Institut de Robòtica
i Informàtica Industrial

CSIC

UNIVERSITAT POLITÈCNICA
DE CATALUNYA
BARCELONATECH

Knowledge, Reasoning and Planning

Adrià Colomé
Postdoctoral researcher

Palma, July 2022

Institut de Robòtica i Informàtica Industrial

- Personel:

>30 Doctors

18 Support personel

>40 PhD students

6 Administration

>25 Other students

>20 Robots

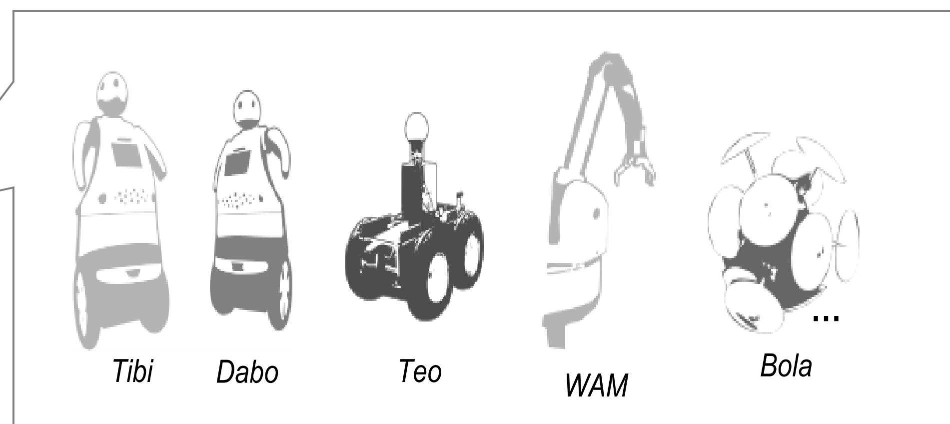

Integrating Knowledge, Reasoning and Planning

- What is knowledge?

  Facts, information, and skills acquired through experience or education;

  the theoretical or practical understanding of a subject.

- *Reasoning?*

  Inferring new information from previously known data connected logically.

  The argument's conclusion must be proper if and only if the premises are true.

- *Planning?*

  Realization of strategies or action sequences, typically for execution by intelligent agents, autonomous robots and unmanned vehicles.

Integrating Knowledge, Reasoning and Planning

Integrating Knowledge, Reasoning and Planning

DeepMind

StarCraft II

DEMONSTRATION

Blizzard Entertainment

Institut de Robòtica i Informàtica Industrial

CSIC

UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH

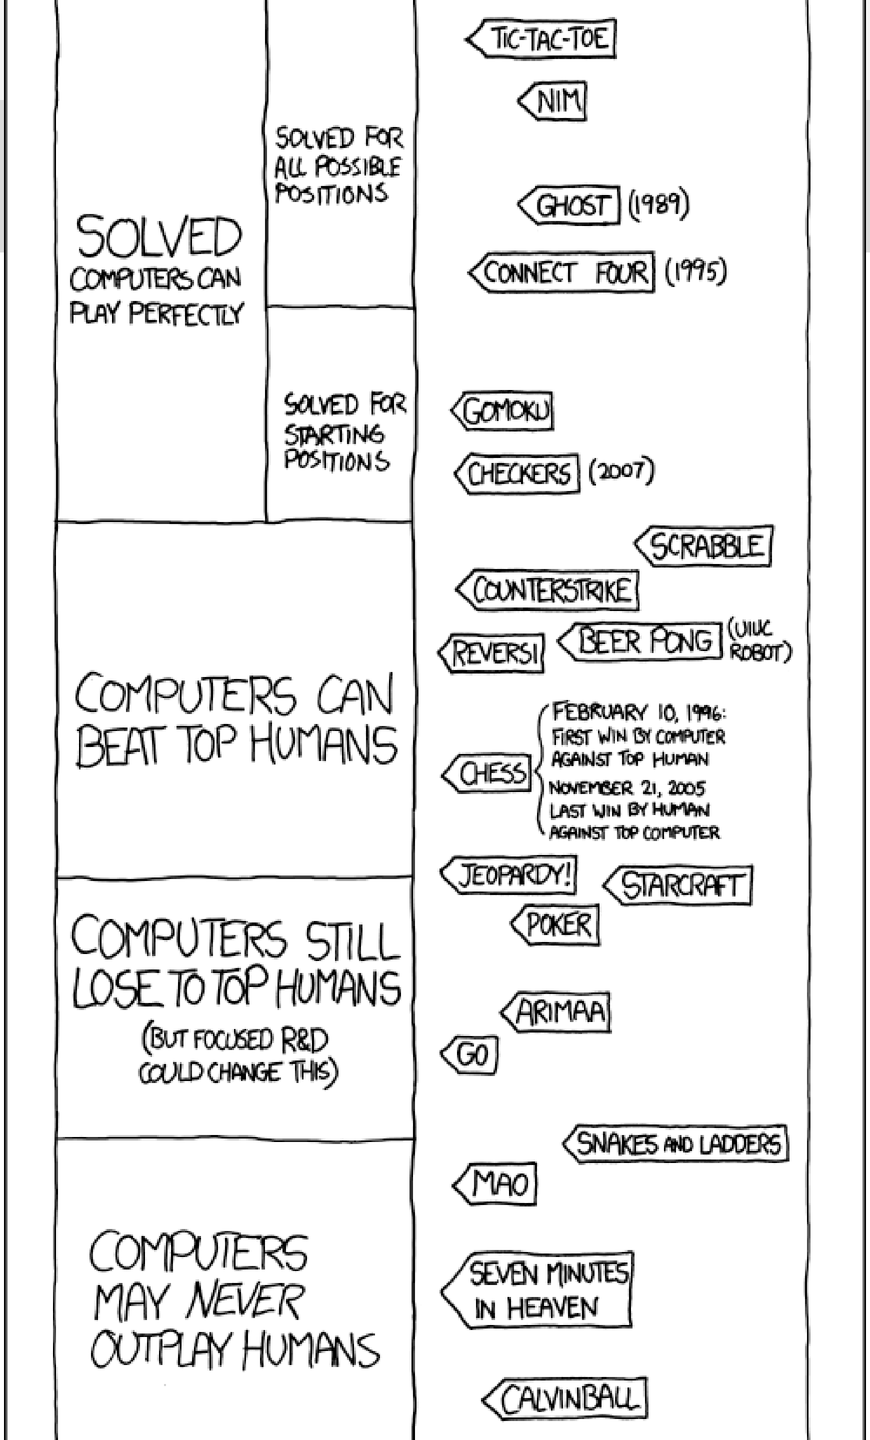

XKCD comics

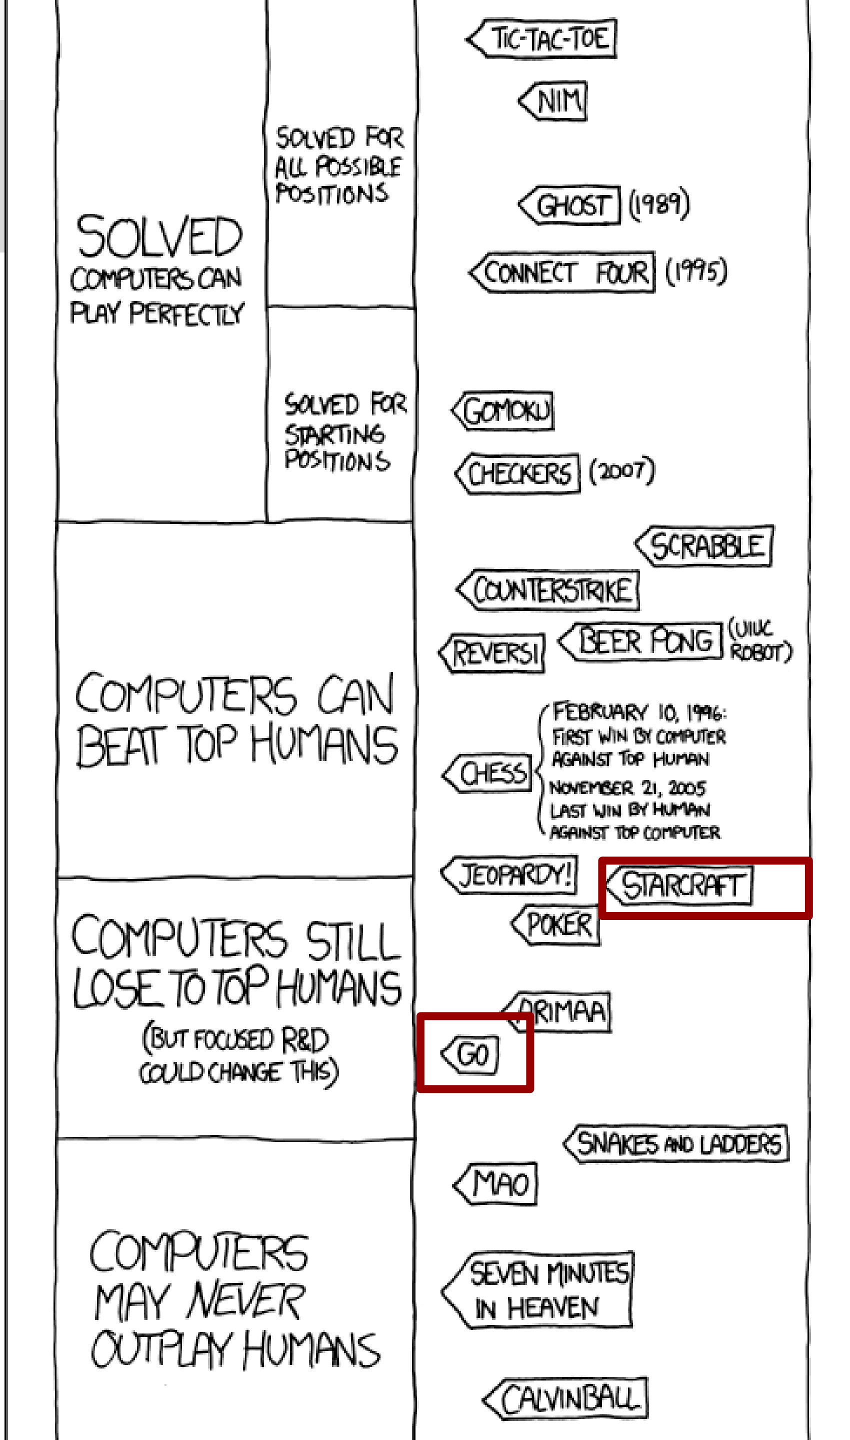
SOLVED
COMPUTERS CAN PLAY PERFECTLY
SOLVED FOR ALL POSSIBLE POSITIONS
TIC-TAC-TOE
NIM
GHOST (1989)
CONNECT FOUR (1995)
SOLVED FOR STARTING POSITIONS
GOMOKU
CHECKERS (2007)
COMPUTERS CAN BEAT TOP HUMANS
SCRABBLE
COUNTERSTRIKE
REVERSI
BEER PONG (UIUC ROBOT)
CHESS
FEBRUARY 10, 1996: FIRST WIN BY COMPUTER AGAINST TOP HUMAN
NOVEMBER 21, 2005 LAST WIN BY HUMAN AGAINST TOP COMPUTER
JEOPARDY!
COMPUTERS STILL LOSE TO TOP HUMANS
(BUT FOCUSED R&D COULD CHANGE THIS)
STARCRAFT
POKER
ARIMAA
GO
COMPUTERS MAY NEVER OUTPLAY HUMANS
SNAKES AND LADDERS
MAO
SEVEN MINUTES IN HEAVEN
CALVINBALL

