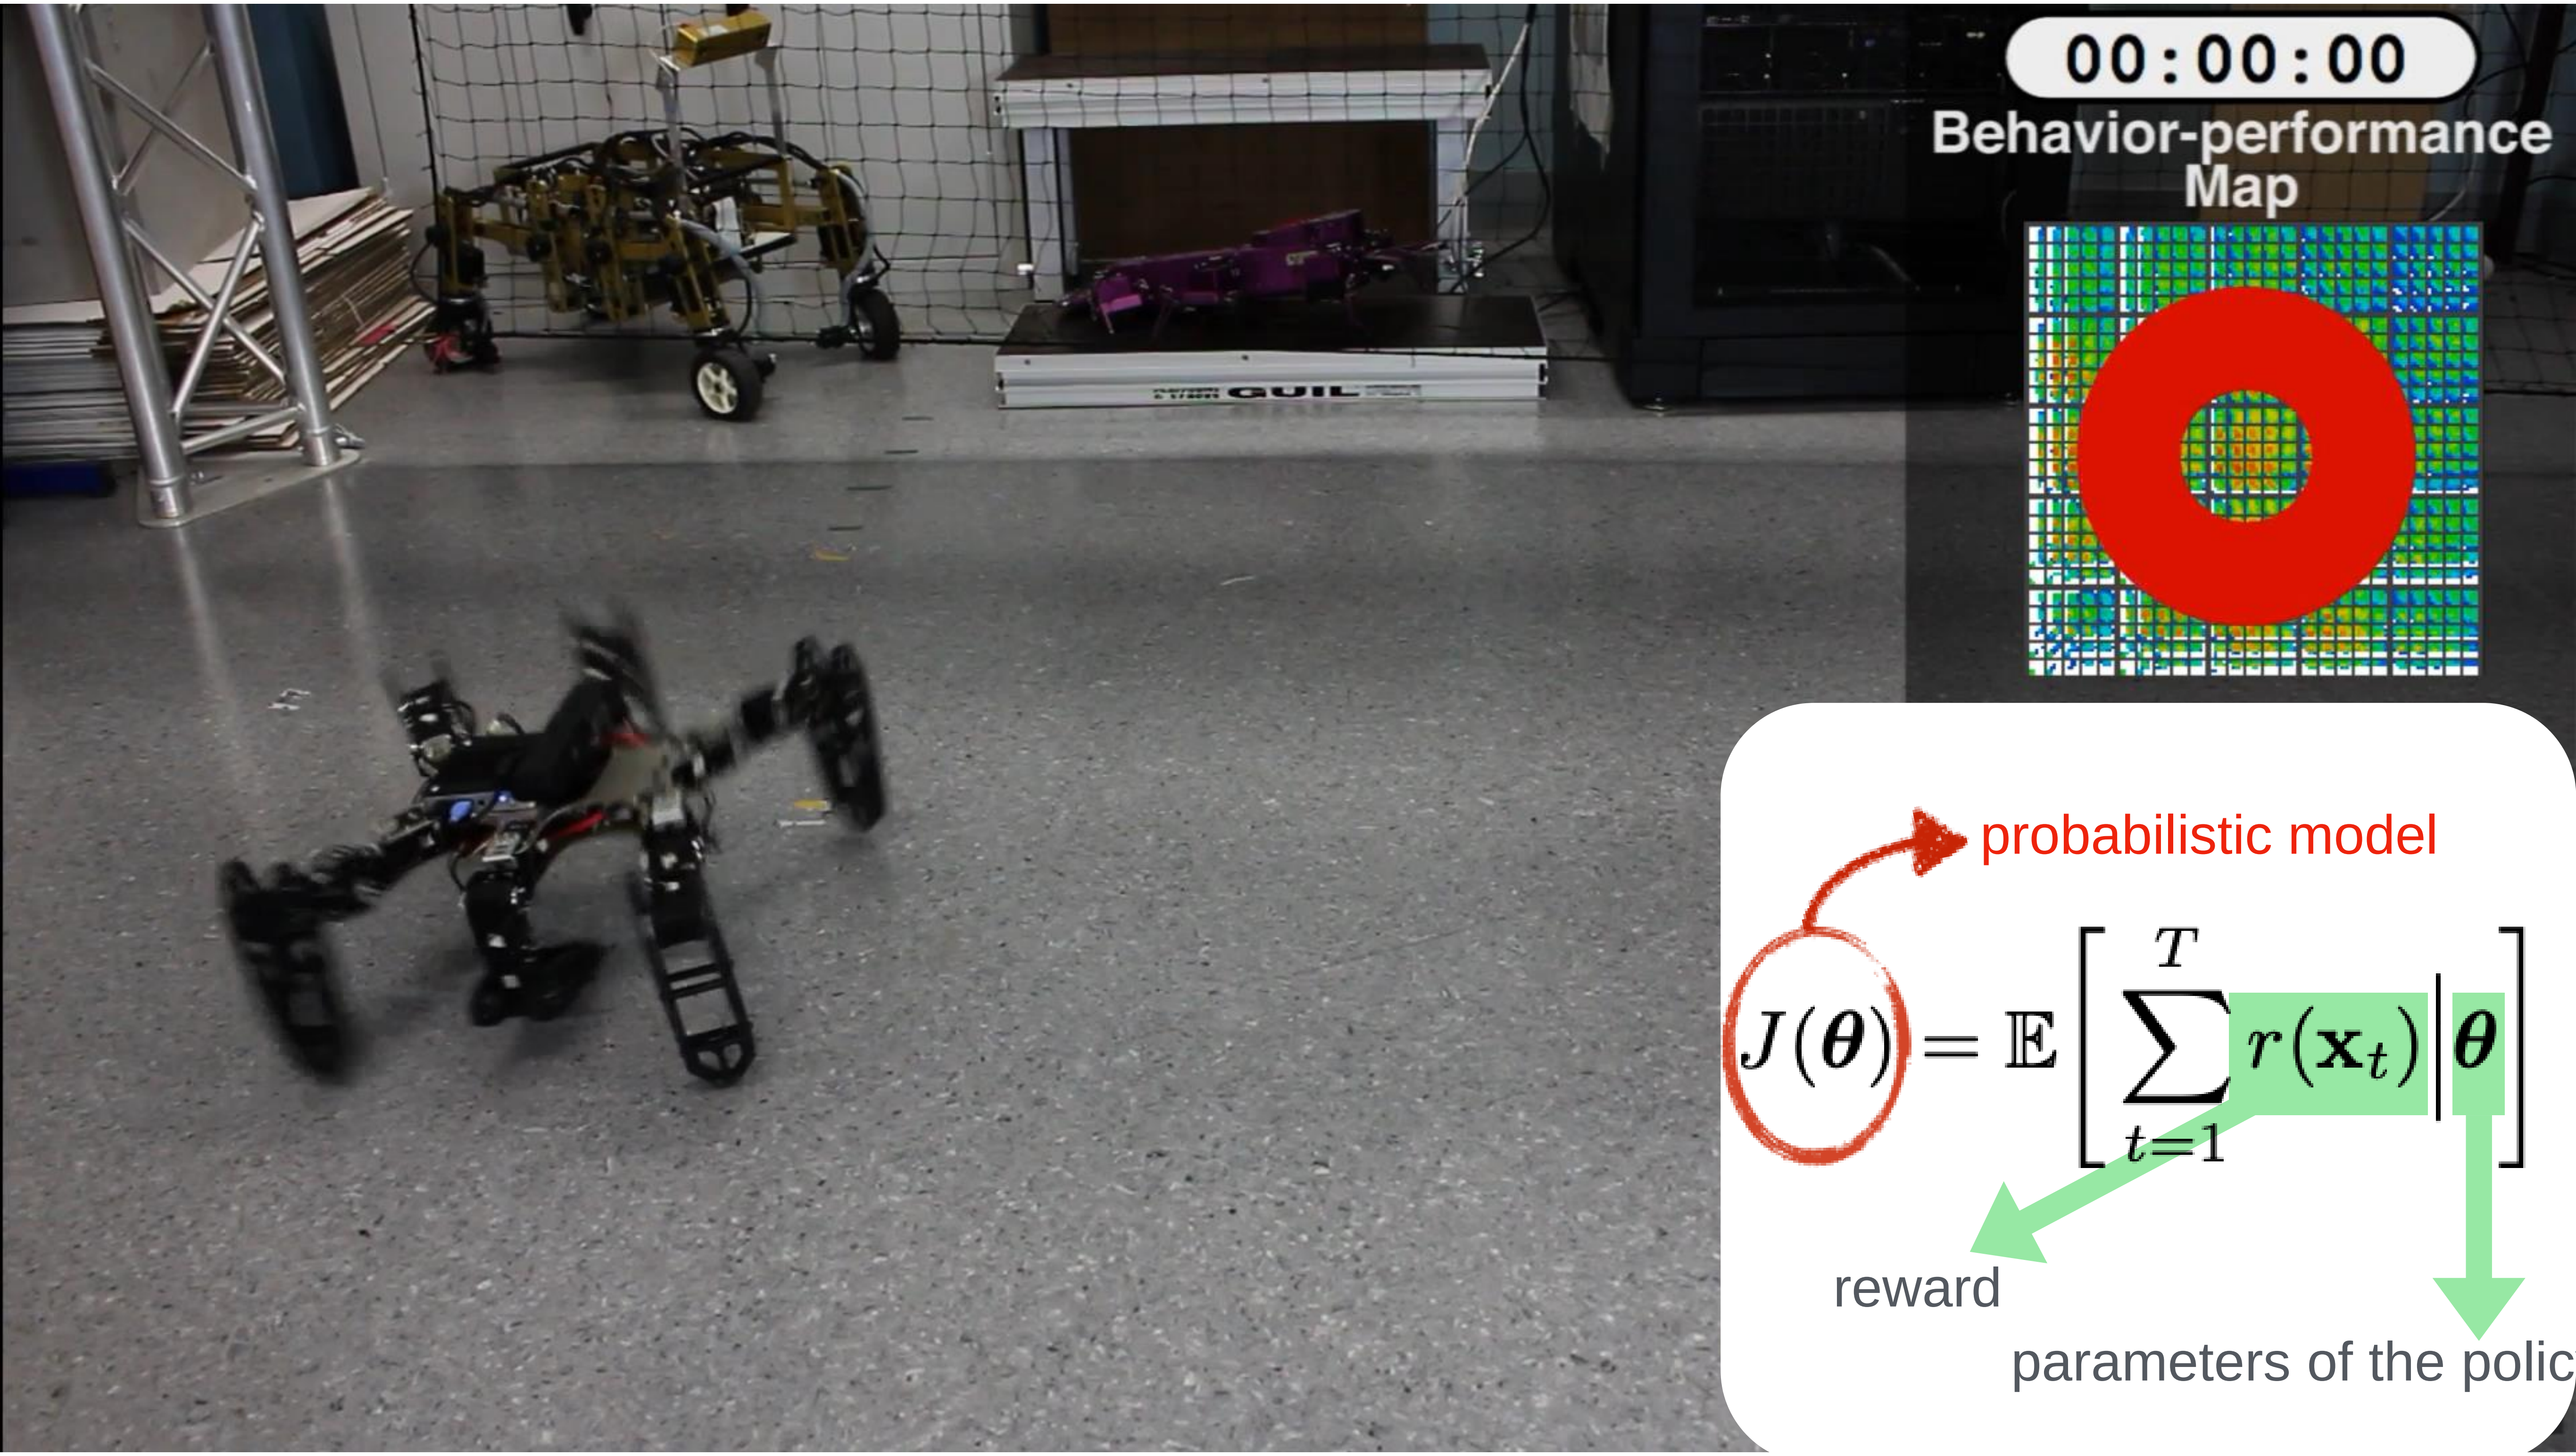

- Controller : periodical signals (36 parameters)
- <u>No information about the damage</u>

Cully, A. and Clune, J. and Tarapore, D. and Mouret, J.-B. (2015). Robots that can adapt like animals. Nature. Vol 521 Pages 503-507.

