Mind everywhere, AI of Things, and the future of engineering

Balazs Kegl, Noah's Ark Research Lab, Paris

Joint work with Abdelhakim Benechehab, Youssef Attia El Hili, Othman Gaizi, Jonas Gonzalez, Gabriel Hurtado, Giuseppe Paolo, Albert Thomas, Chunchun Yang

www.huawei.com

HUAWEI TECHNOLOGIES CO., LTD.

Part I

AI in engineering systems

AI Embodiment Through 6G: Shaping the Future of AGI

Lina Bariah and Mérouane Debbah

- Additionally, the paper examines **the limitations of current AI models,** particularly Large Language Models (LLMs), in achieving true AGI.
- It argues that while LLMs have made significant strides in natural language understanding, they still **lack the depth of knowledge representation and reasoning** required for AGI.
- The authors suggest that **sensory grounding in the real world is necessary** but not sufficient for LLMs to achieve robust understanding.
- They propose that a blend of technologies, including **generative language models, internet-of-senses, reinforcement learning, and edge intelligence**, integrated harmoniously via 6G networks, will be pivotal in realizing the vision of AGI.

18:22 | 0.6KB/s

Description

#7 Anna Ciaunica: embodied cognition and depersonalization

8	139	Oct 24
Likes	Views	2023

I, Scientist

00:00:00 Intro: empirical vs armchair philosopher. Visual vs tactile understanding.
00:06:08 How subjective experience rises from physical matter. Entering wine into your body vs seeing a