Integrating Knowledge, Reasoning and Planning

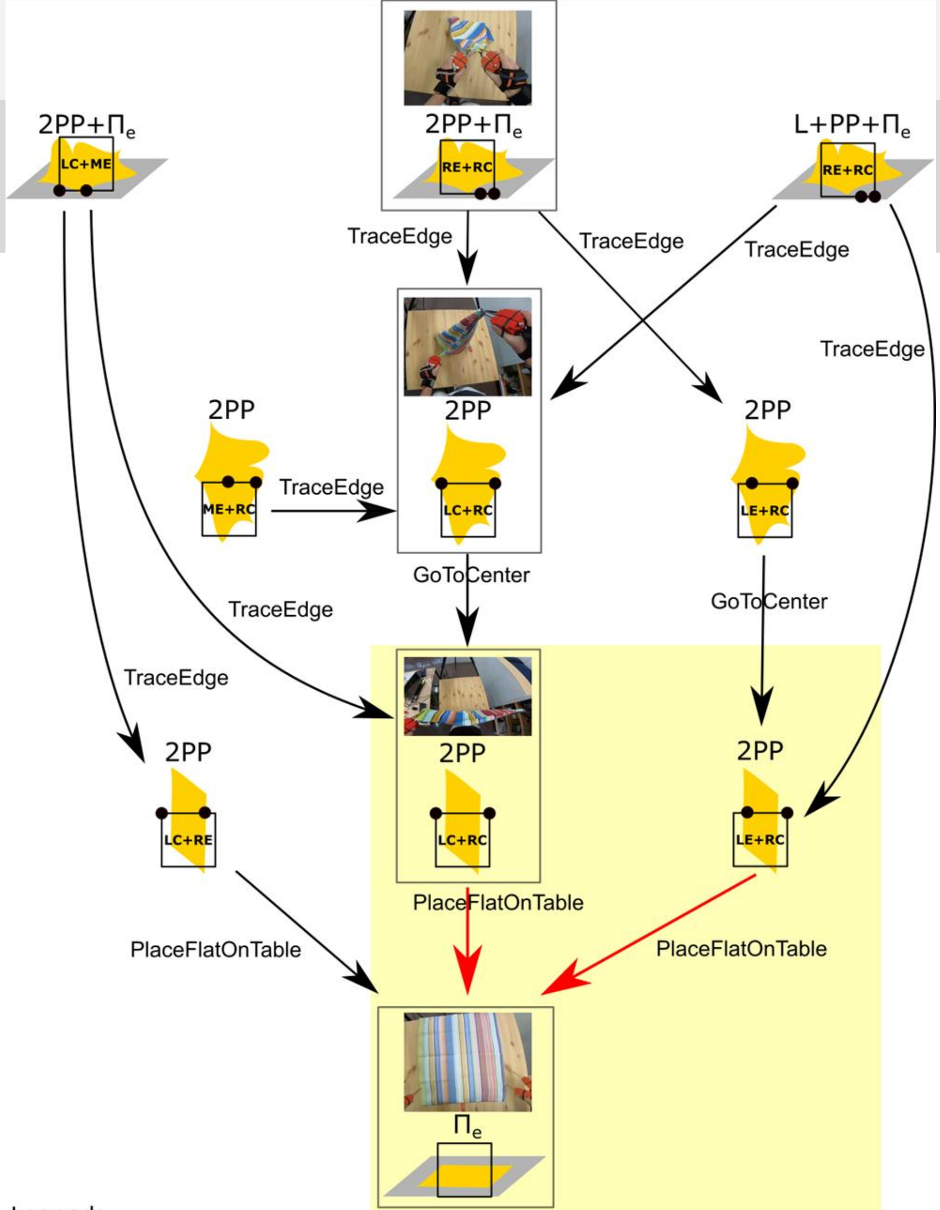

Legend:
State = { Grasp type = Geometric virtual fingers (P: point, L:Line, Π:Plane),
Cloth configuration = Flat or Crumpled ,
Grasp location: On the table: or hanging:
}

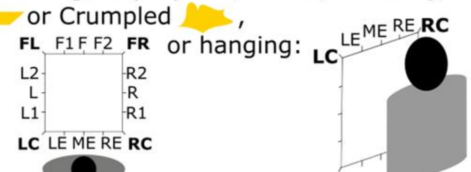

Integrating Knowledge, Reasoning and Planning


DALL·E
An Impressionist oil painting of sunflowers in a purple vase...
DALLE-2
FIRST LOOK

Integrating Knowledge, Reasoning and Planning

“Nikola Tesla playing soccer against Leo Messi, with vibrant colors”

KRP and data complexity

- Alphastar trained hundreds of agents with 200y of gameplay each.
- OpenAI 45000y of gameplay training could not beat top DOTA human.

KRP and data complexity

- Alphastar trained hundreds of agents with 200y of gameplay each.
- OpenAI 45000y of gameplay training could not beat top DOTA human.

What if we do not
have centuries of data
for training?

KRP and data complexity

OpenAI disbands its robotics research team

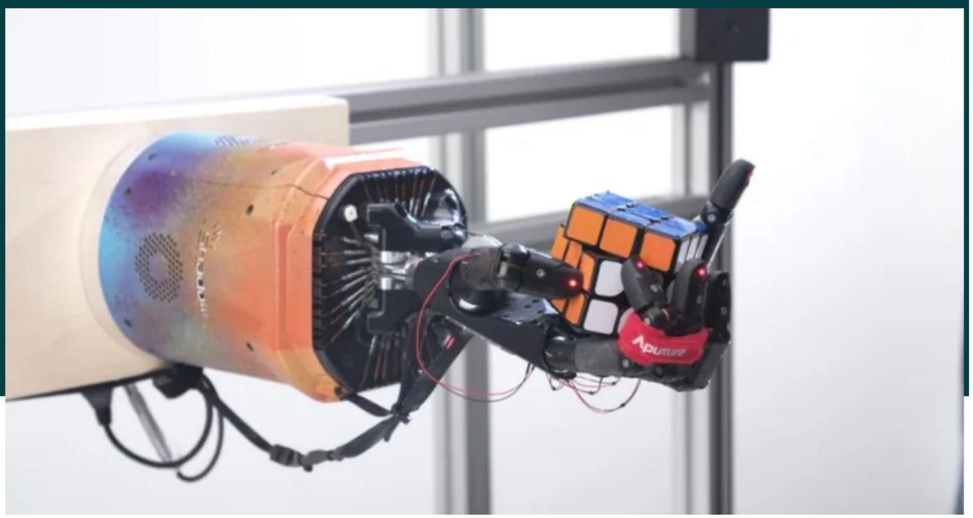

Image Credit: OpenAI

"There are actually plenty of domains that are very rich with data. Ultimately that was holding us back, in the case of robotics"