Strategy 2: modelling the expected return

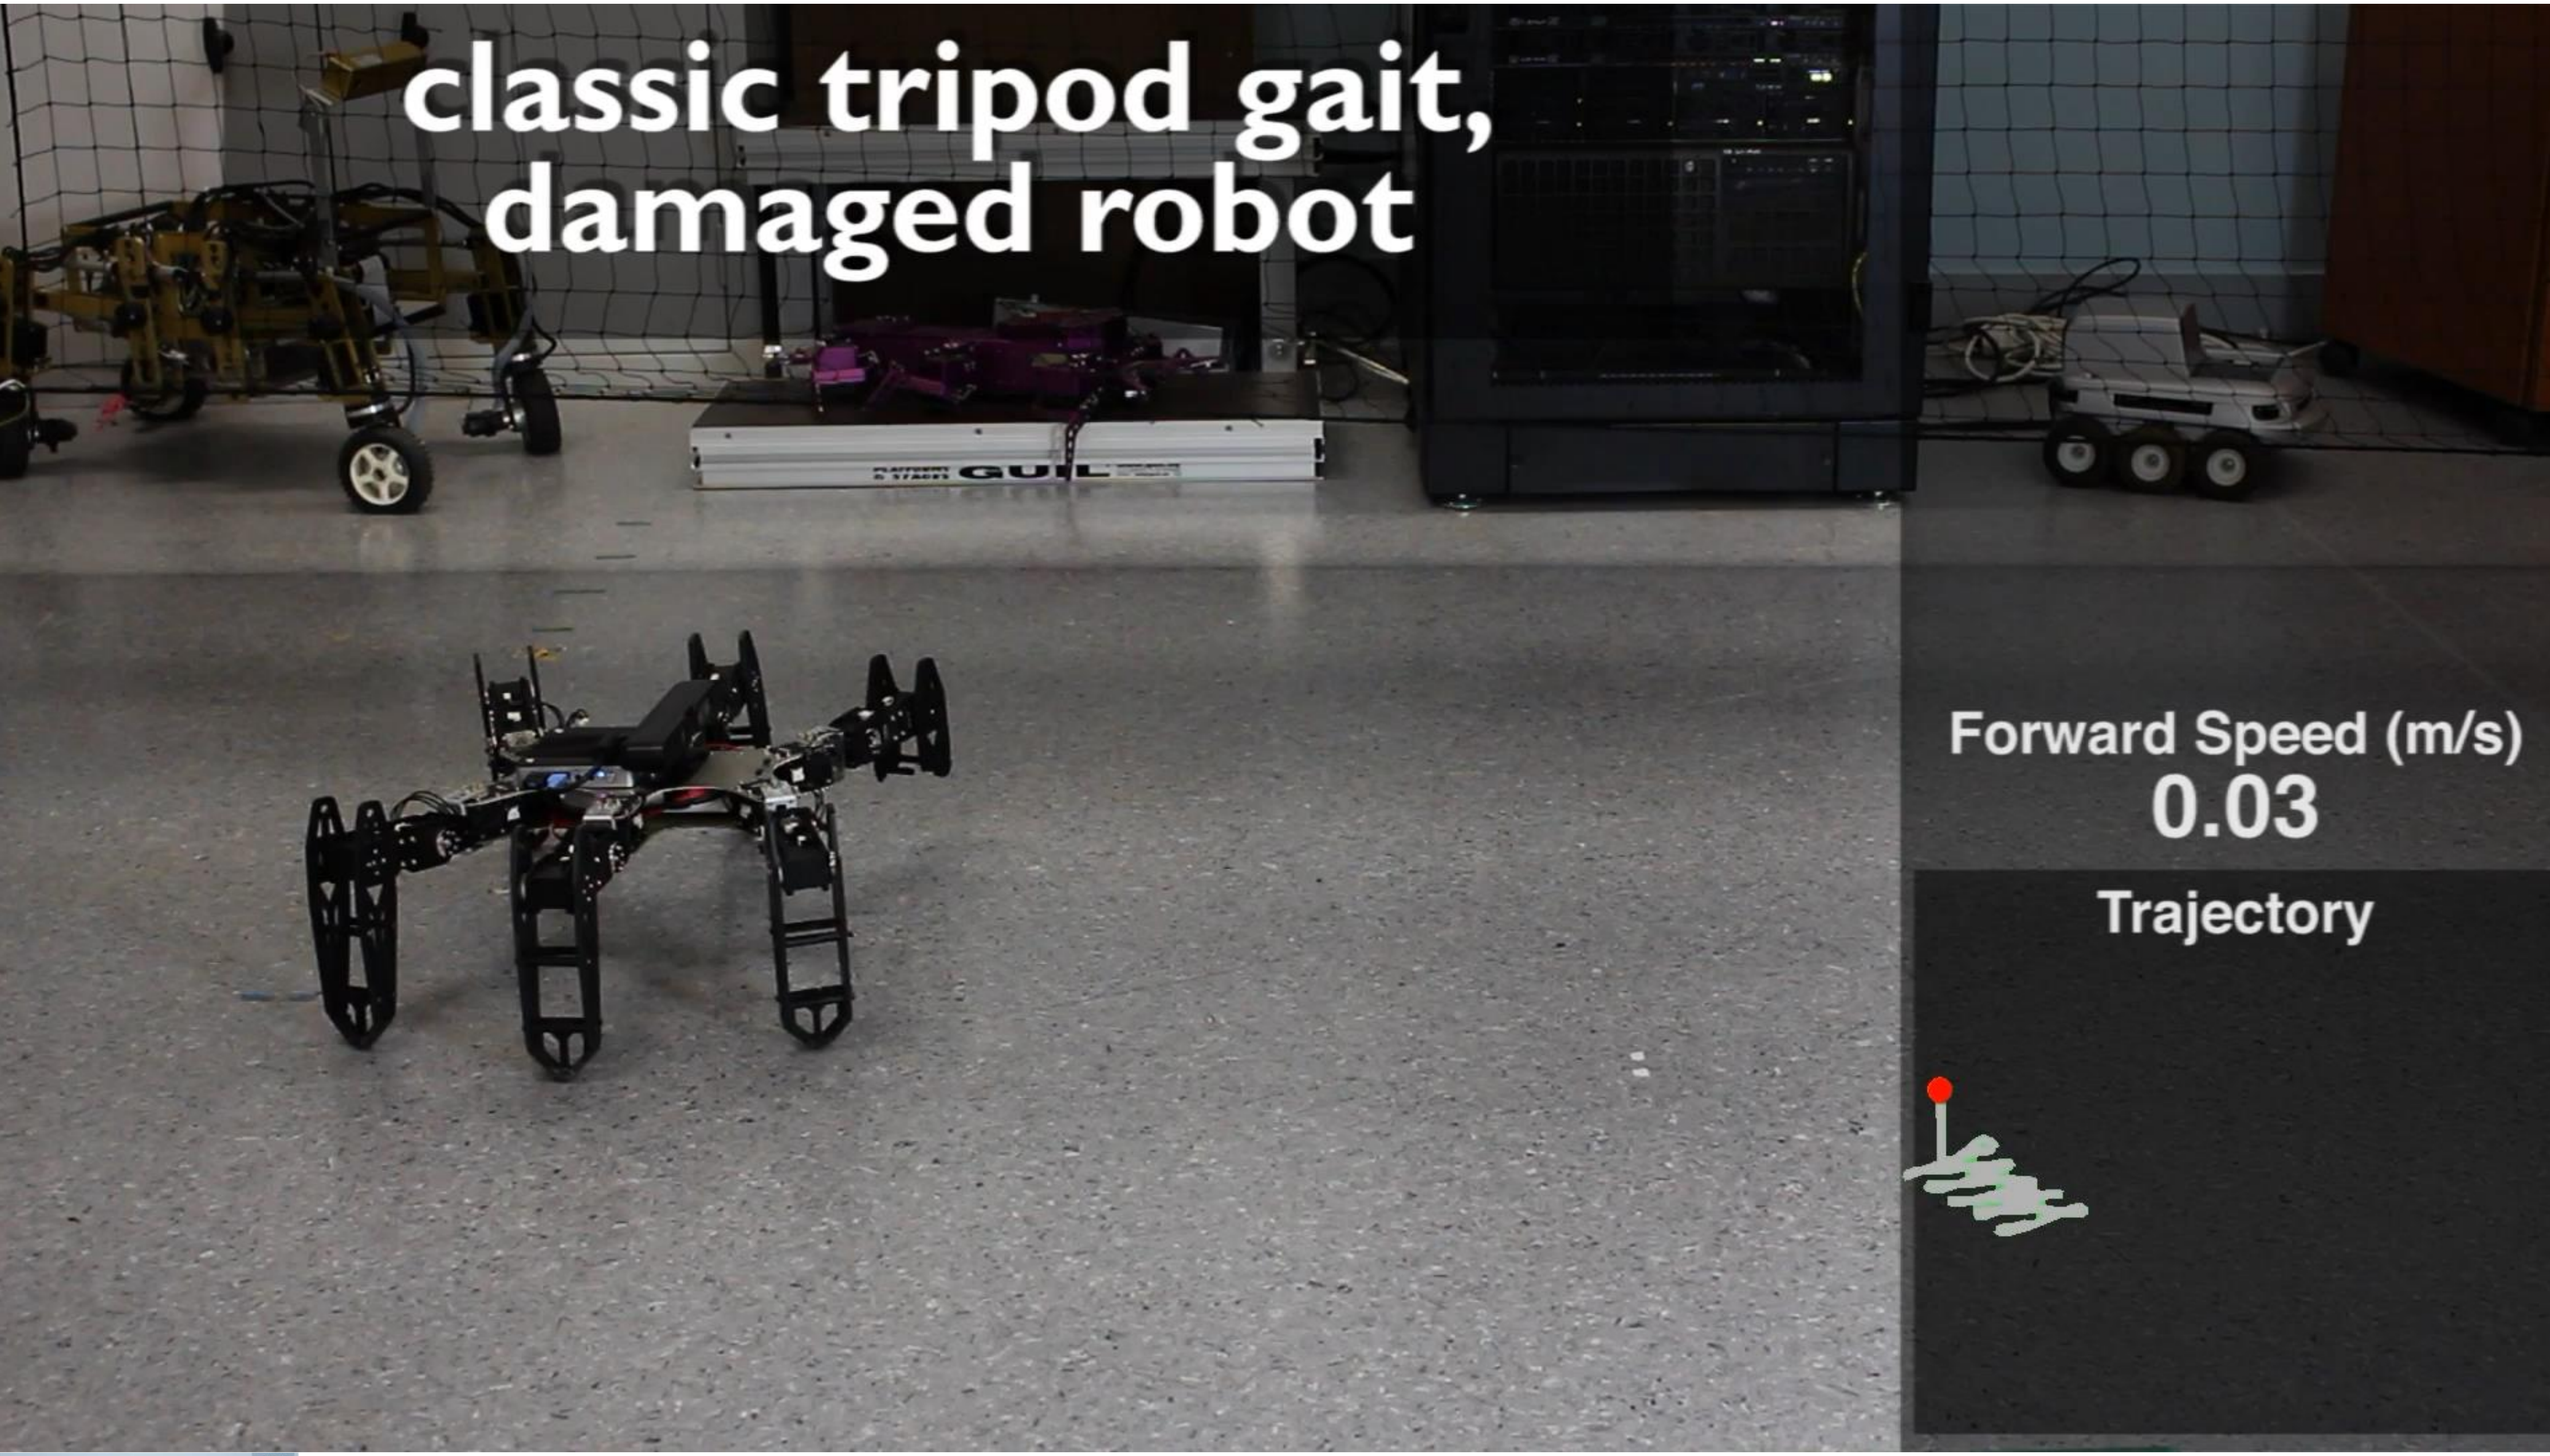

- Controller : periodical signals (36 parameters)
- <u>No information about the damage</u>

Cully, A. and Clune, J. and Tarapore, D. and Mouret, J.-B. (2015). Robots that can adapt like animals. **Nature**. Vol 521 Pages 503-507.