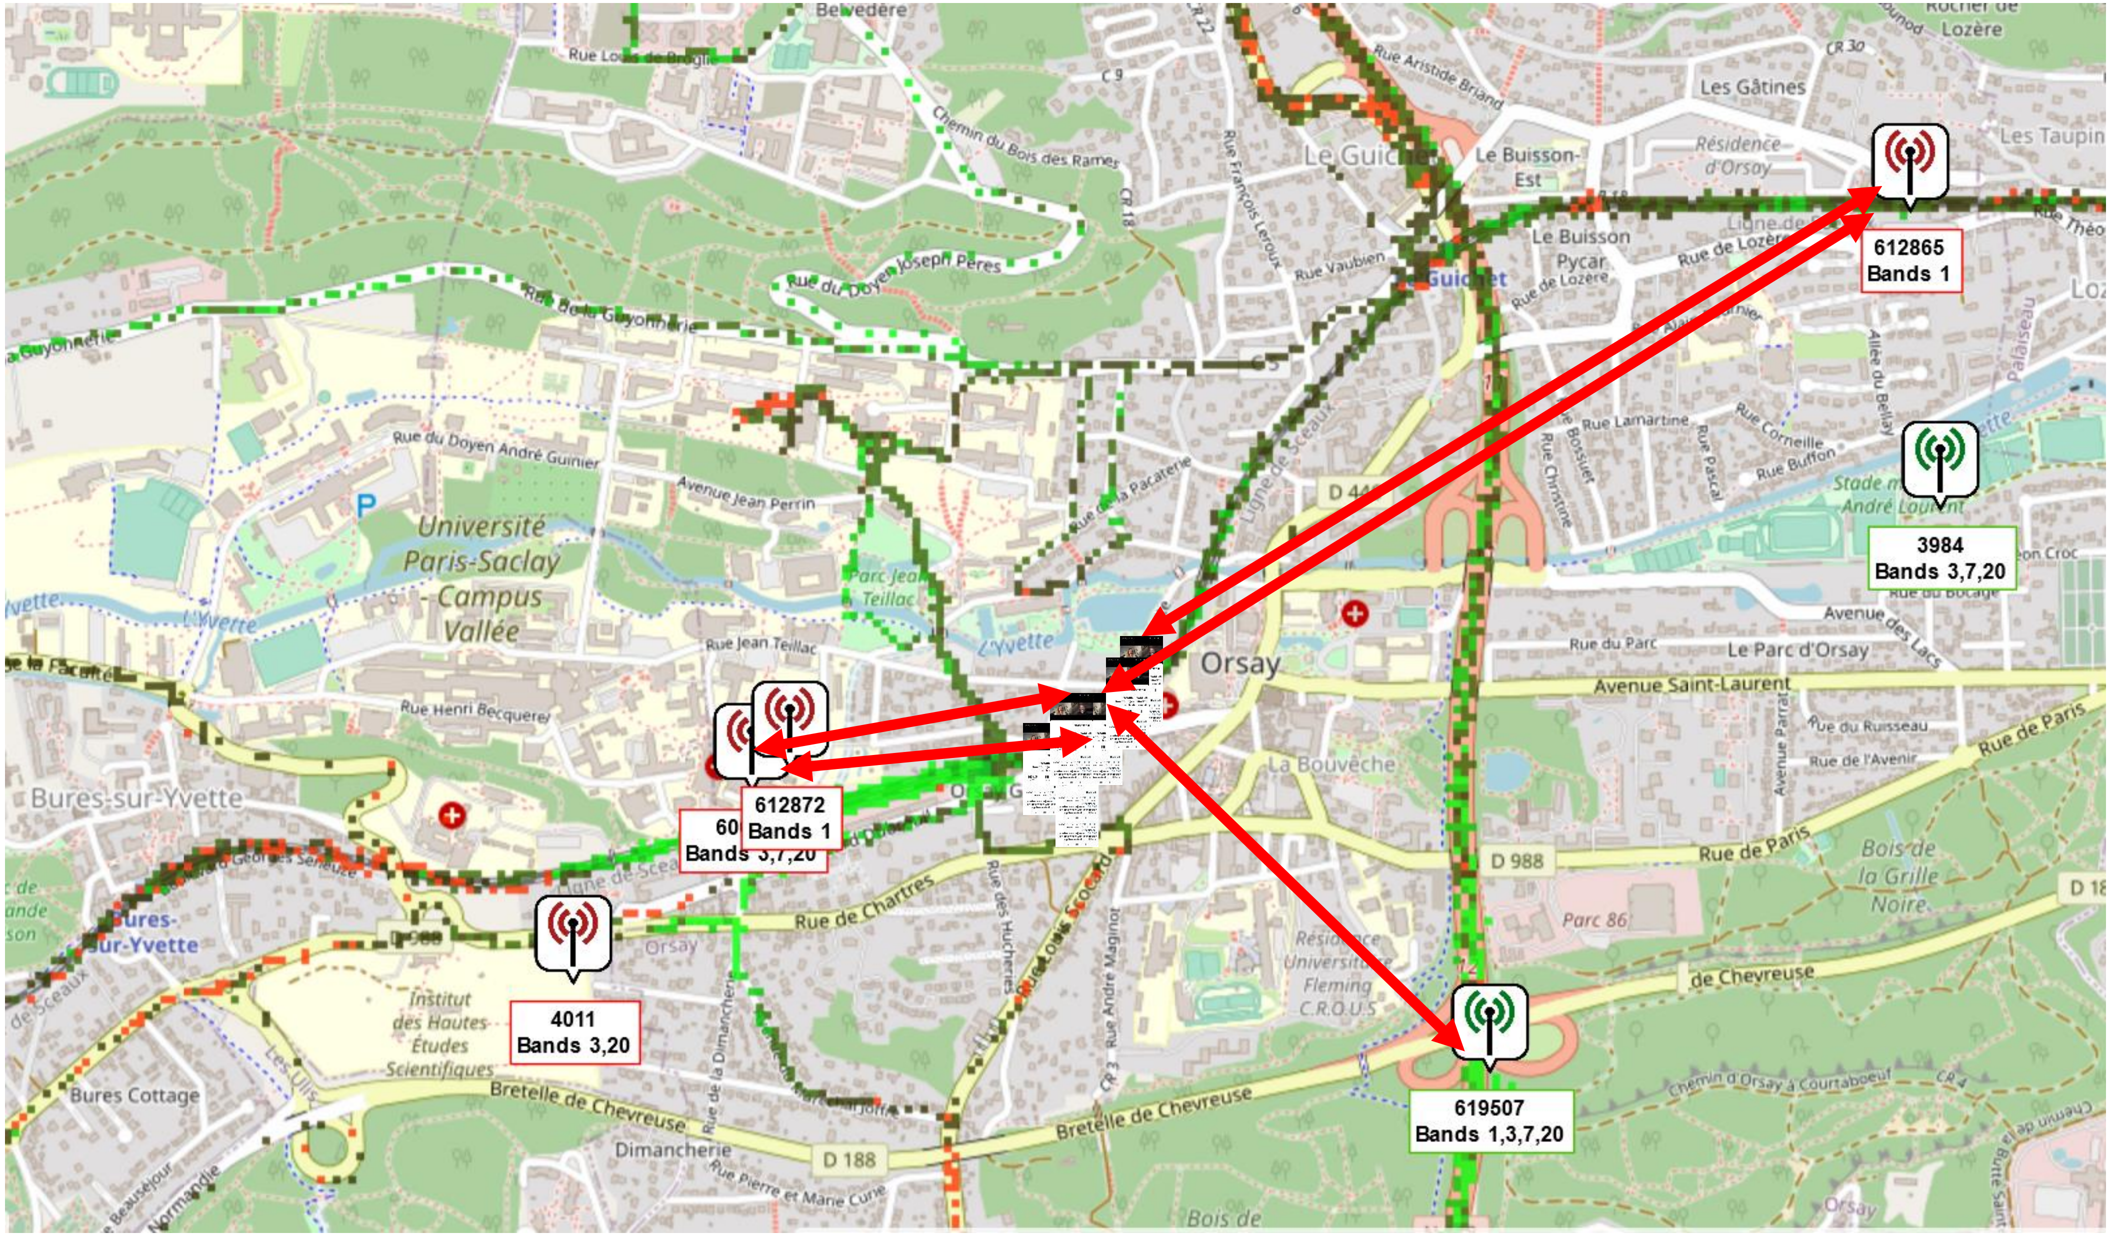
612865
Bands 1
3984
Bands 3,7,20
612872
Bands 1
60
Bands 3,7,20
4011
Bands 3,20
619507
Bands 1,3,7,20