Wojciech Zaremba, cofounder of OpenAI

KRP and data complexity

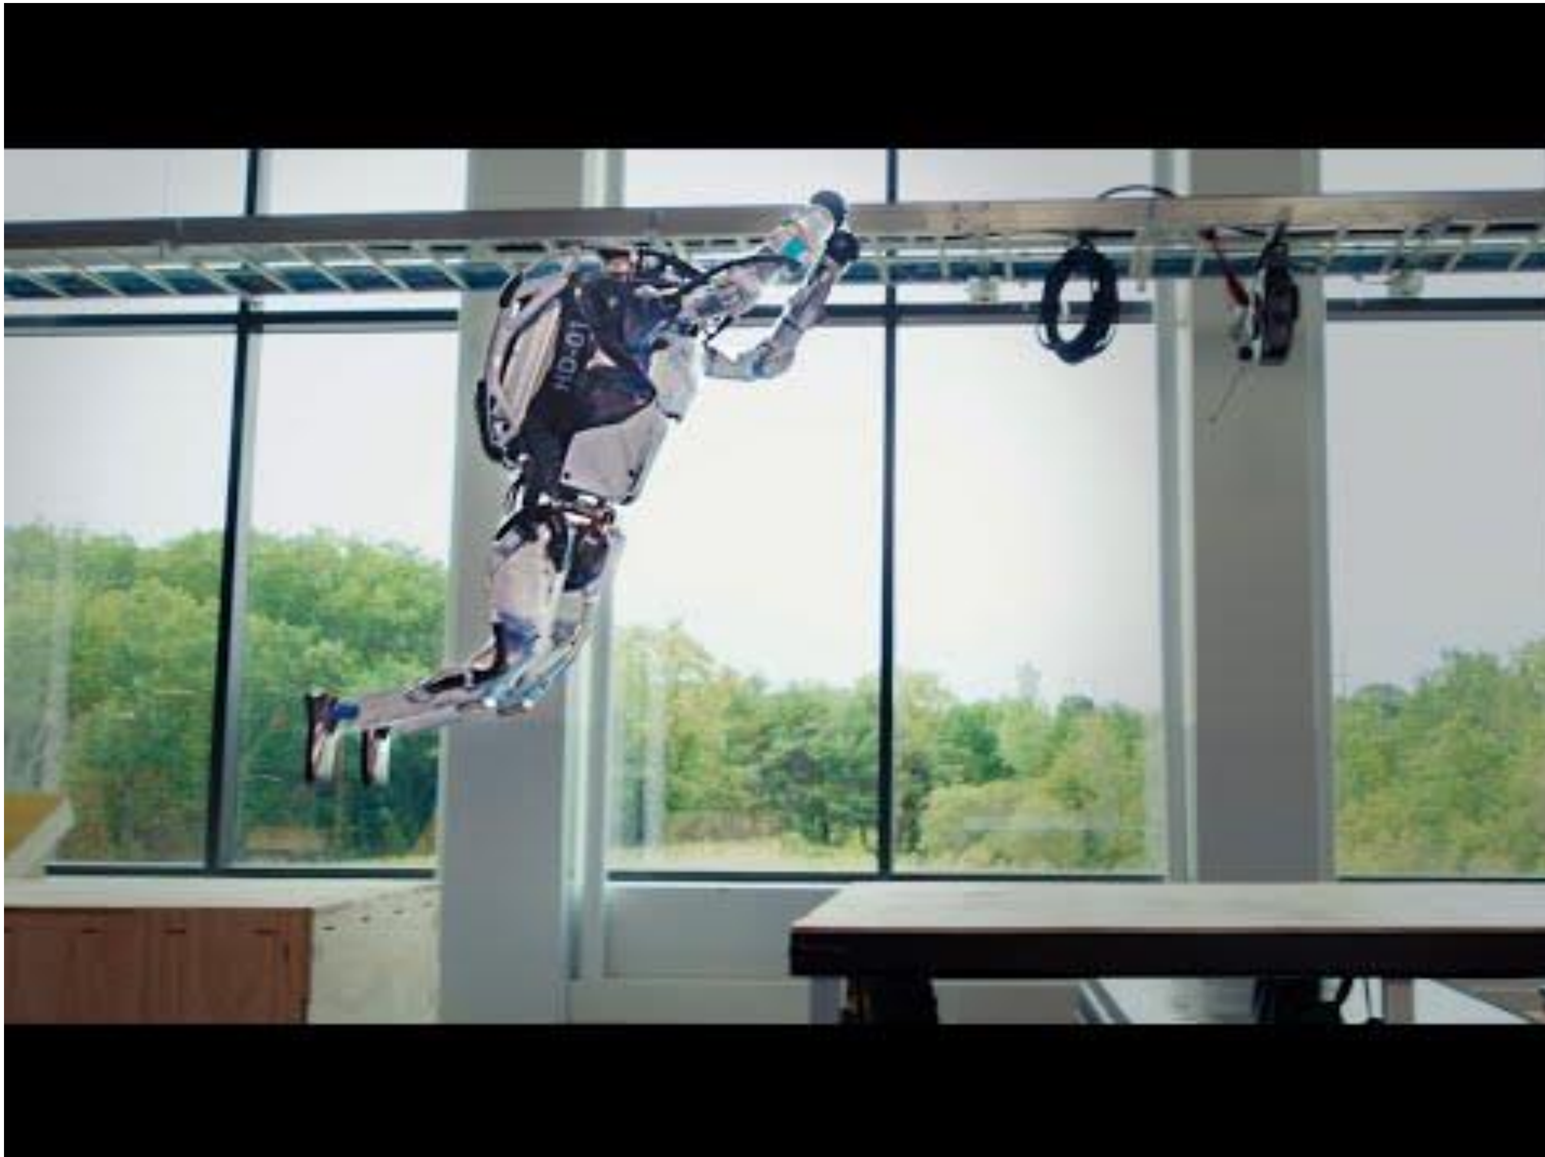

KRP and data complexity

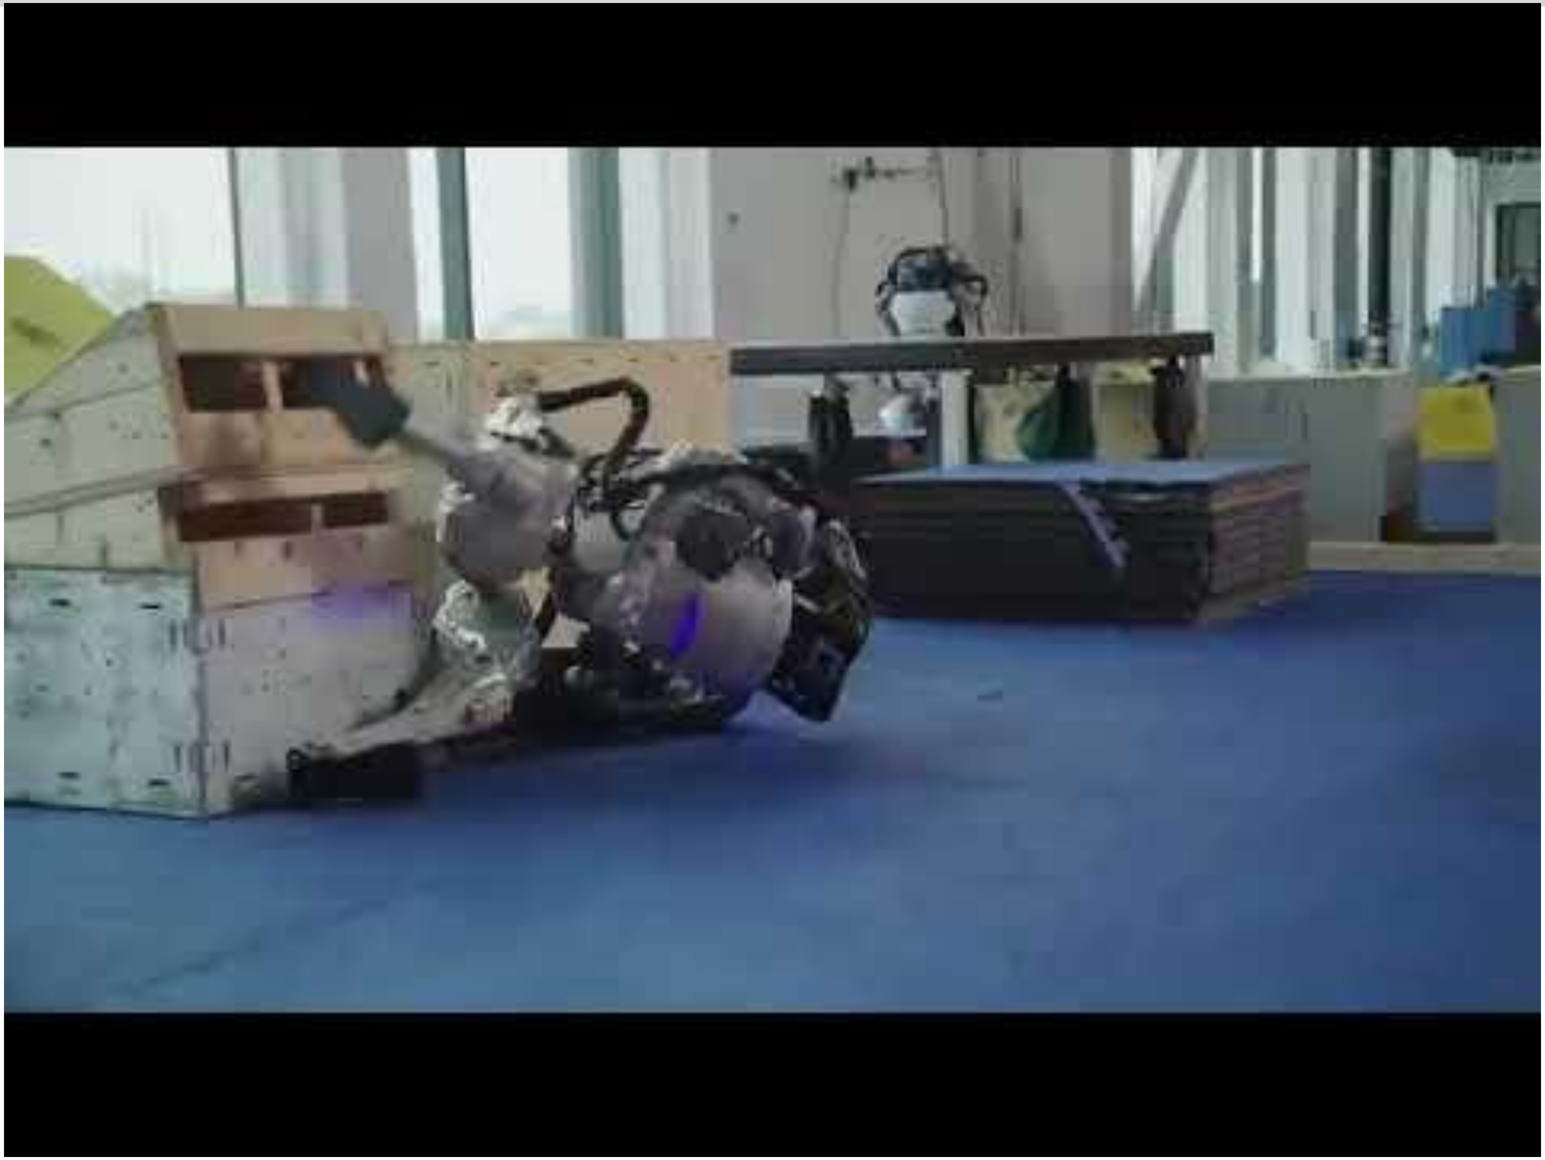

KRP and data complexity

1943: First NN by modelling how a brain works with an electric circuit

1958: First perceptron. Research in NN grows.

1969: Marvin Minsky kills research on perceptrons in his book.

AI winter

1982: US and JP start working on NN again, backpropagation is rediscovered

2000: Research on NN loosing steam. Not enough data/too costly.

AI research focused on Non-NN methods.

2006: Deep belief network. NN now called DL.

2012: Major breakthrough. AlexNet halves the competitors' error in ImageNet Large Scale Visual Recognition Challenge with a CNN. DL becomes the hot topic and the only option for certain fields. A new generation of DL researches is born.

Methodologies vs Data Volume

Big Data *is a combination of* ***insane volumes of structured, semi-structured, and unstructured data*** *that are too complex to be analyzed and processed by traditional data-processing techniques. (…)* ***Big Data*** *is quite different from the traditional concept of* ***small data*** *in terms of volume, velocity, variety, and veracity.* ***Small data****, on the contrary,* ***is data small enough to be conveniently stored on a single machine****, particularly local servers, or a laptop, and is easily accessible.*

Methodologies vs Data Volume

Many fields are now snowed under with an avalanche of data*, which raises considerable challenges for computer scientists.* ***Meanwhile, robotics (among other fields) can often only use a few dozen data points*** *because acquiring them involves a process that is expensive or time-consuming.* ***How can an algorithm learn with only a few data points?***

Micro-Data Learning: The Other End of the Spectrum by Jean-Baptiste Mouret (Inria)