Planning (MCTS) + repertoire learning (priors)

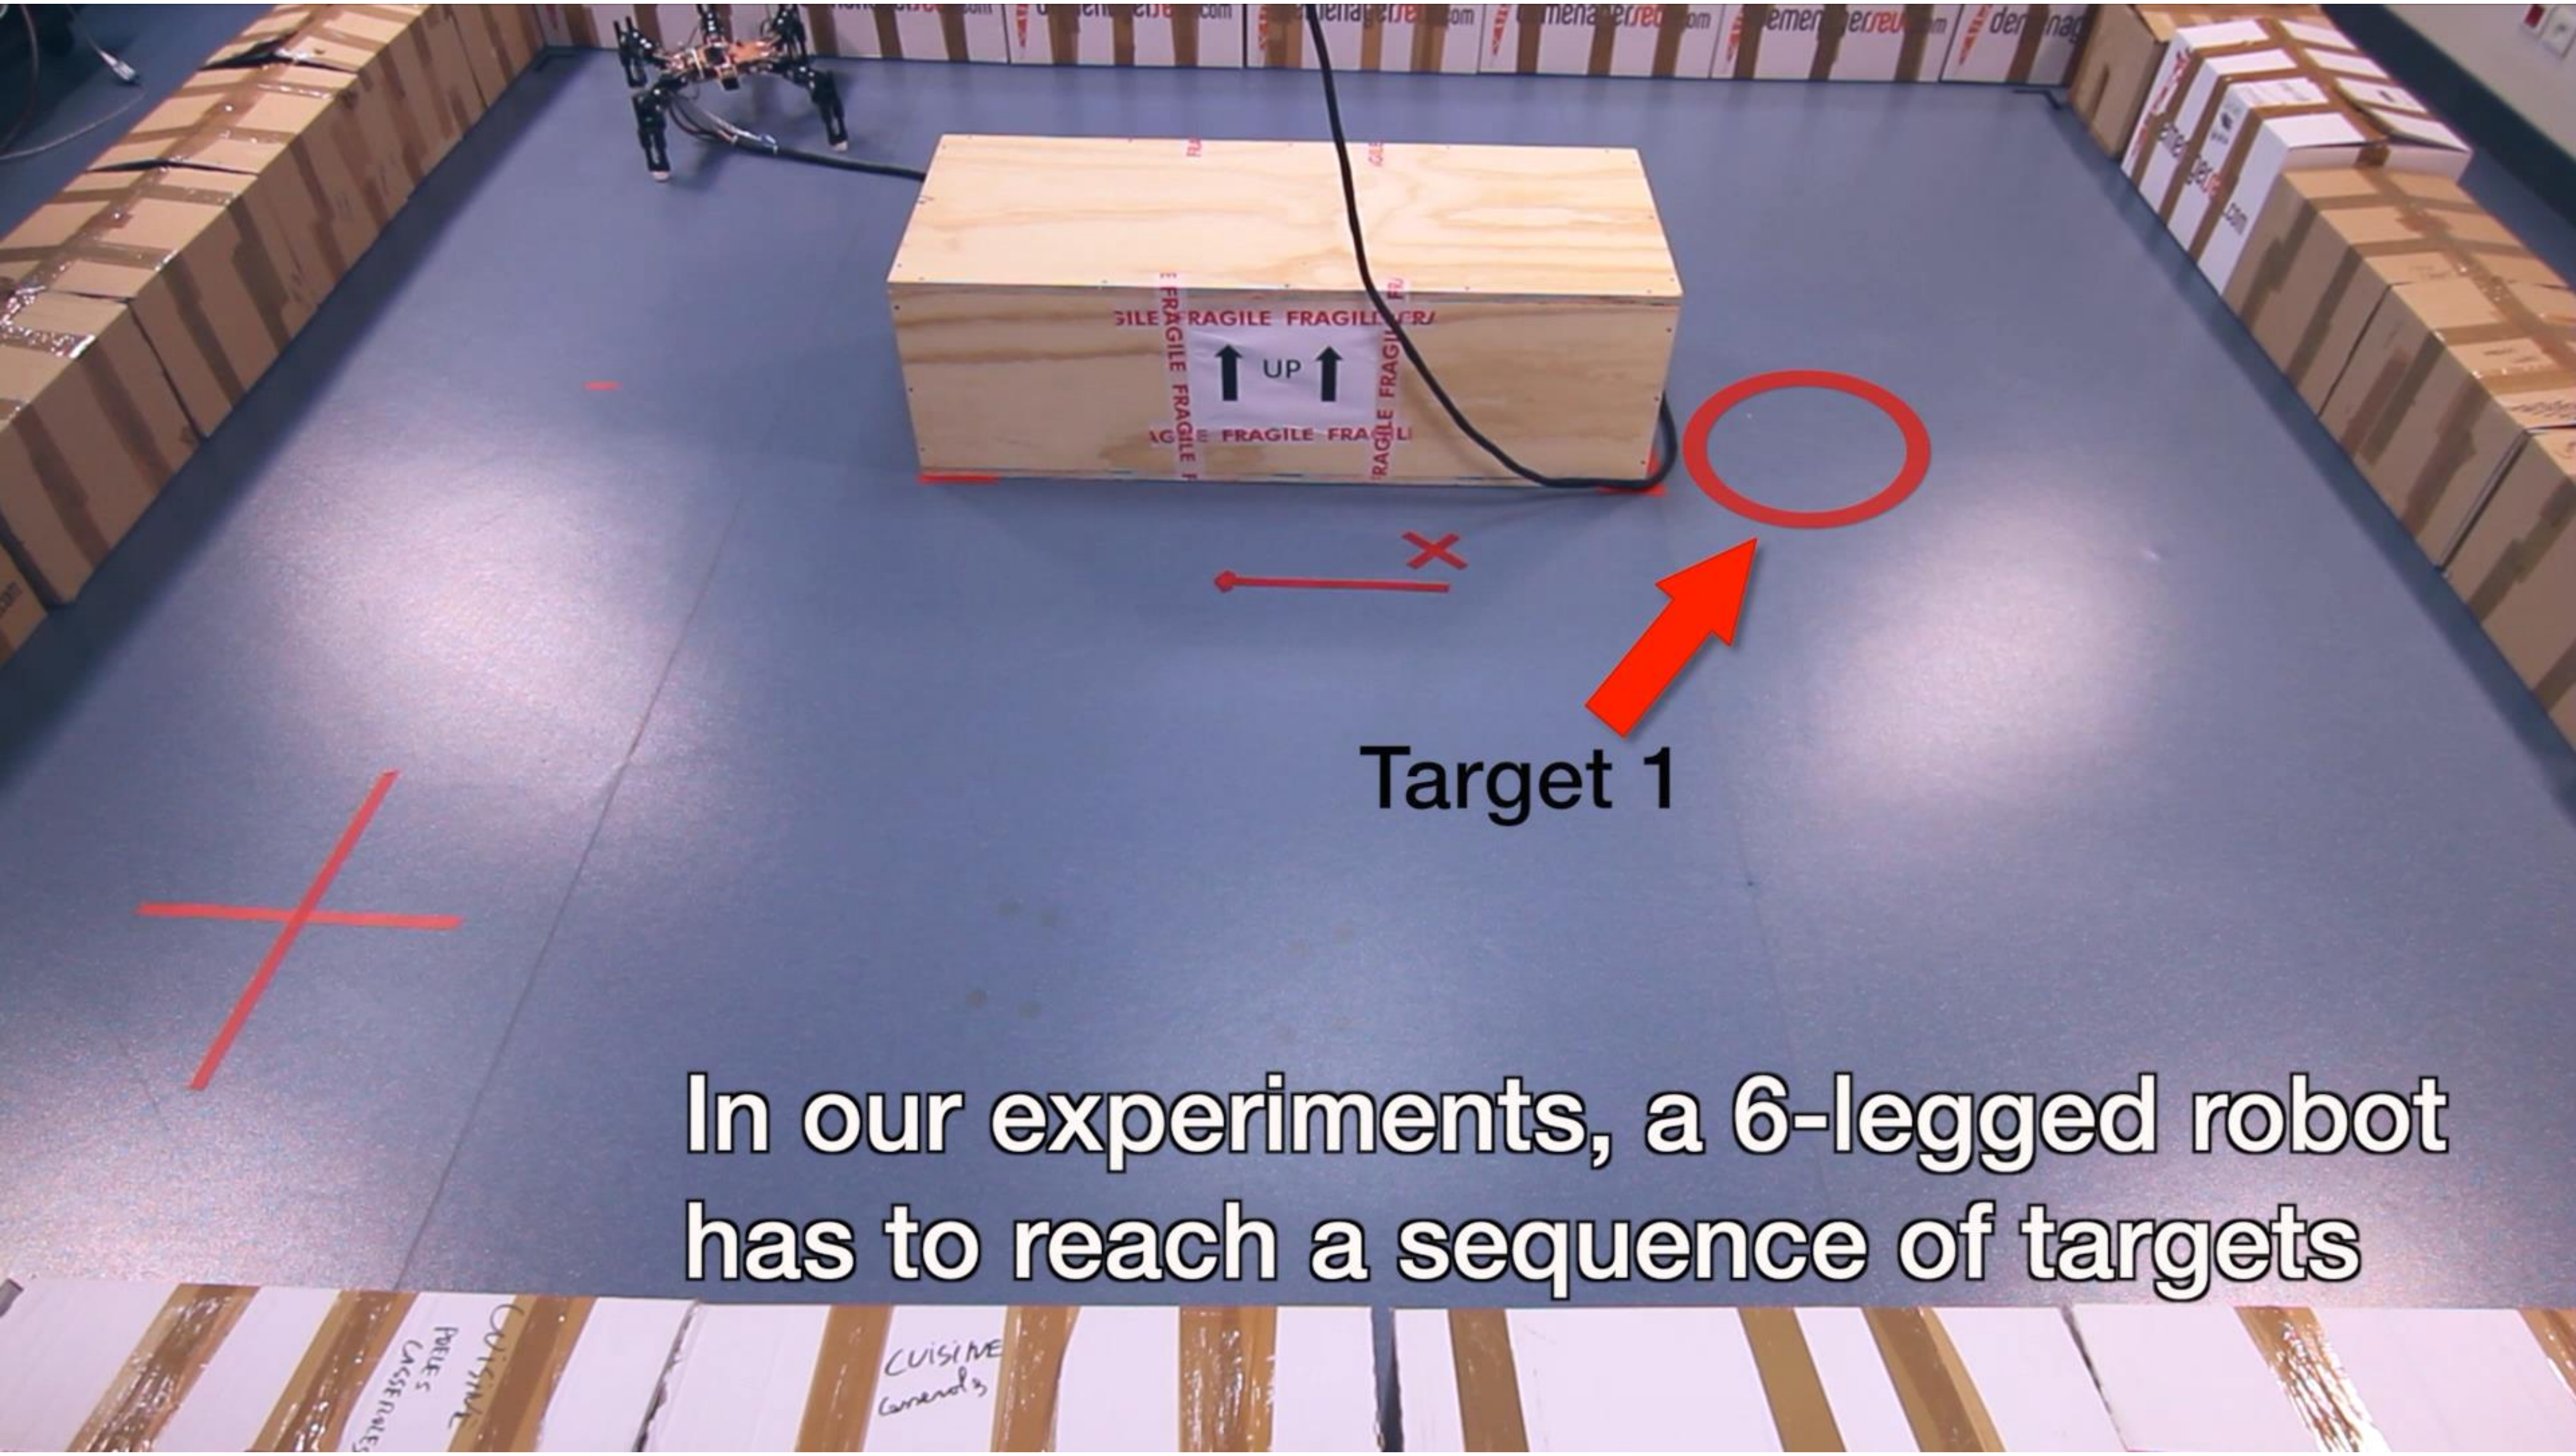

K. Chatzilygeroudis, V. Vassiliades, and J.-B. Mouret (2018). Reset-free Trial-and-Error Learning for Data-Efficient Robot Damage Recovery. Robotics and Autonomous Systems.

What do you I want to achieve?

Fast adaptation in the real world: a few minutes

➠ Much more autonomy

➠ Versatile robots for versatile missions

➠ Animals vs “superhuman”