A typical engineering control system

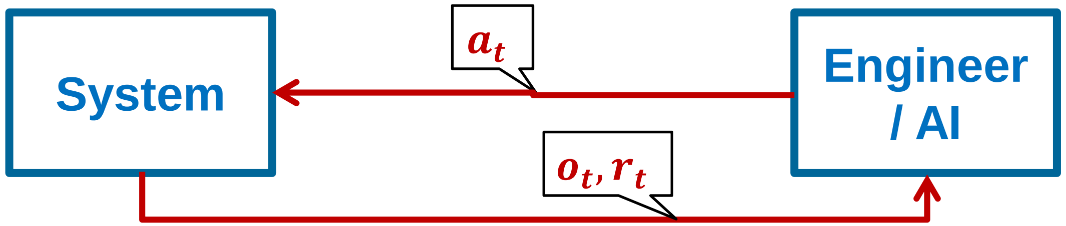

Engineer or AI observes

system states and **performance indicators**,

tunes some parameters time to time,

to **optimize the performance indicators**

Engineering systems = ~$10s of trillions per year

Our use cases

- **Self-driving engineering systems**
  - Wireless parameter tuning
  - Wi-Fi parameter tuning
  - Data center cooling
  - Energy, smart cities, etc.
- **Making them**
  - Safer, better, more reliable, more energy efficient
- **These are only the tip of the iceberg**