Methodologies vs Data Volume

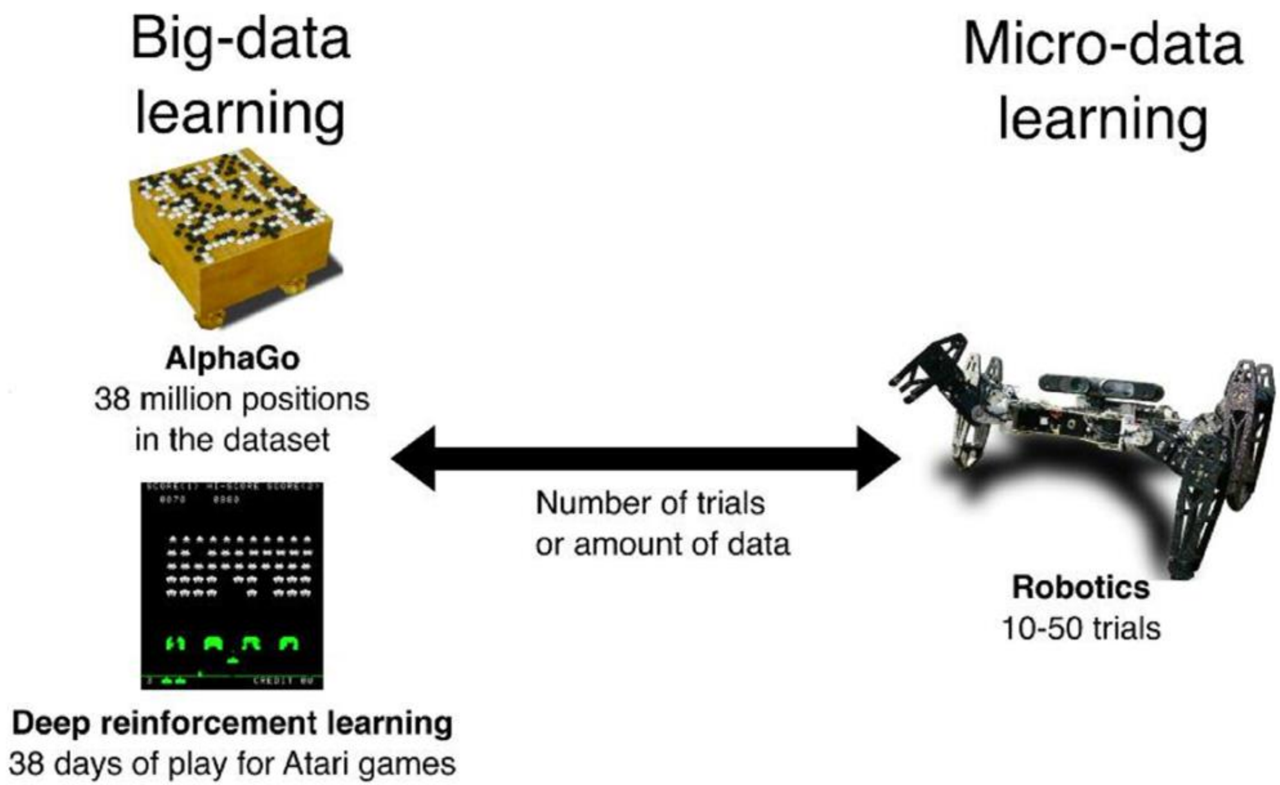

Methodologies vs Data Volume

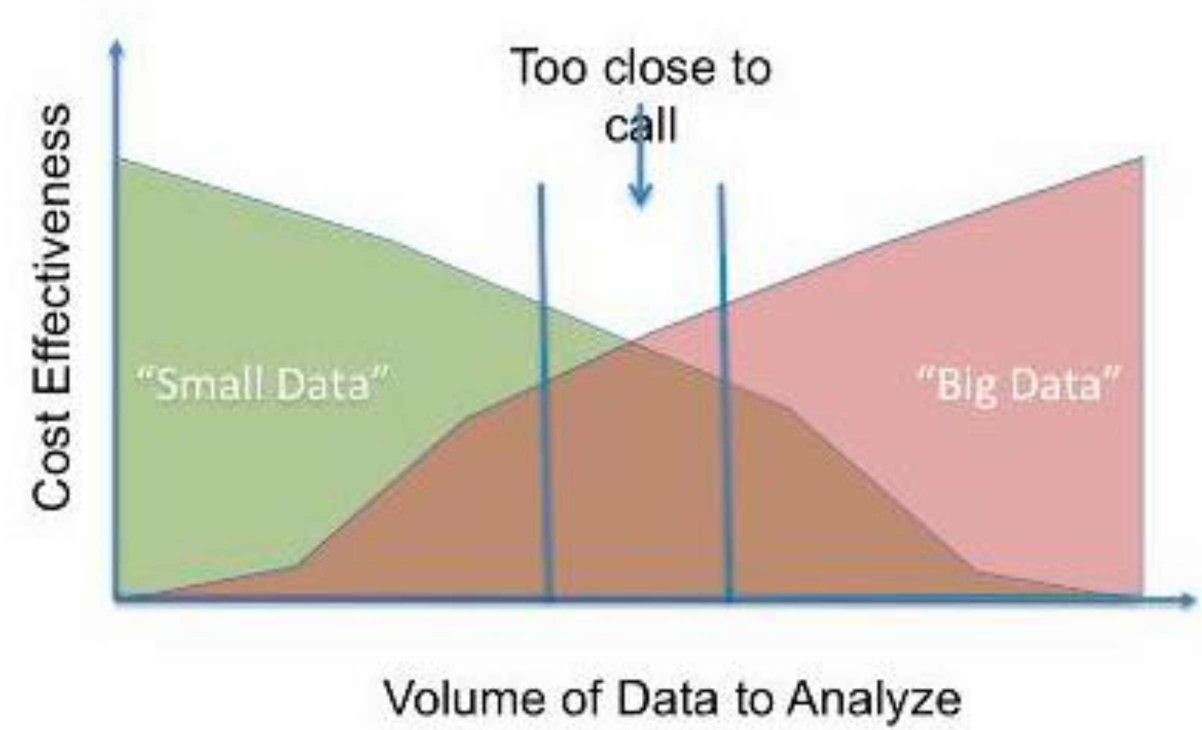

Methodologies vs Data Volume

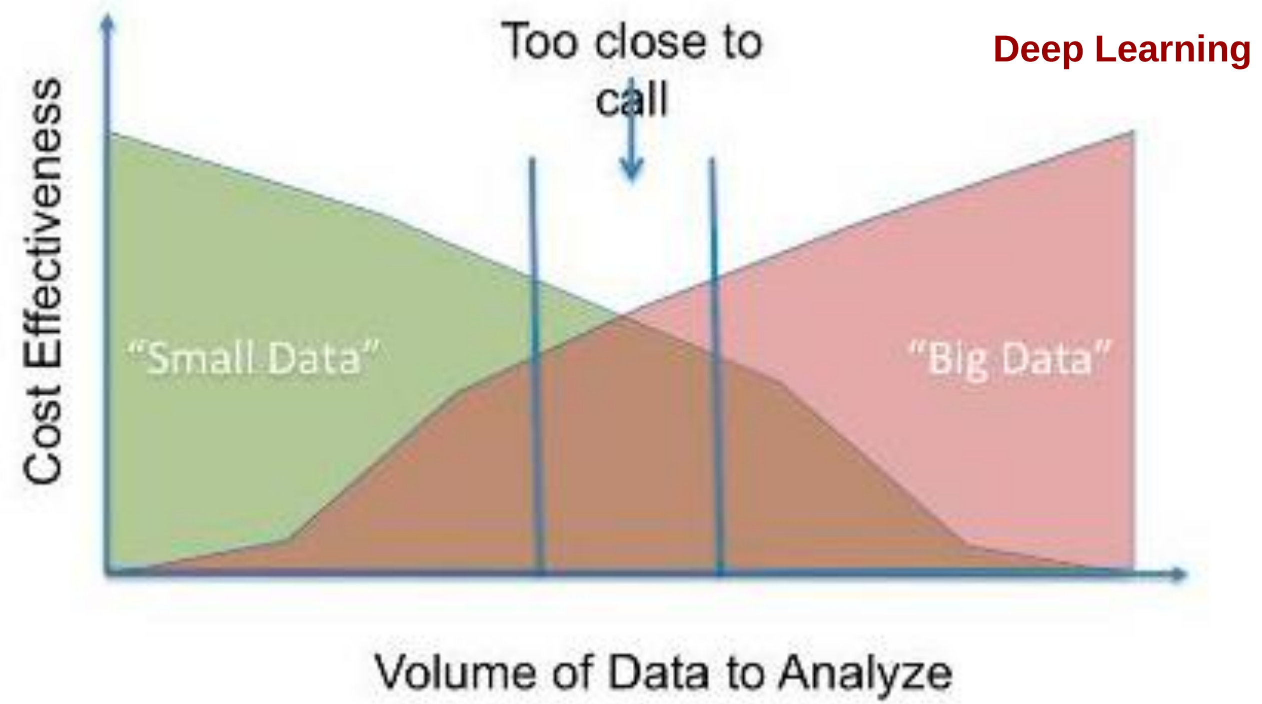

Methodologies vs Data Volume

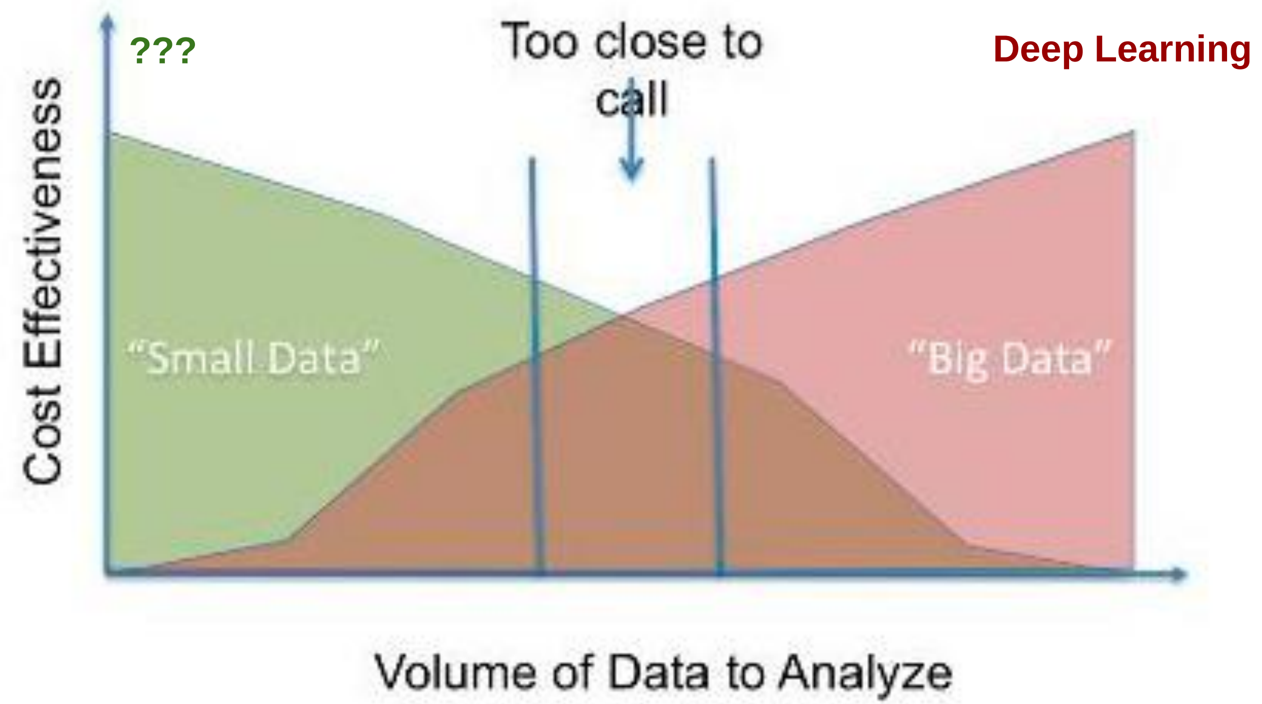

Bayesian Optimization