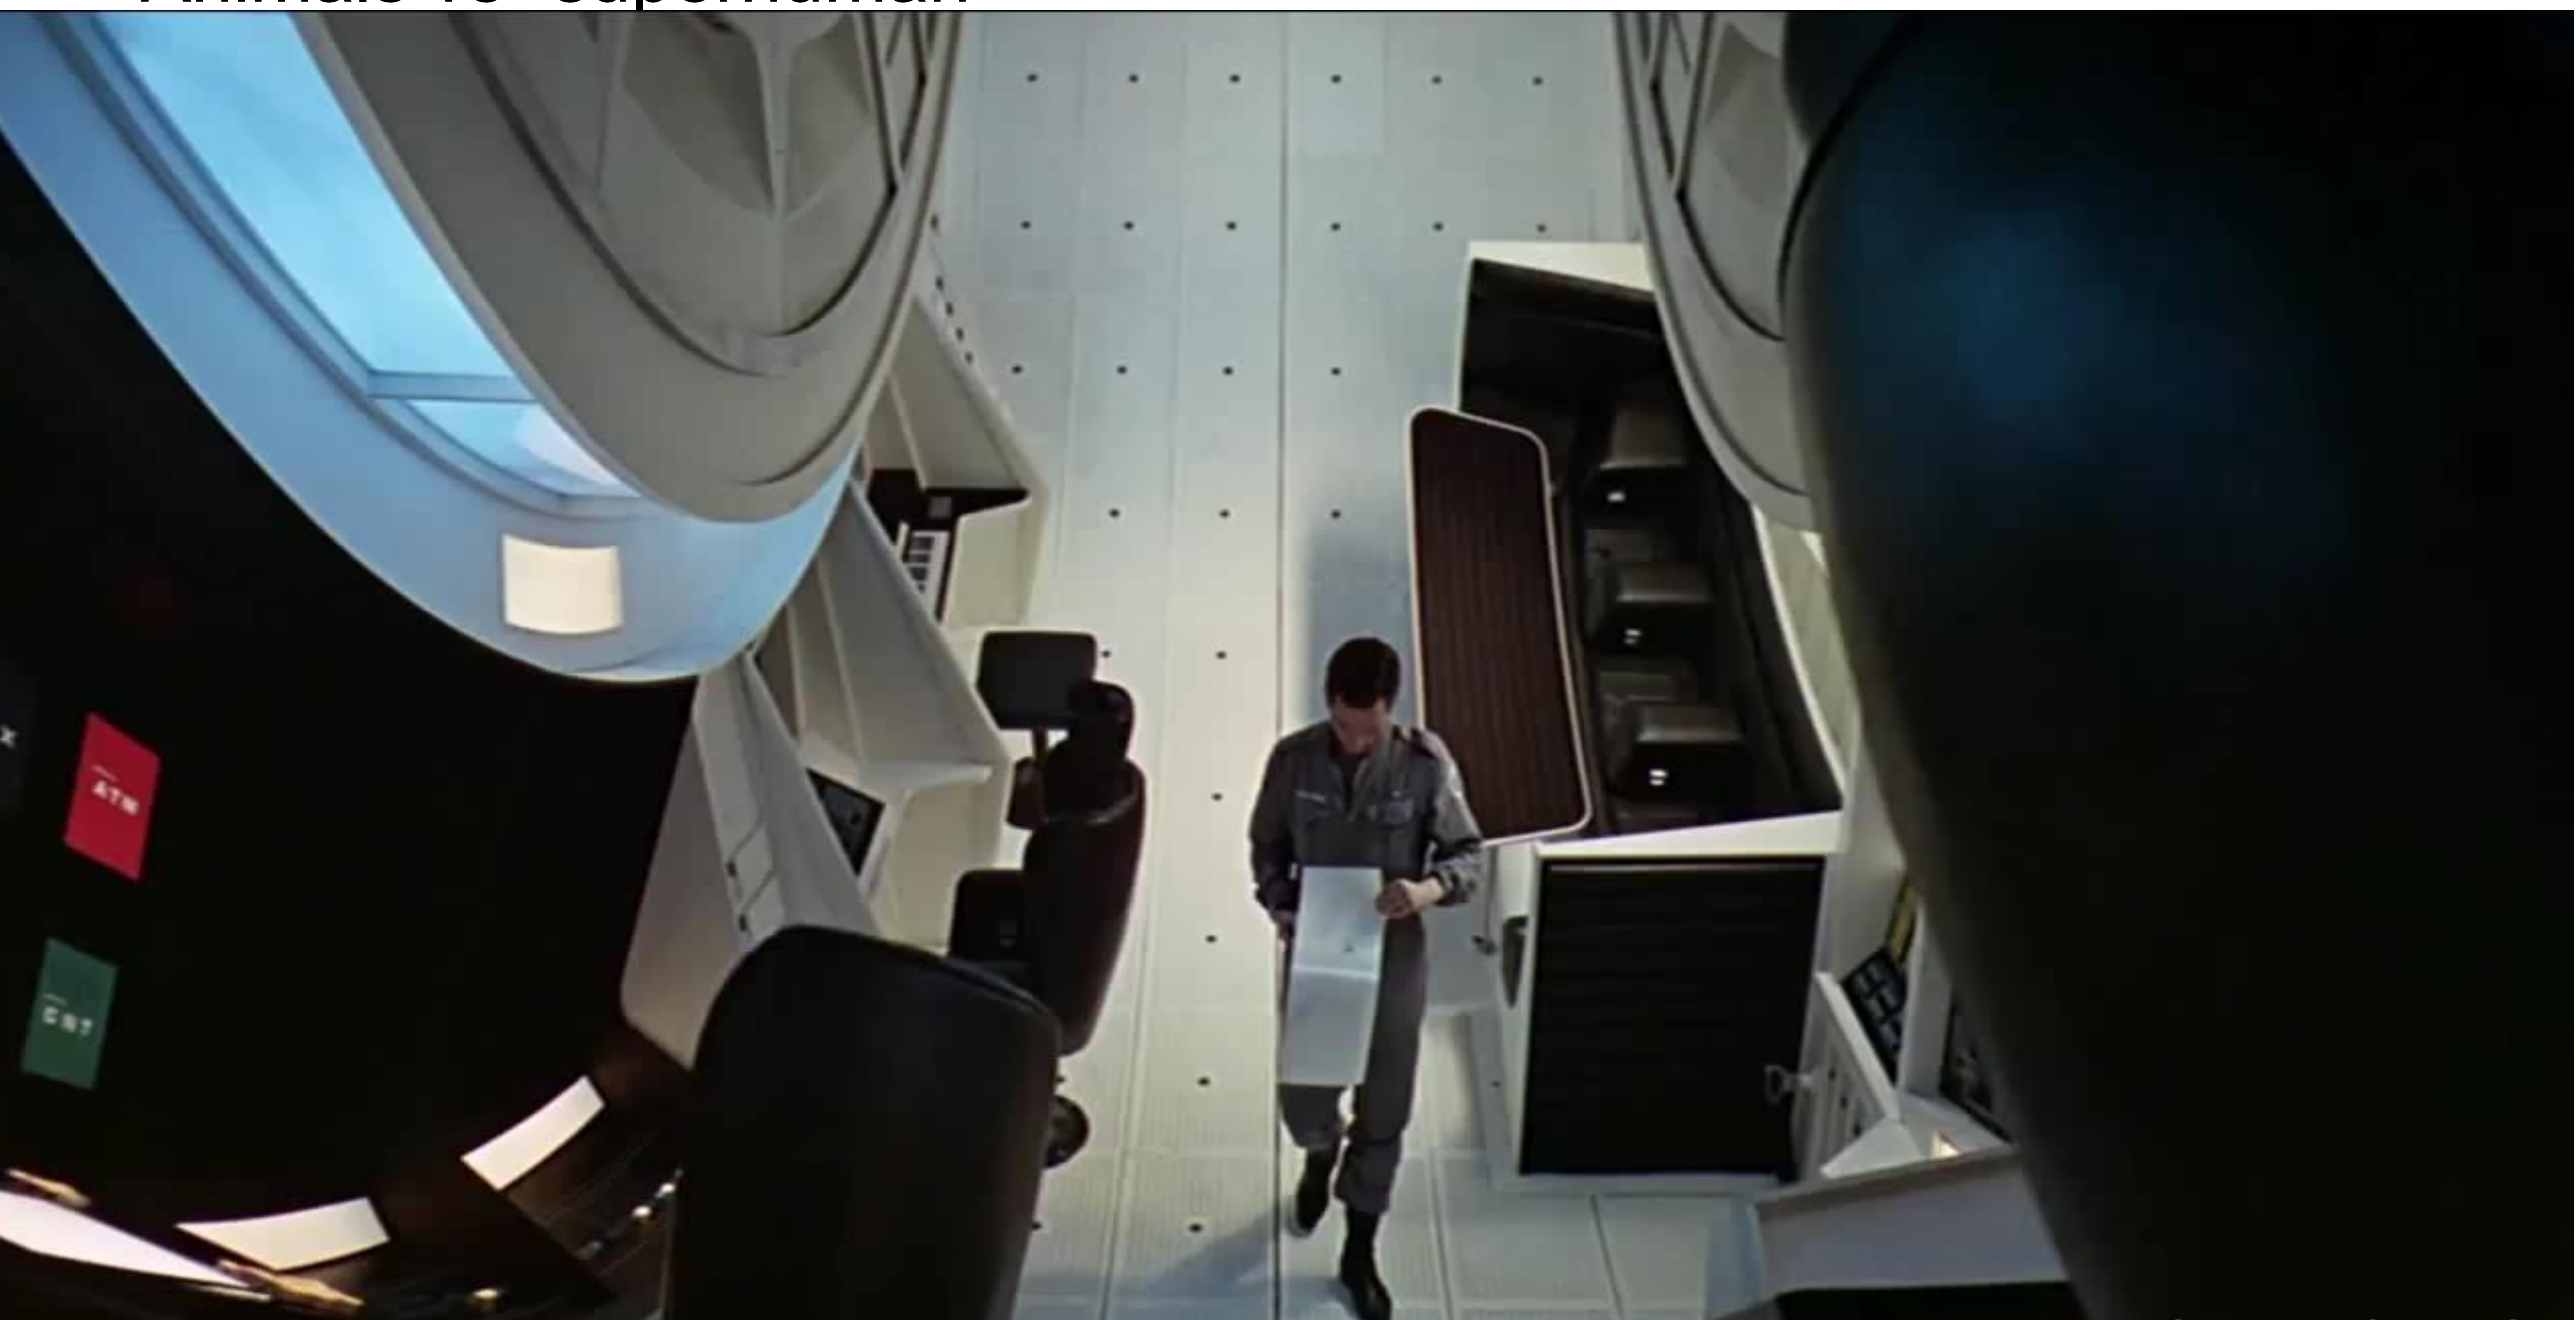

What do you I want to achieve?

Fast adaptation in the real world: a few minutes

➡ Much more autonomy

➡ Versatile robots for versatile missions

➡ More animals than “superhuman”

What do we need in term of science?

Amount of data

« Big-Data » ⟷ « Micro-data »

Simulated world

Real world

1-20 trials
A few minutes

1. Algorithms that can adapt with little data: most current learning does not really work on real machines & systems

2. Algorithms that can explore with rare rewards (e.g., door opening): curious machines / exploration

3. Prior knowledge is a key: from big data, models, engineering, …

4. Probabilistic predictive models of the world: know what you do not know
   - ➟ data-efficiency (search in the model)
   - ➟ “passive” acquisition (unsupervised)
   - ➟ “optimal” planning under uncertainty

5. Safety during learning: for the robot & for the environment

6. Continuous / long-term learning / avoid catastrophic forgetting

What do we need from society?