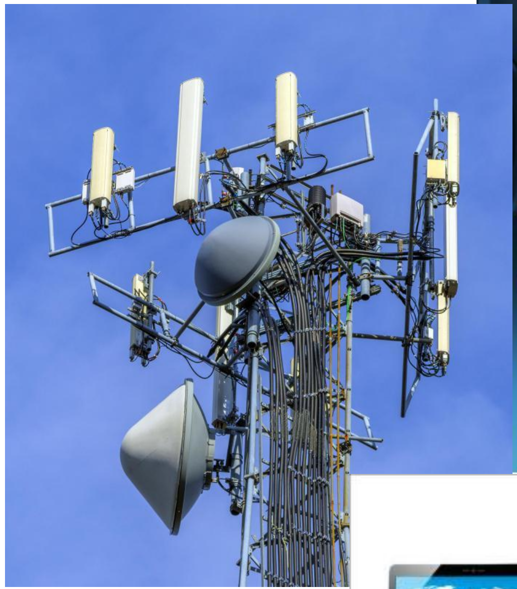

What is AI (in this context)?

Learn the system behavior
based on **historical data**
and use it for better control

Controlled engineering system: organizational constraints

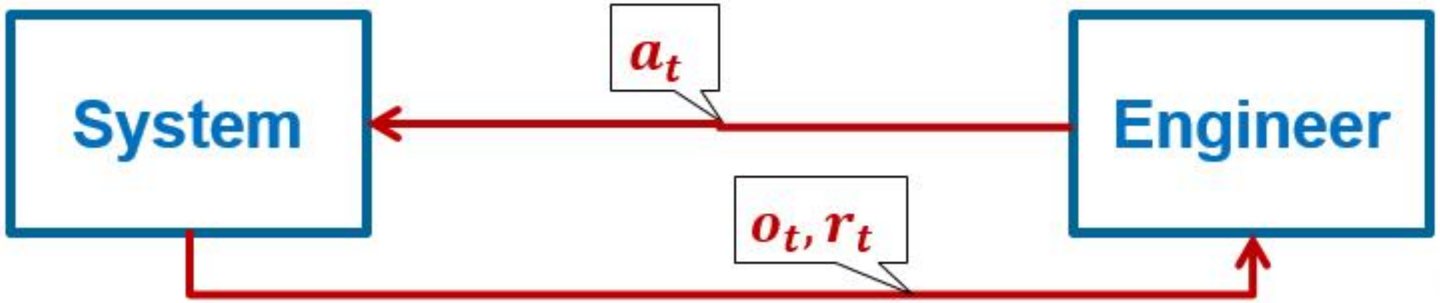

- **Offline (batch)**: system traces (logs)
- **Micro-data**: physical systems, slow white-box simulators
- **Safety**: we cannot "lose" while learning
- **Multi-agent**: multiple interdependent copies of the same system