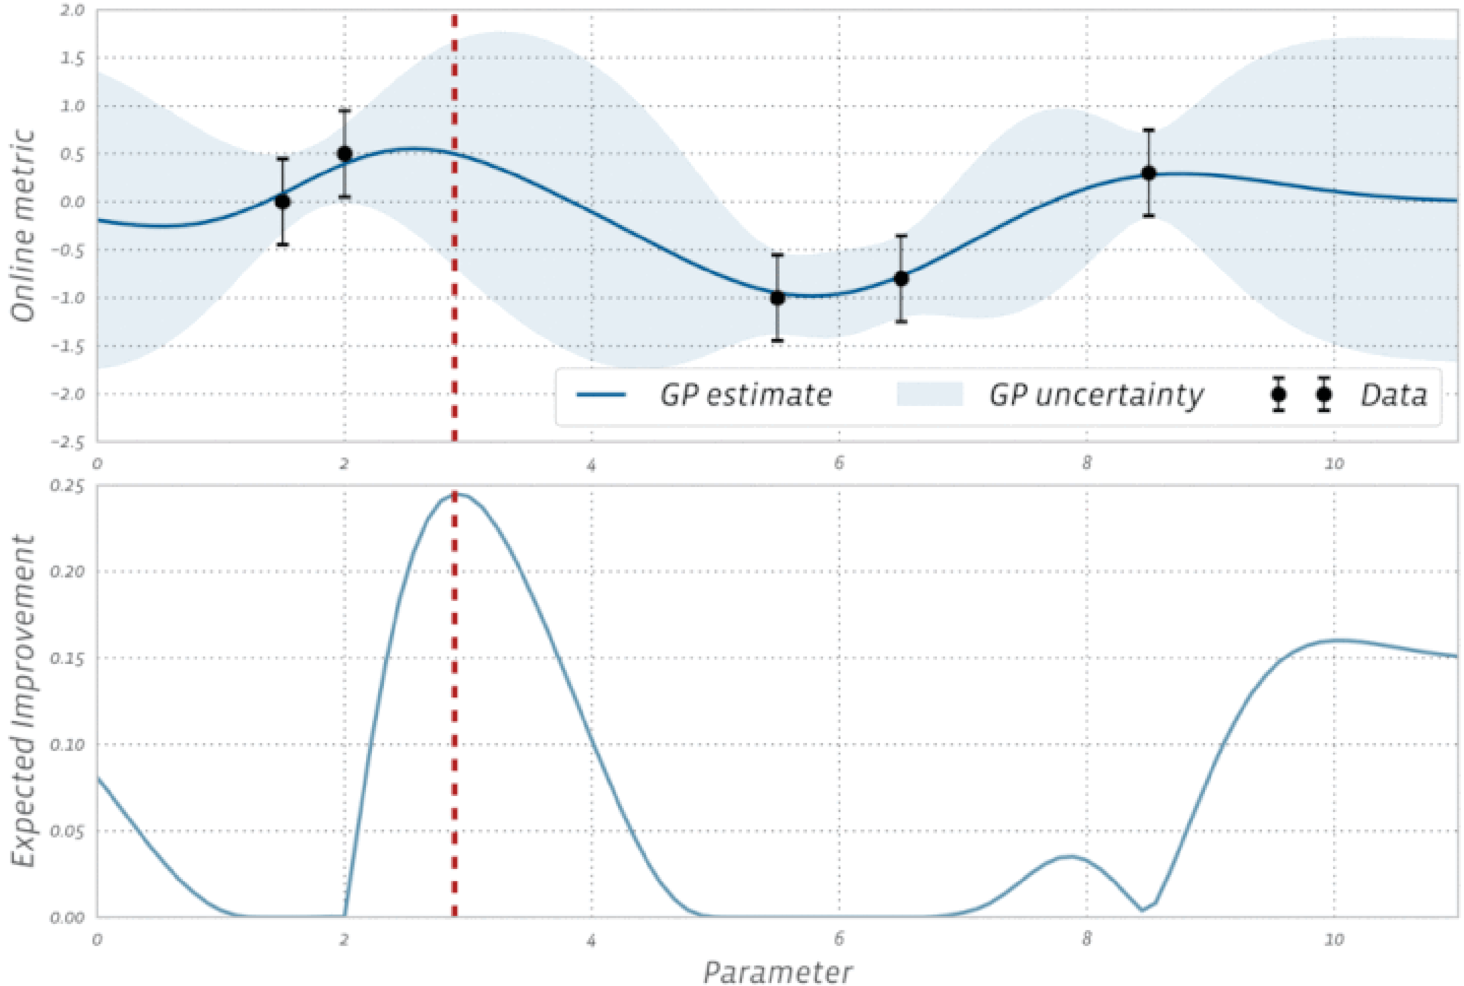

Source: B. Letham et al. Efficient tuning of online systems using Bayesian optimization. Meta

Bayesian Optimization

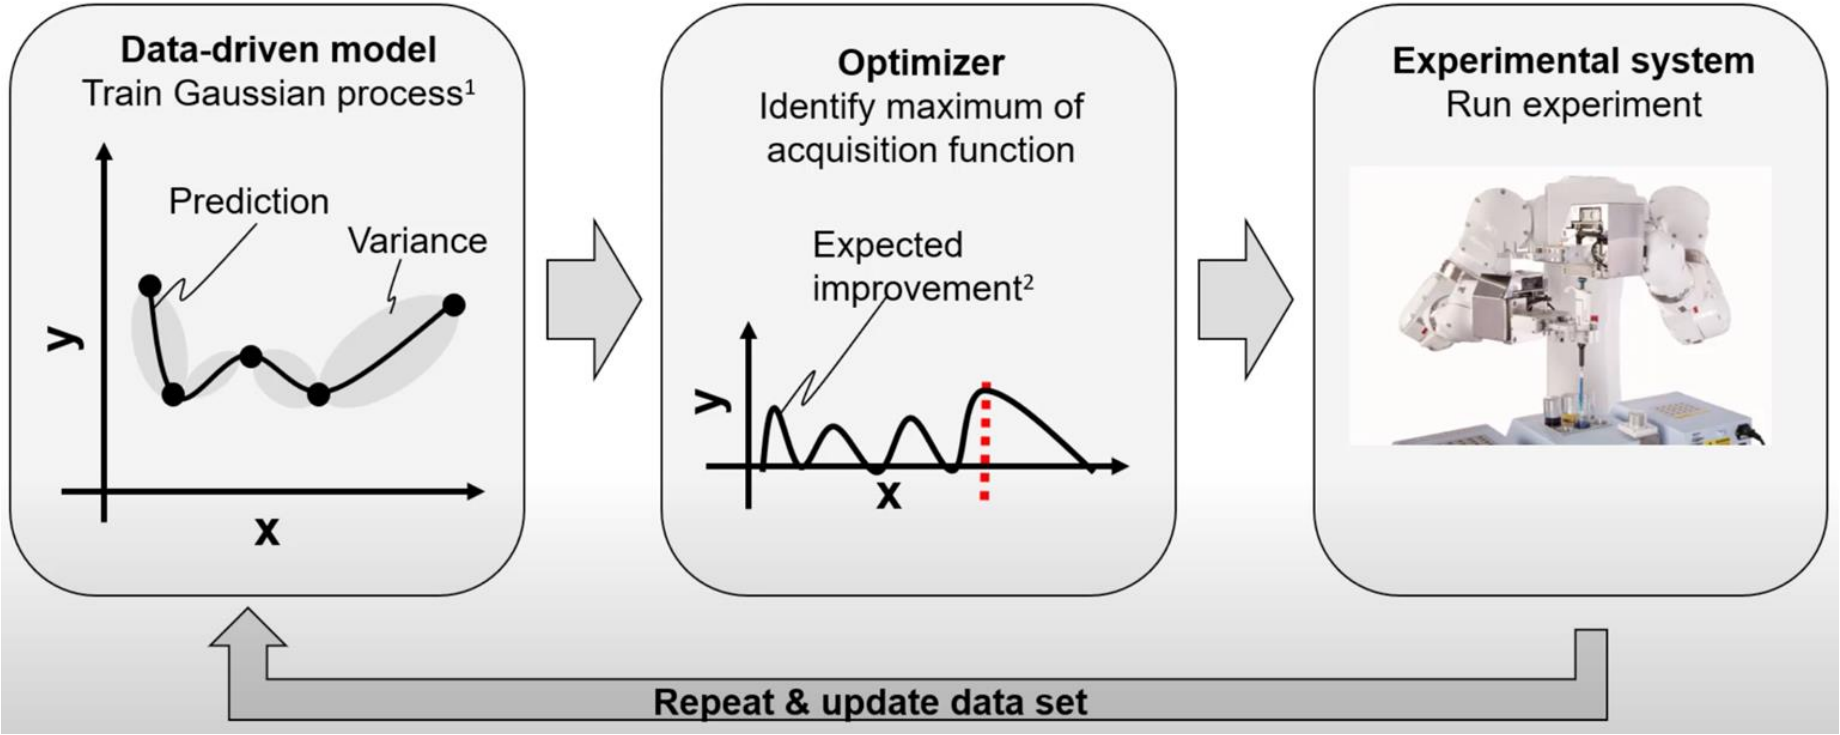

Source: A.M. Schweidtmann et al: *Maximizing the acquisition function of Bayesian optimization to guaranteed global optimality*