Focus on the important questions:

1. Autonomous robots are far from being able to “take over the world”
2. but “dumb” robots will be deployed and are more <u>dangerous</u>
   - bad <u>context awareness</u> = bad decisions / accidents / biased decisions

   ➡ a dumb robot is more dangerous than a “smart” one
   - robots ~ pets: they should NOT be dangerous for the world

Difficulty = (Complexity of task) x (Complexity of Environment) x (Cost of mistake)

Big Hero 6 (2014)

Conclusion

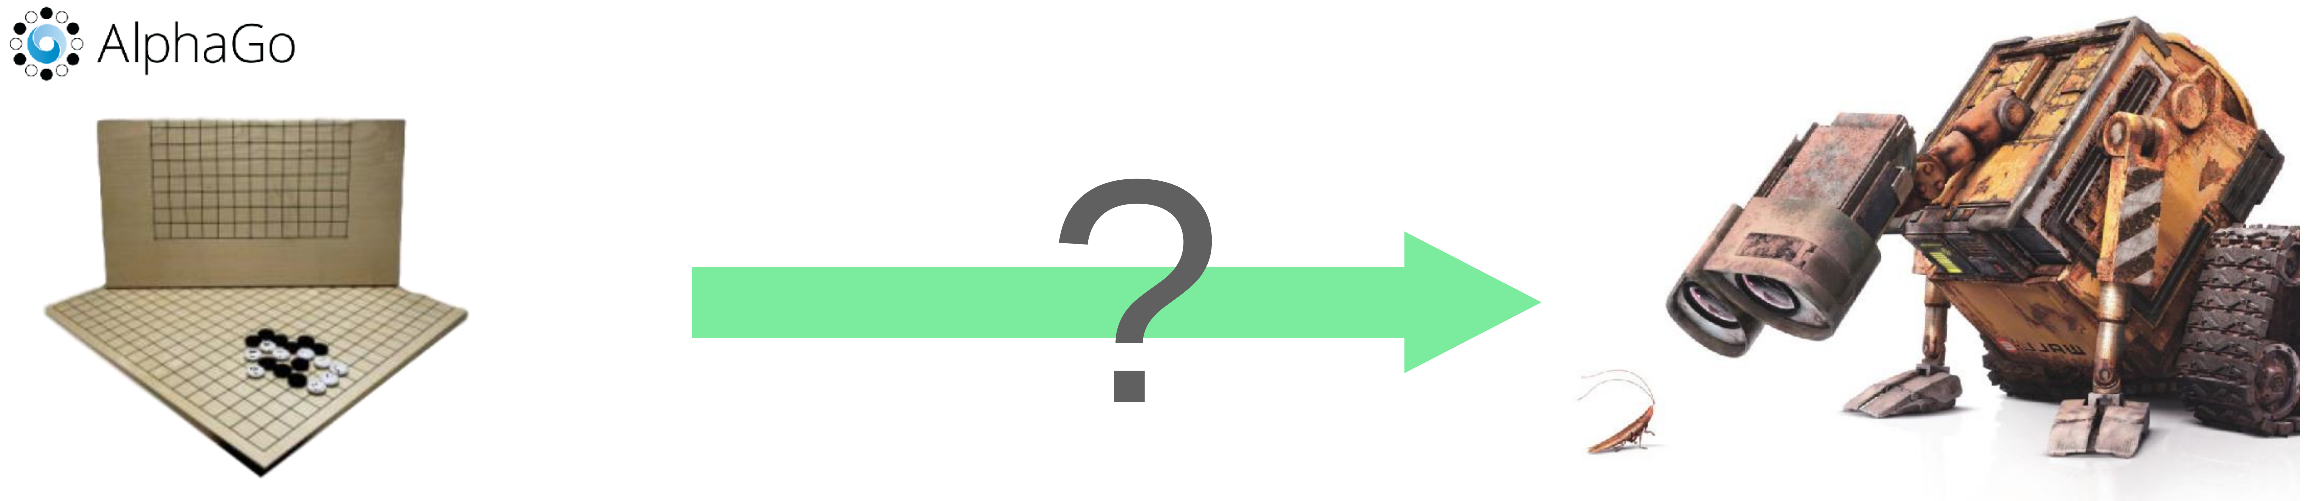

We achieved a lot in machine learning (e.g., deep learning)

… but this is not enough for robots (yet)

➡ we need algorithms that allow robots to adapt in a few minutes / a few trials

We should be more afraid by “dumb robots” than from “smart robots”

resibots

- Videos & papers: http://members.loria.fr/jbmouret
- ERC Project: http://www.resilient-robots.net
- code (C++11): http://github.com/resibots
- contact: jean-baptiste.mouret@inria.fr