Part II

Single-agent intelligence

Iterated offline/batch RL

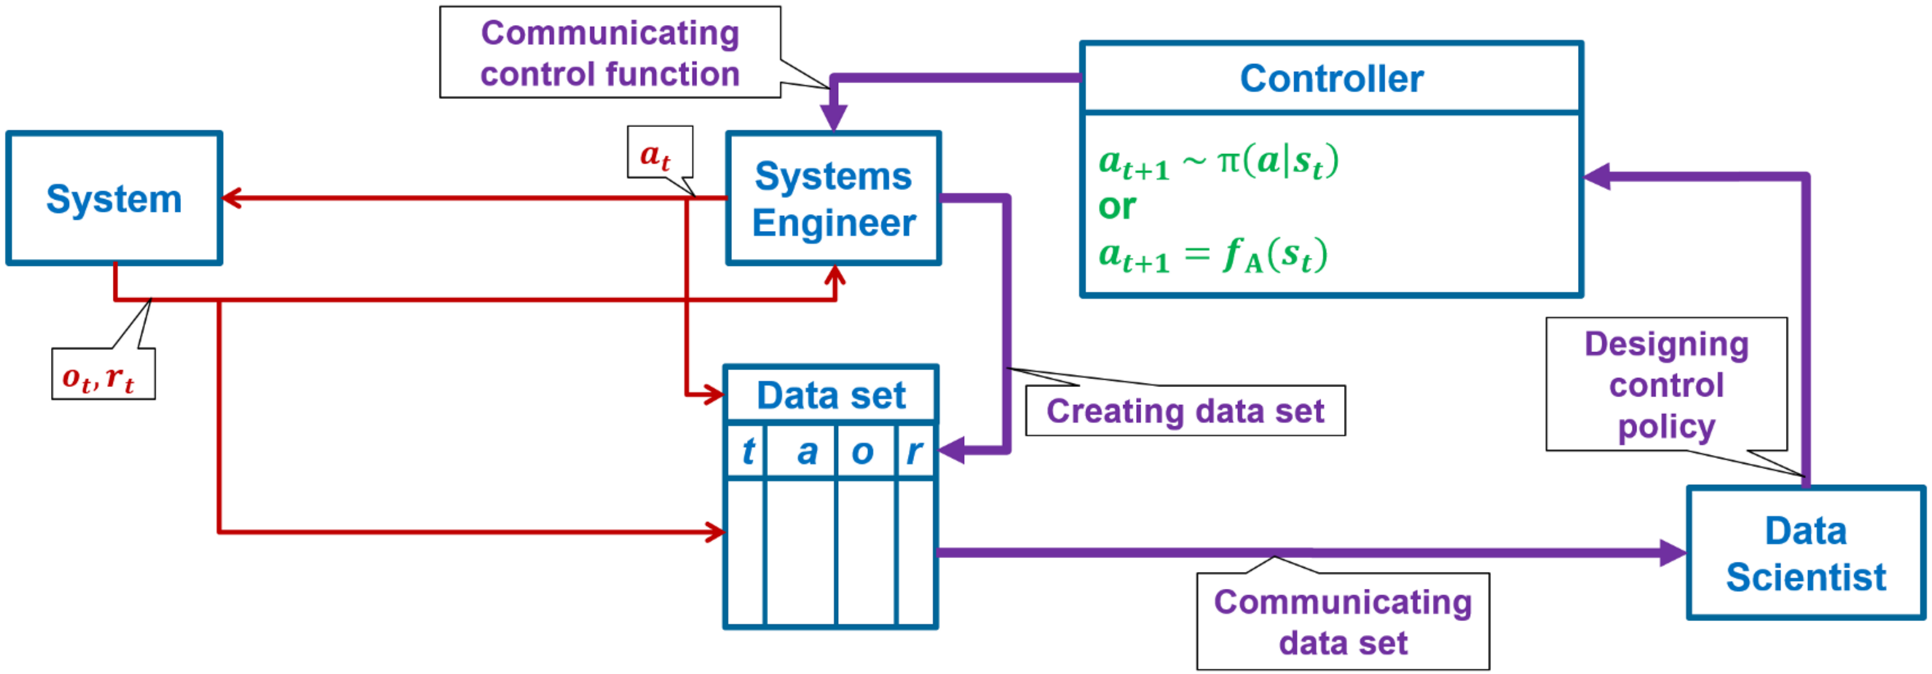

- **Realistic:**
  - **Fits the organizational scenario** we can hope to implement
  - **Technically doable**
  - **Not well-studied** in research (cf trillion dollar market)