Bayesian Optimization

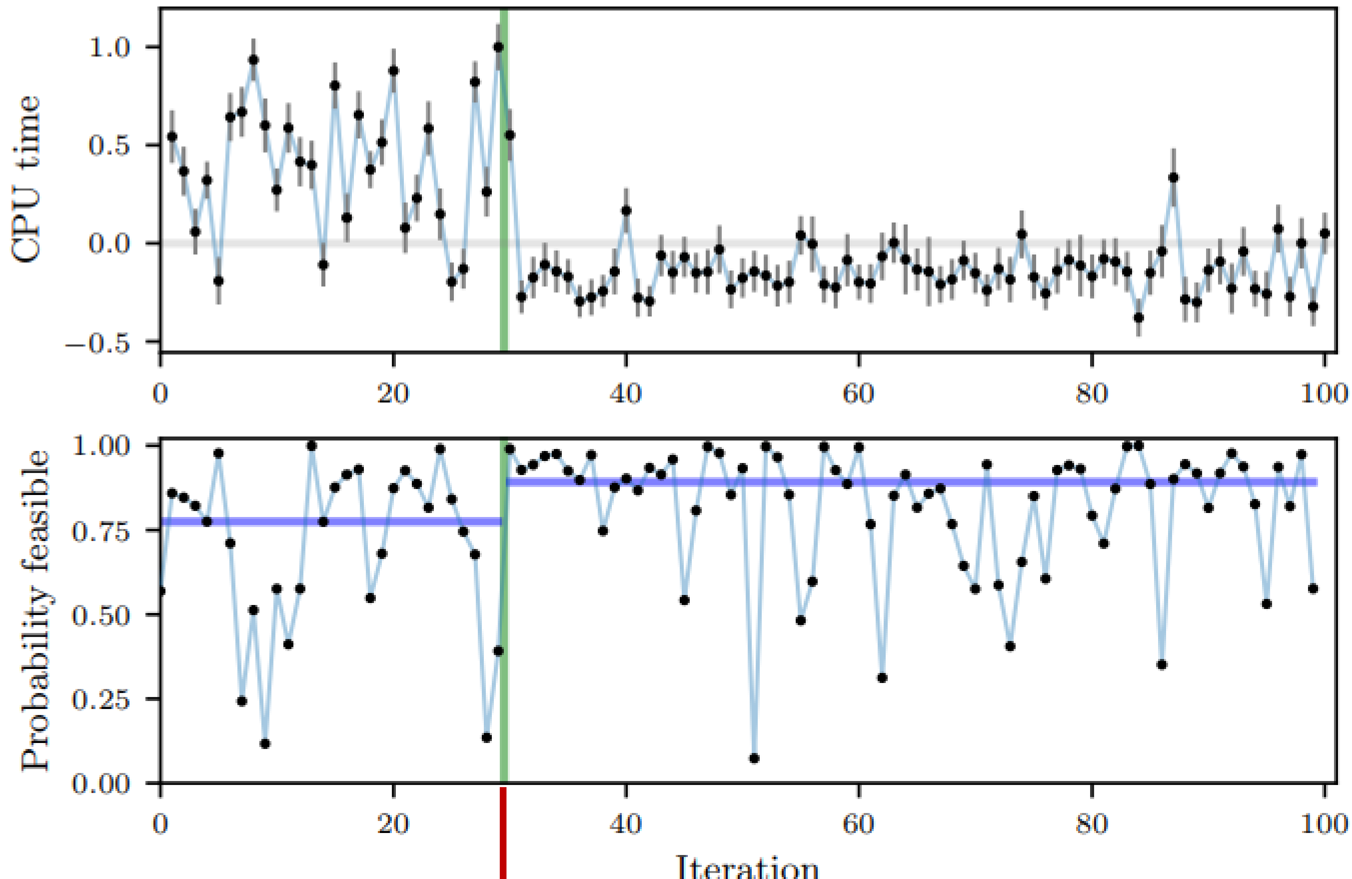

Random samples

Bayesian Optimization

Reinforcement learning

Reinforcement learning

Reinforcement learning

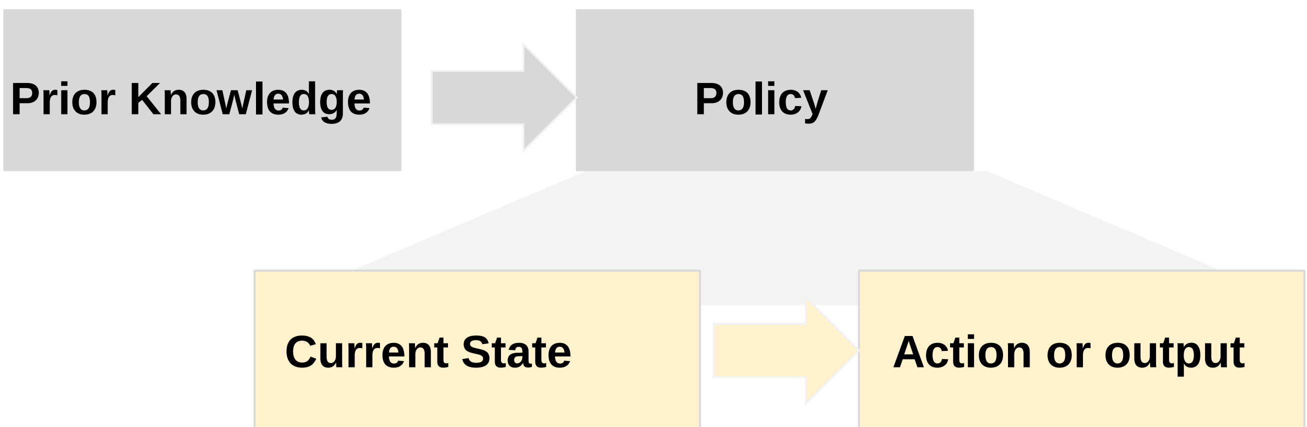

With the best expected outcome according to knowledge

Reinforcement learning

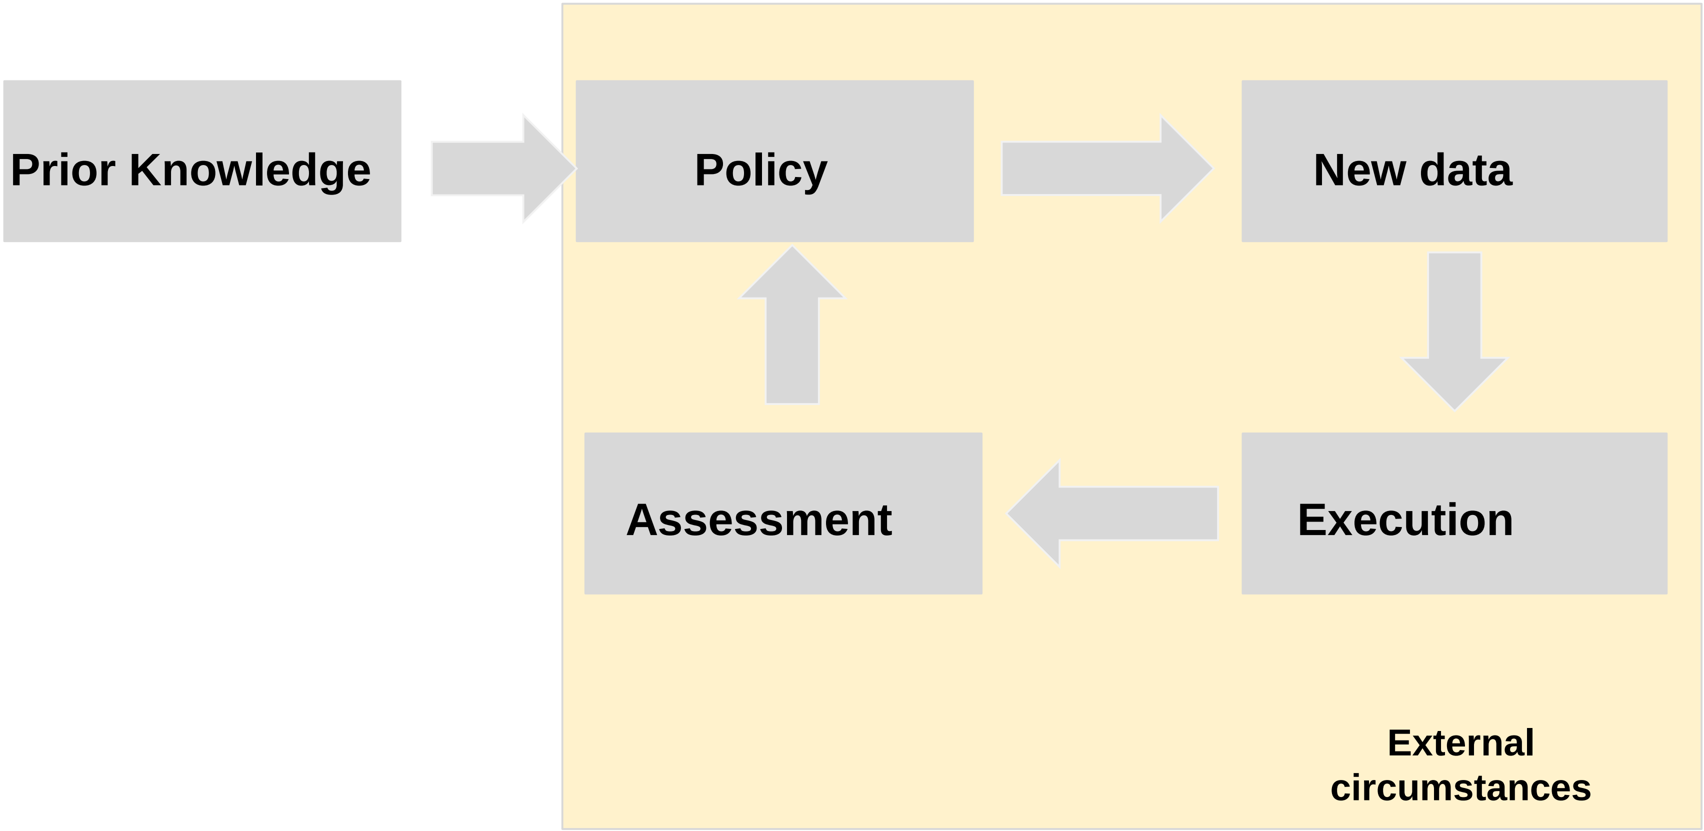

Reinforcement learning

Prior Knowledge → **Policy**

Reinforcement Learning

Given: (S, A, T, R)