Inria informatics mathematics

erc European Research Council
Established by the European Commission

Team


Sylvain Koos


Paul Tonelli


Antoine Cully

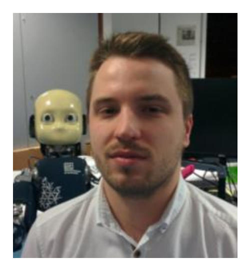
Brice Clément


Rituraj Kaushik

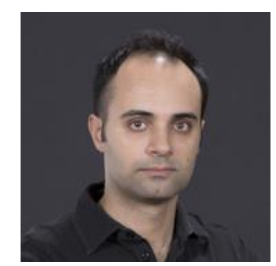
Vassilis Vassiliades

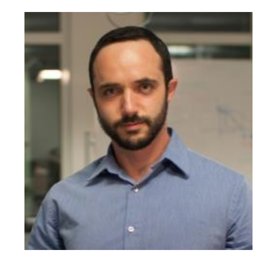
Jonathan Spitz


Danesh Tarapore


J.-M. Jehanno


Dorian Goepp


K. Chatzilygeroudis


Adam Gaier

MAP-Elites: 6-legged locomotion

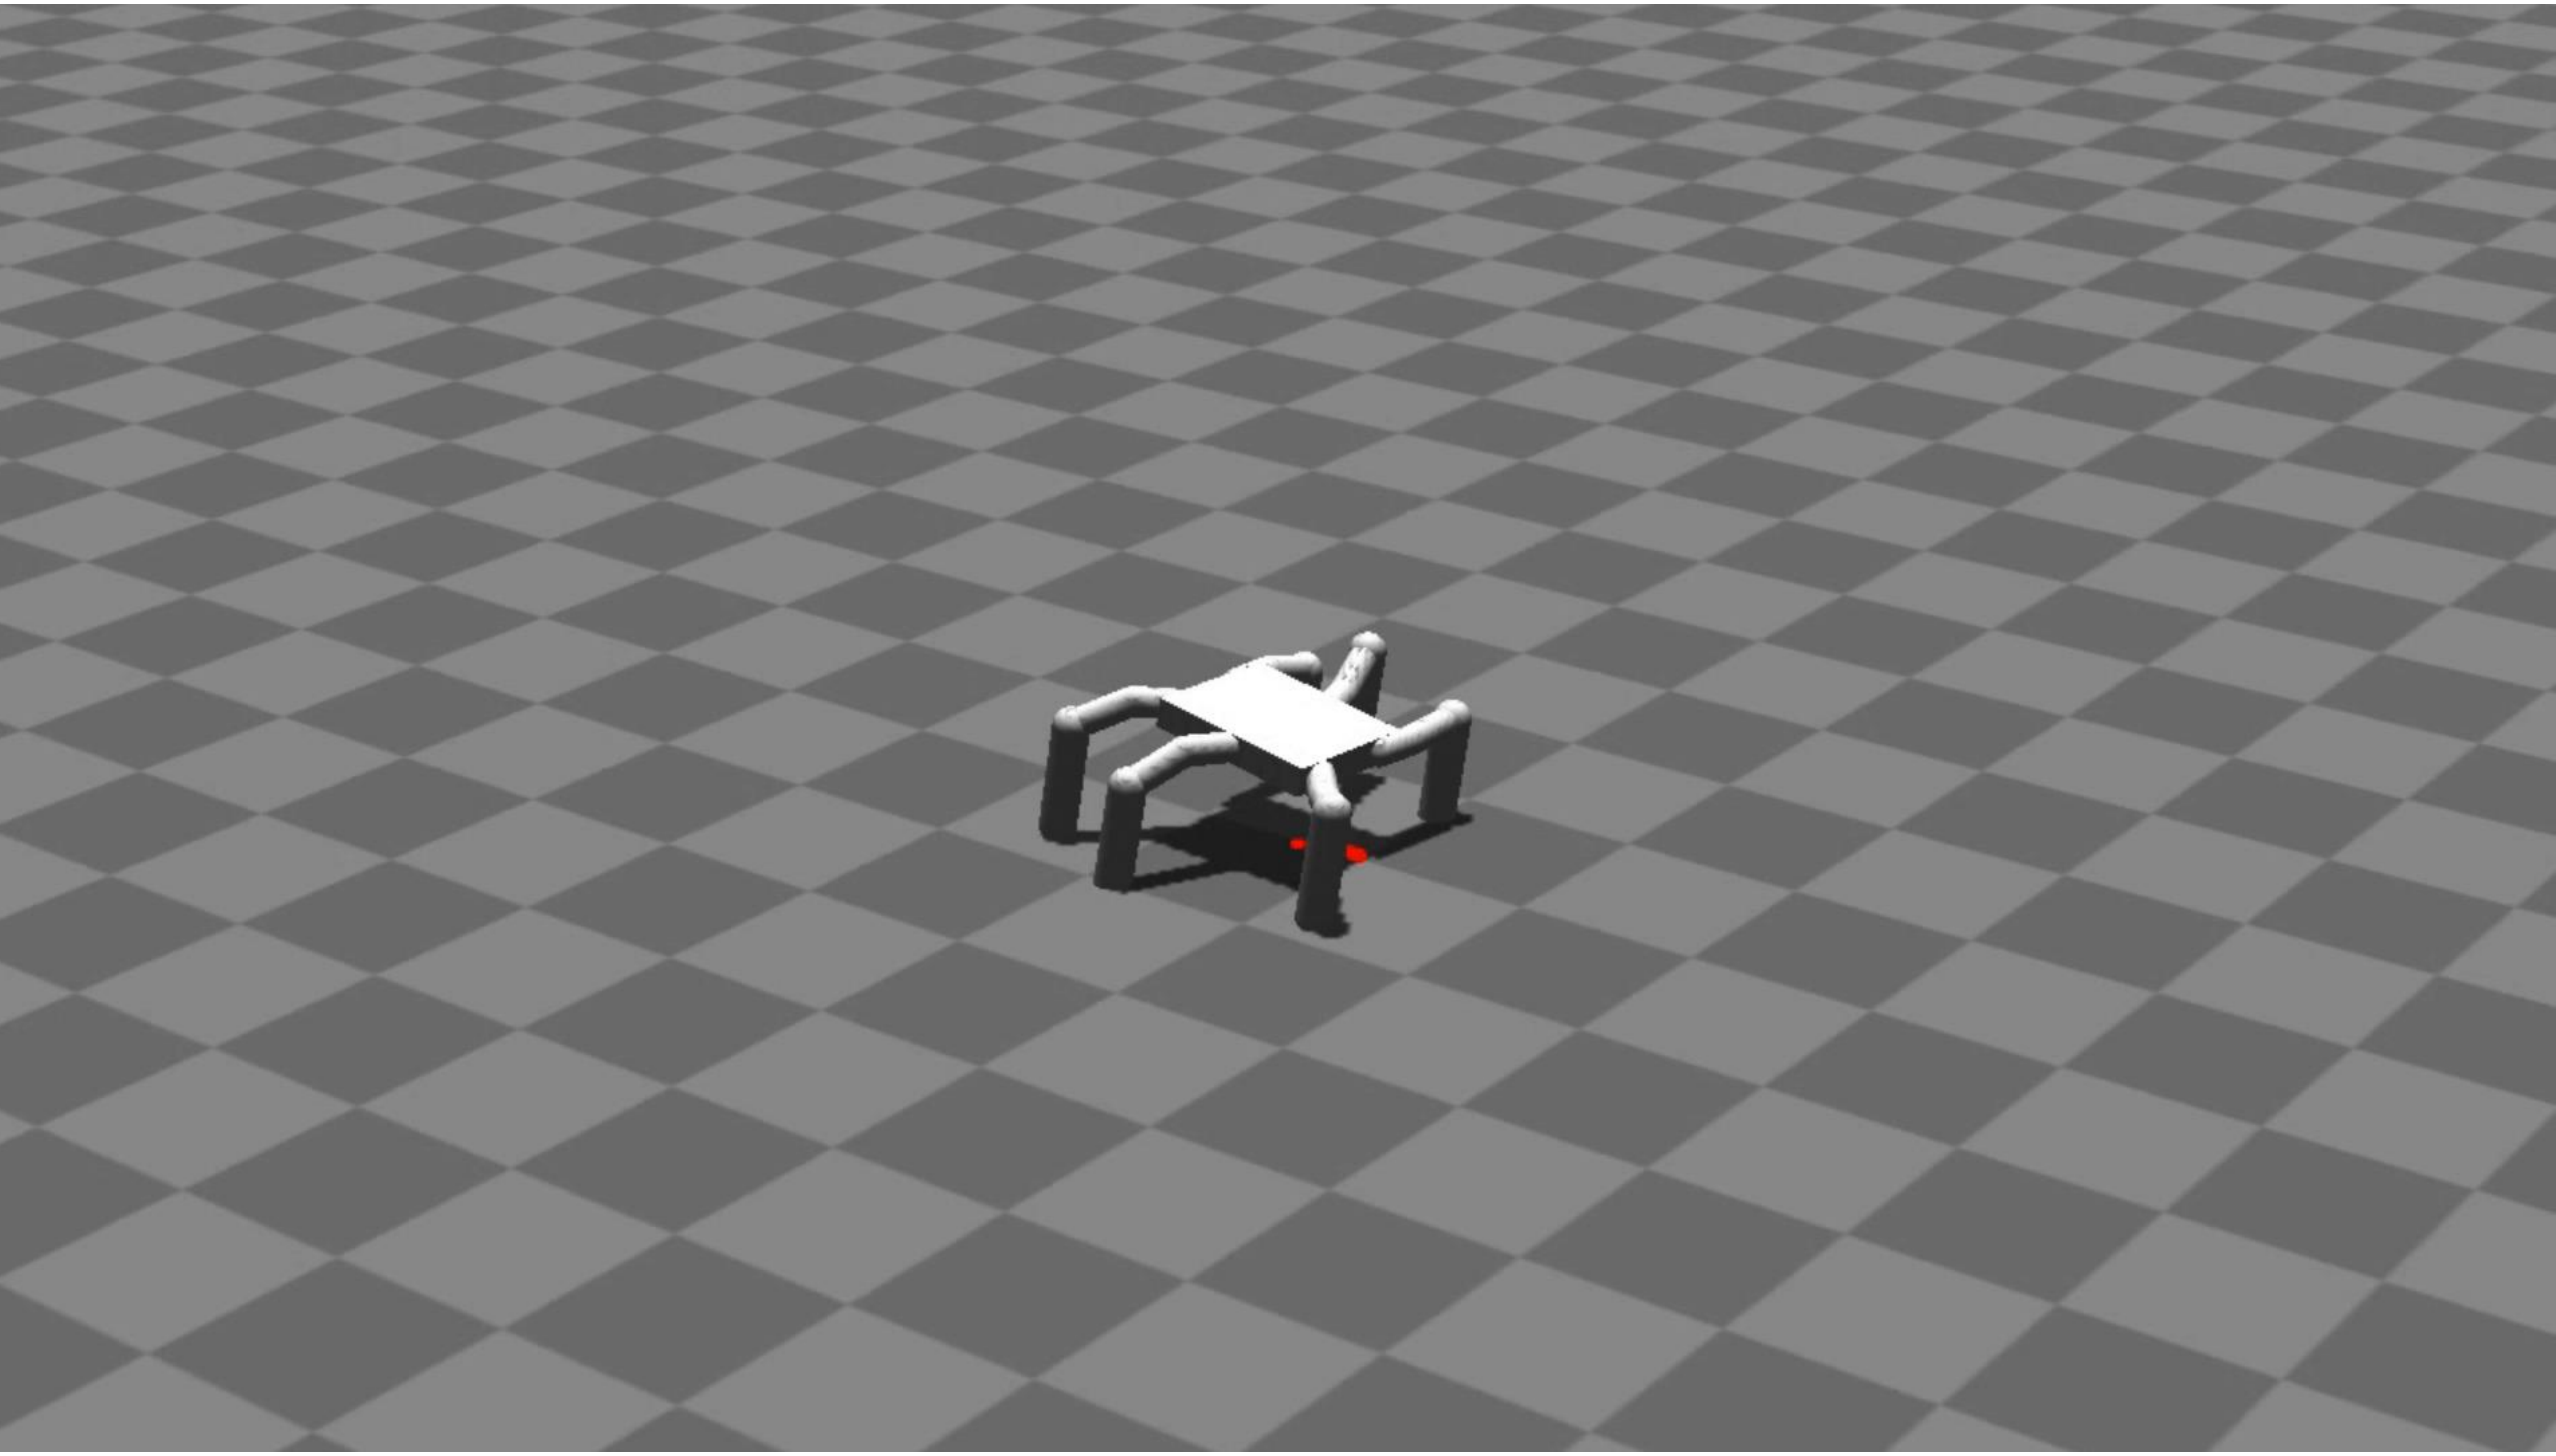

MAP-Elites: 6-legged locomotion