Model-based iterated batch RL

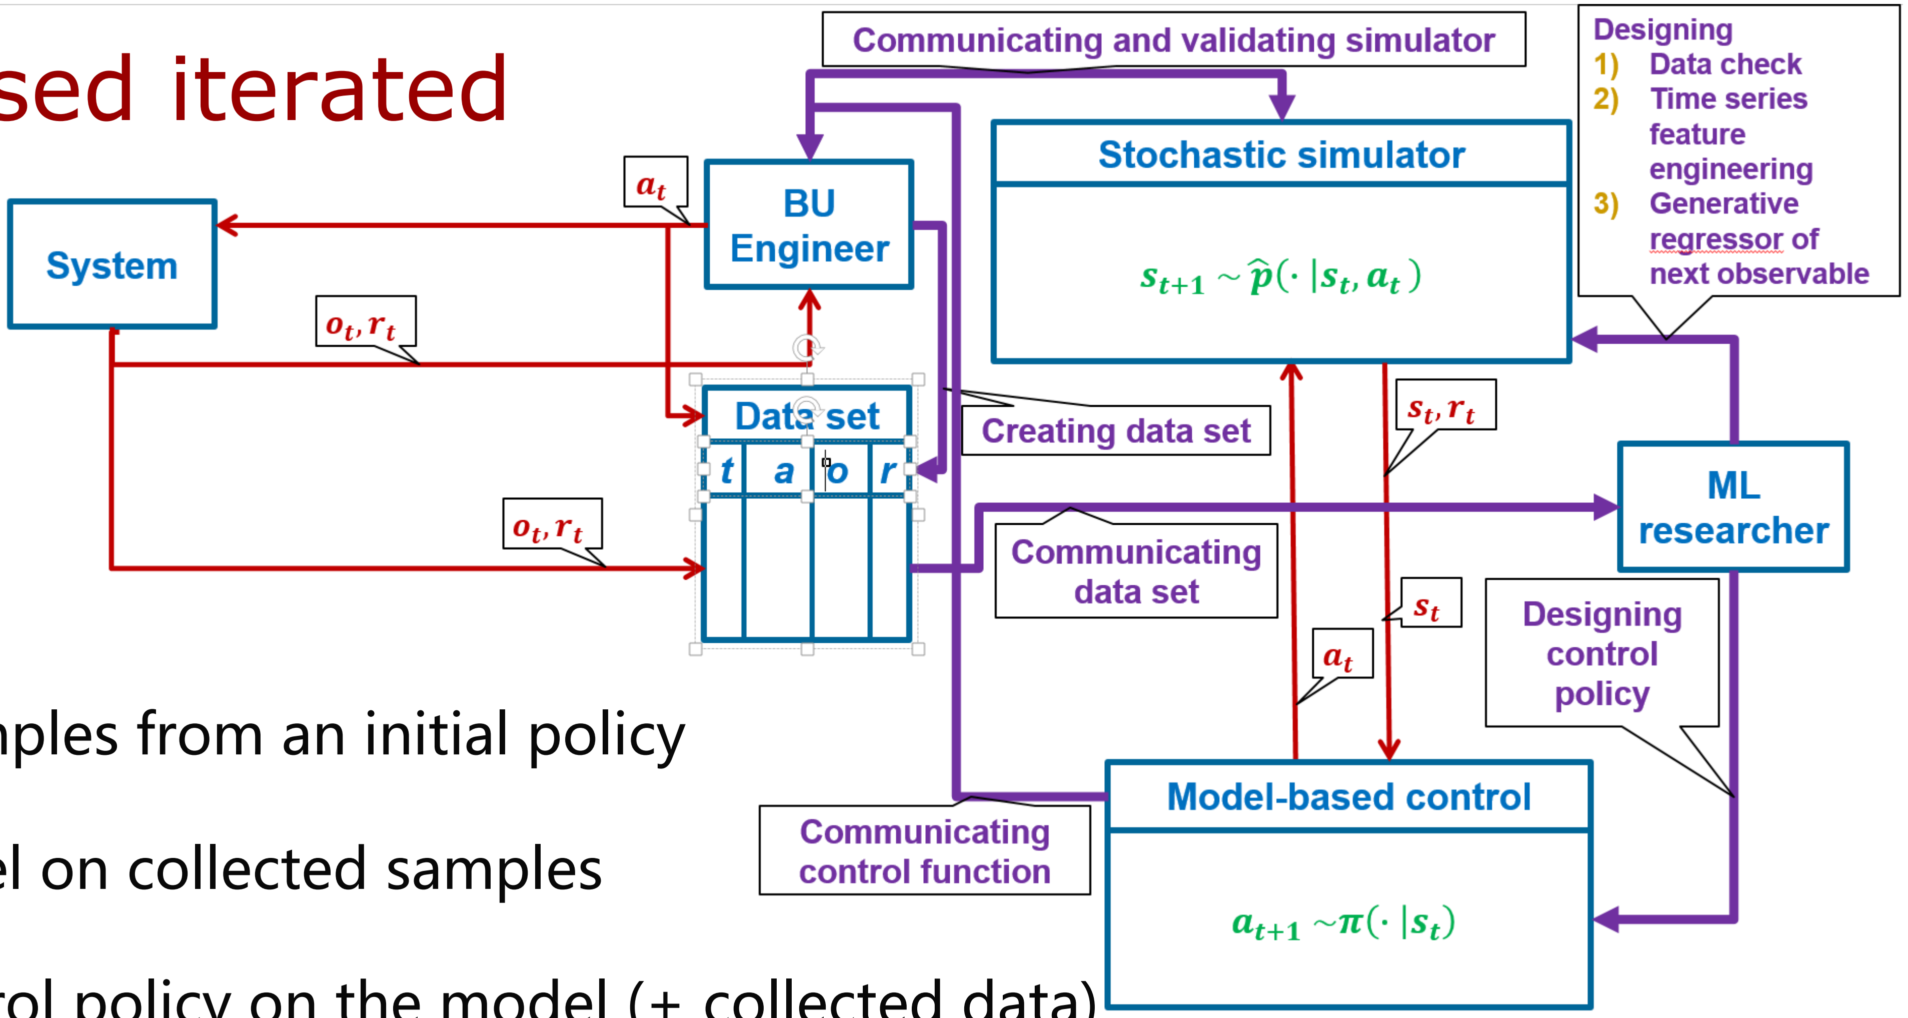

1. **Collect** samples from an initial policy
2. **Train** model on collected samples
3. **Learn** control policy on the model (+ collected data)
4. **Apply** control policy on real system and collect the data, go back to 2.

System model

=

multi-output

probabilistic (generative)

time series forecaster

System model = multi-output time series forecaster

- **Predict (random) future from history of system observables and control actions:**

$$\boldsymbol{o}_{t+1} \sim \boldsymbol{p}\left(\overbrace{\underbrace{\boldsymbol{o}_{t+1}}}^{\boldsymbol{y}} \middle| \overbrace{(\boldsymbol{o}_1, a_1), \dots (\boldsymbol{o}_t, a_t)}^{\boldsymbol{x}}\right)$$

  - We want to **simulate multiple futures** from the model
  - **Density Nets**: output the parameters of a distribution by a neural net
    - Typically **μ and σ of a Gaussian**
  - Learned on single-step transitions