States
Actions
Transitions
Rewards

→ π Policy

Reinforcement Learning

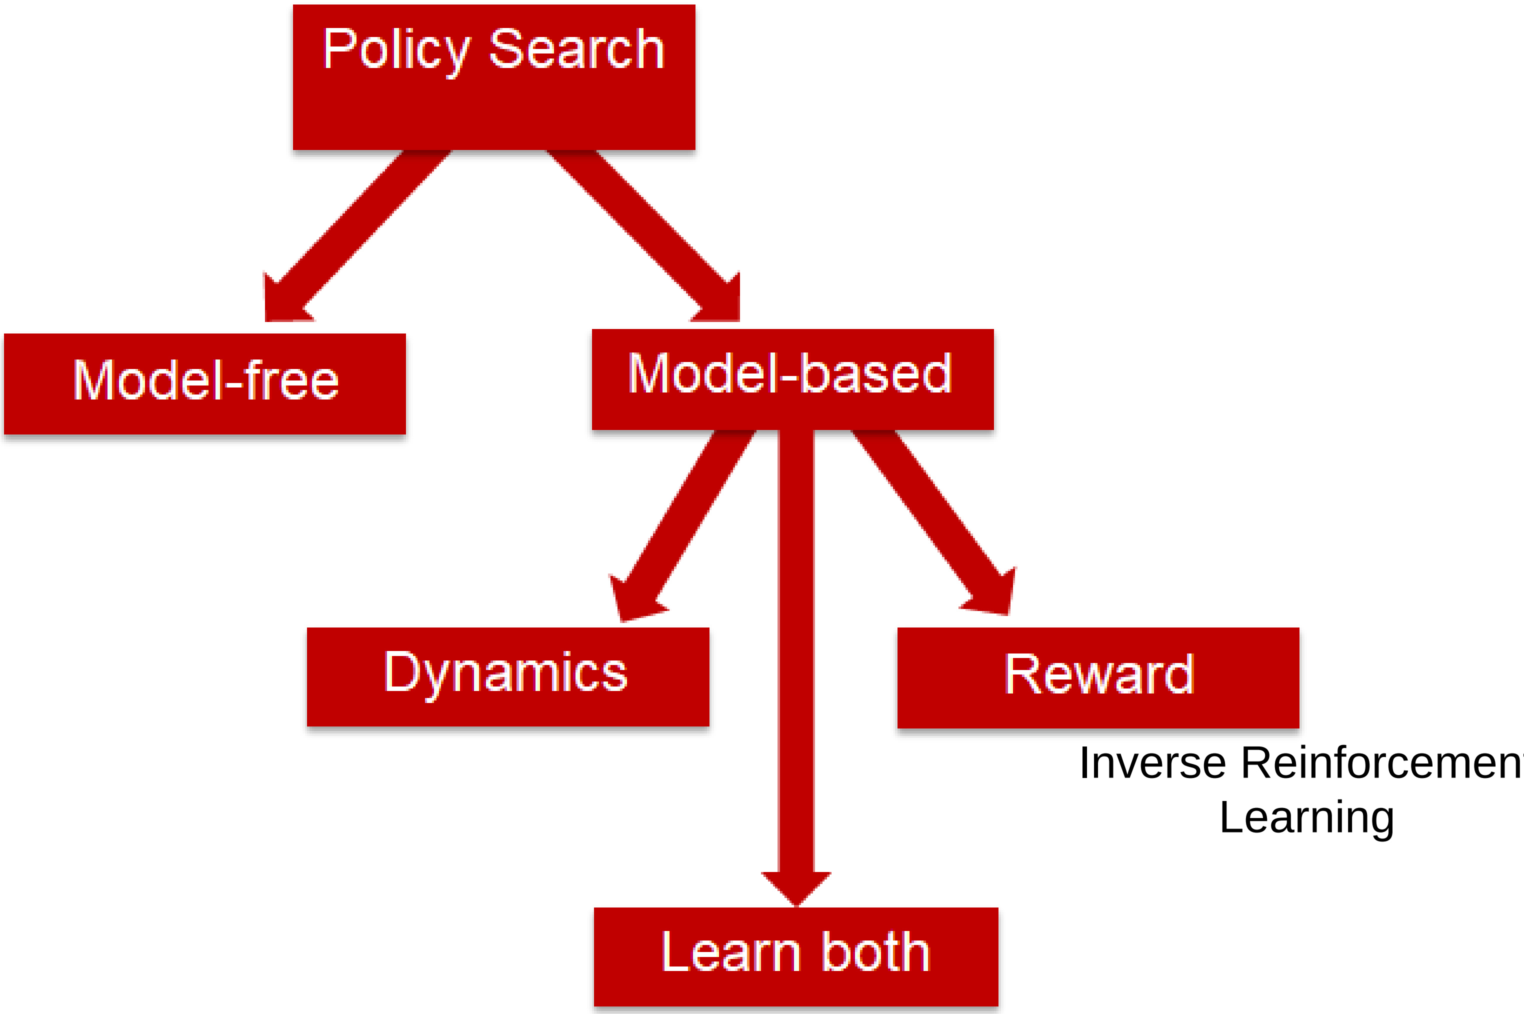

Reinforcement Learning

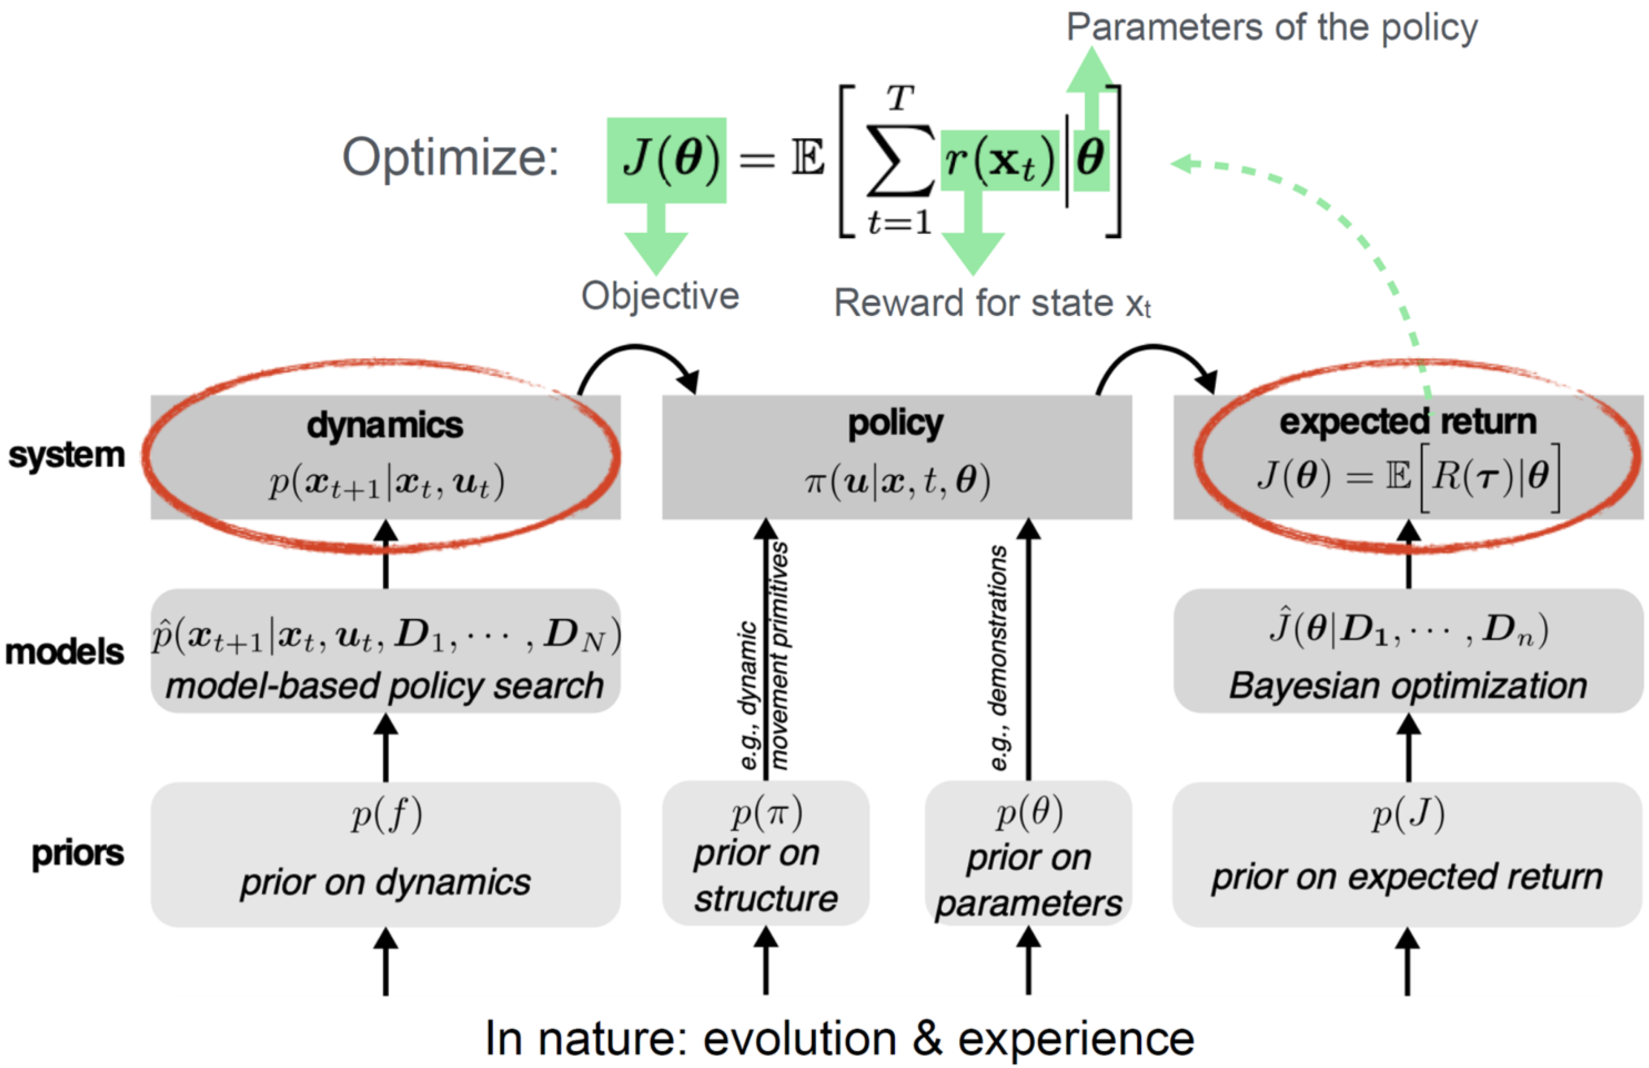

K. Chatzilygeroudis, V. Vassiliades, F. Stulp, S. Calinon, J.-B. Mouret. (2018). A survey on policy search algorithms for learning robot controllers in a handful of trials. arXiv:1807.02303

Transfer Learning

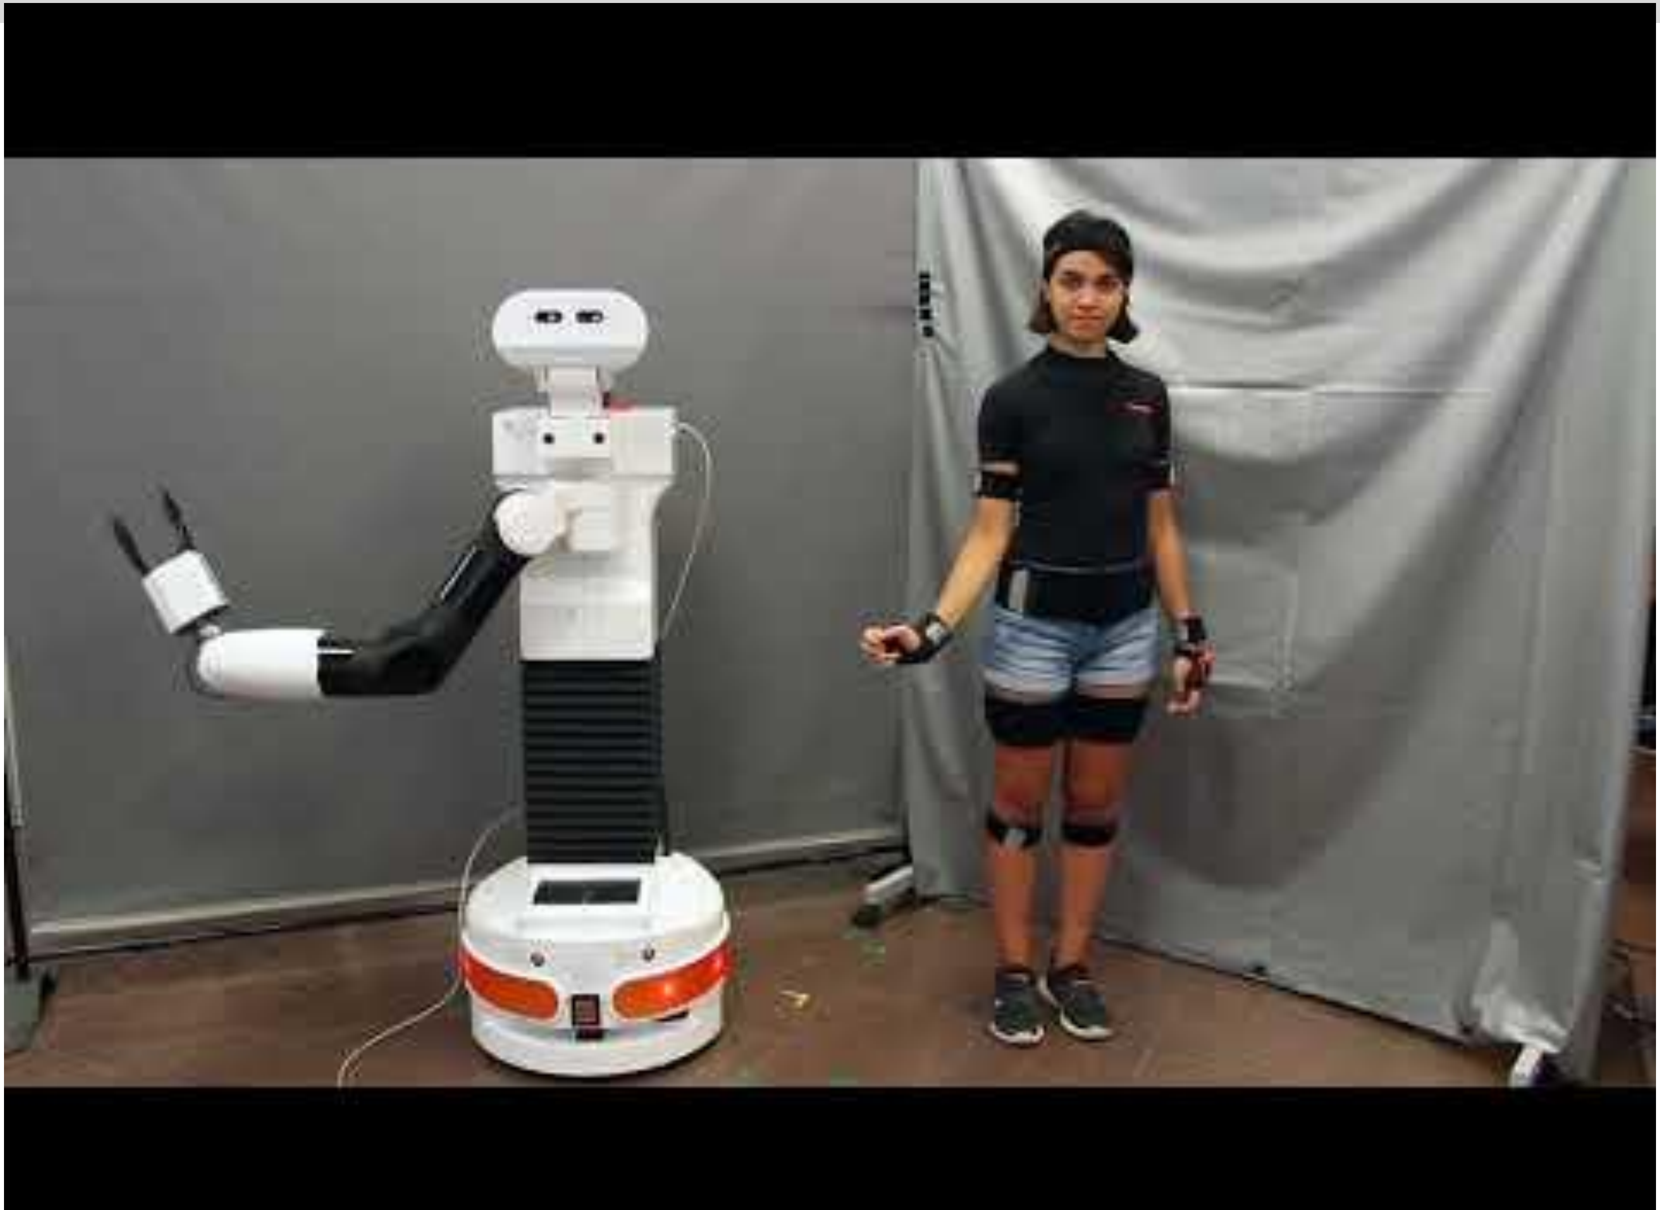

Meta Learning

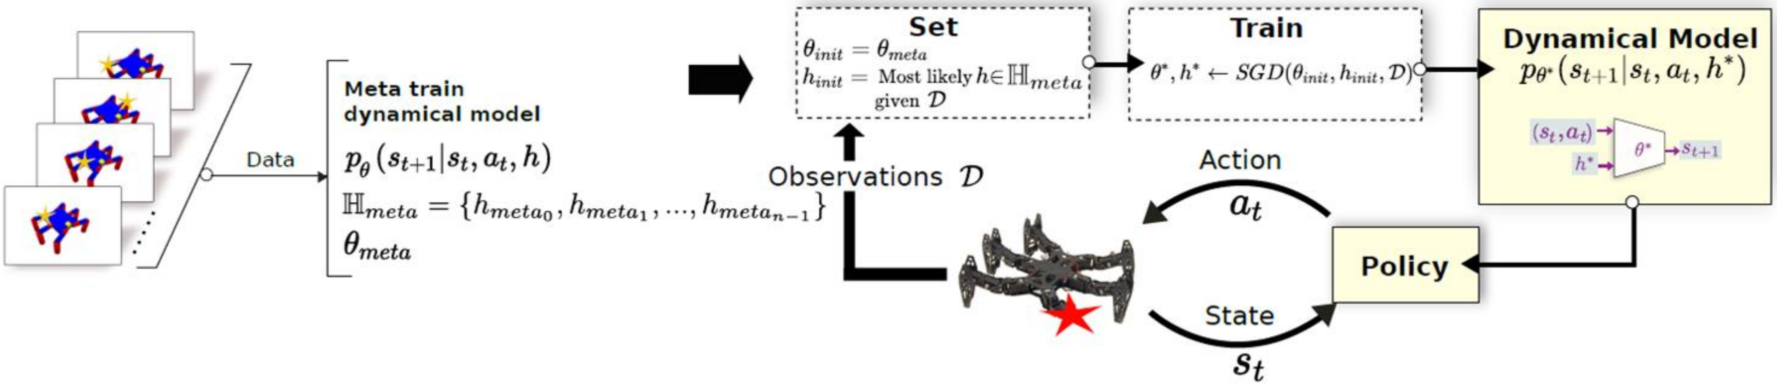

Hindsight Experience Replay (HER)

Assume failed executions actually targeted the observed end-point and add them to data as such

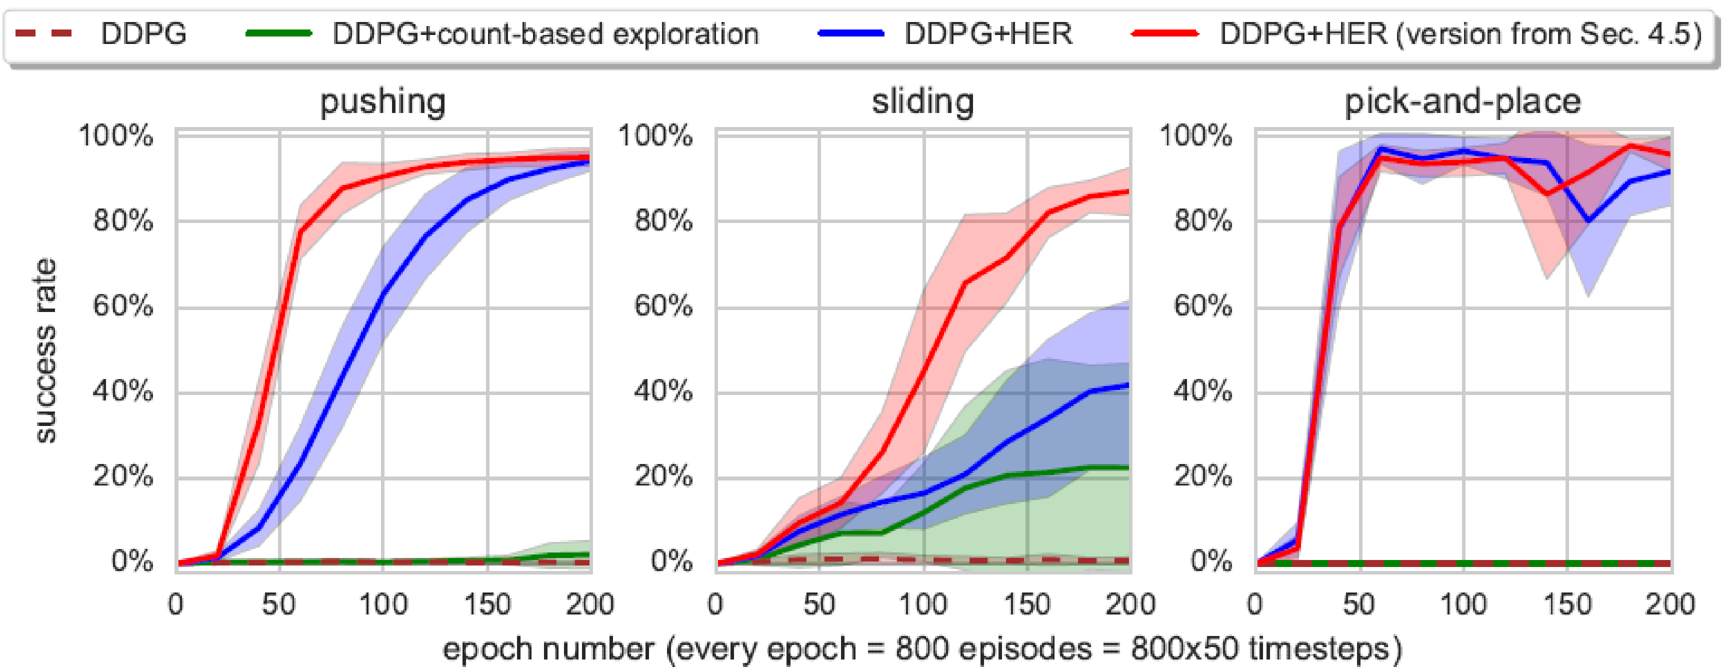

Conclusions

- Know your problem, the type/amount of data you have
  - *DL is suitable for big data*
  - *Other methods (such as GPs) are more suitable for small data*
- If using black-box methods:
  - *Be careful with the data you provide*
  - *Check the output thoroughly*

" if you torture the data long enough,

it will confess to anything "

Thank you!

Questions