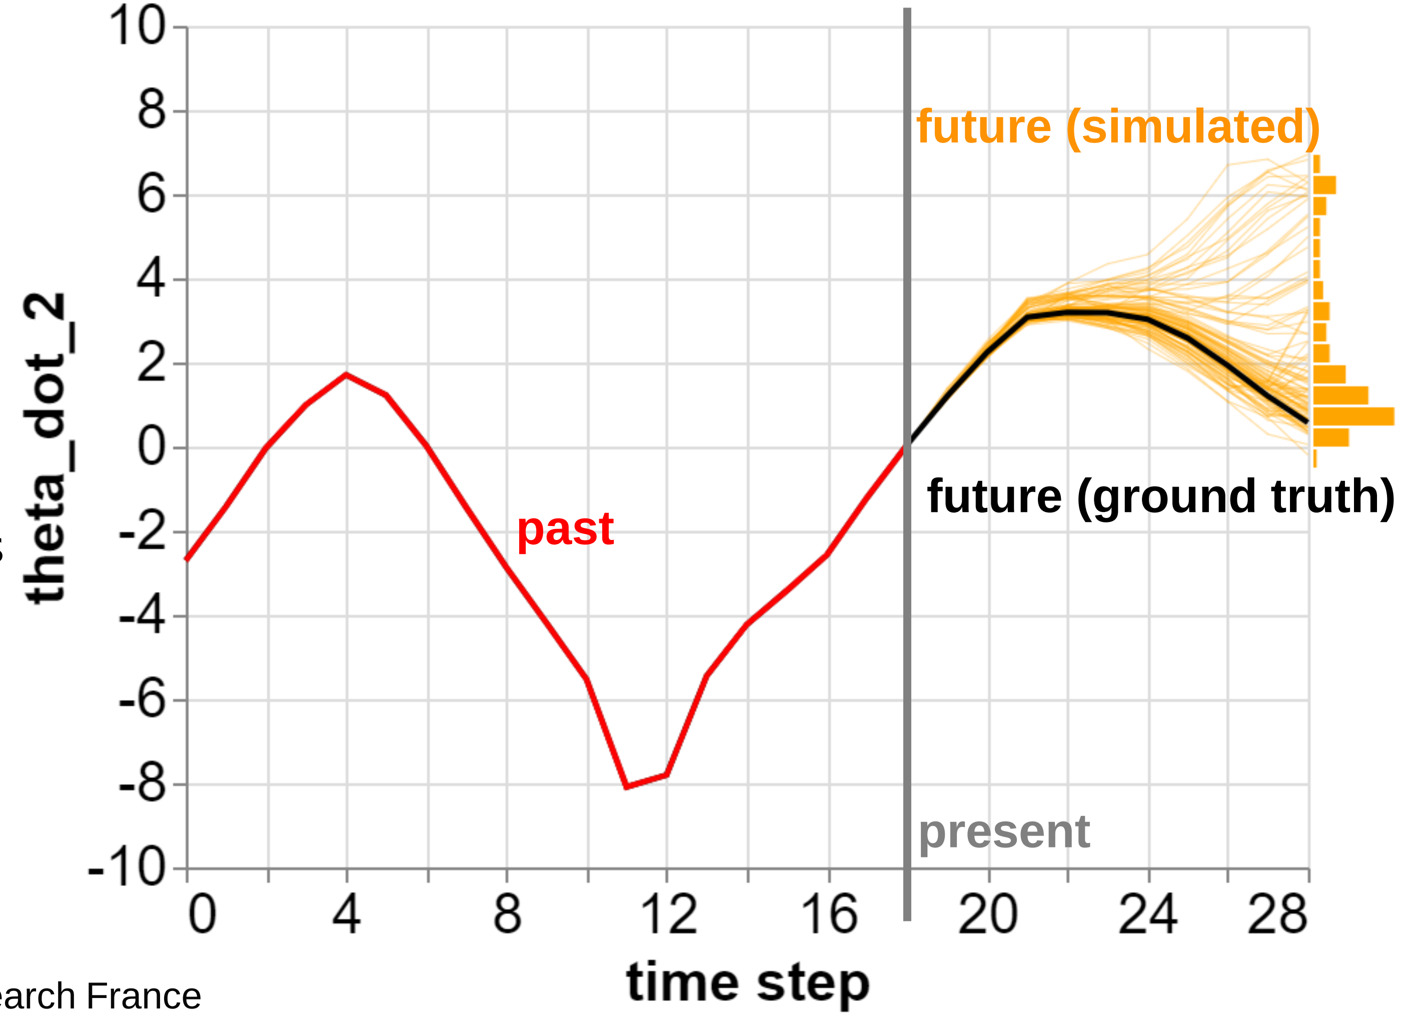

What to do with a good model (simulator)?

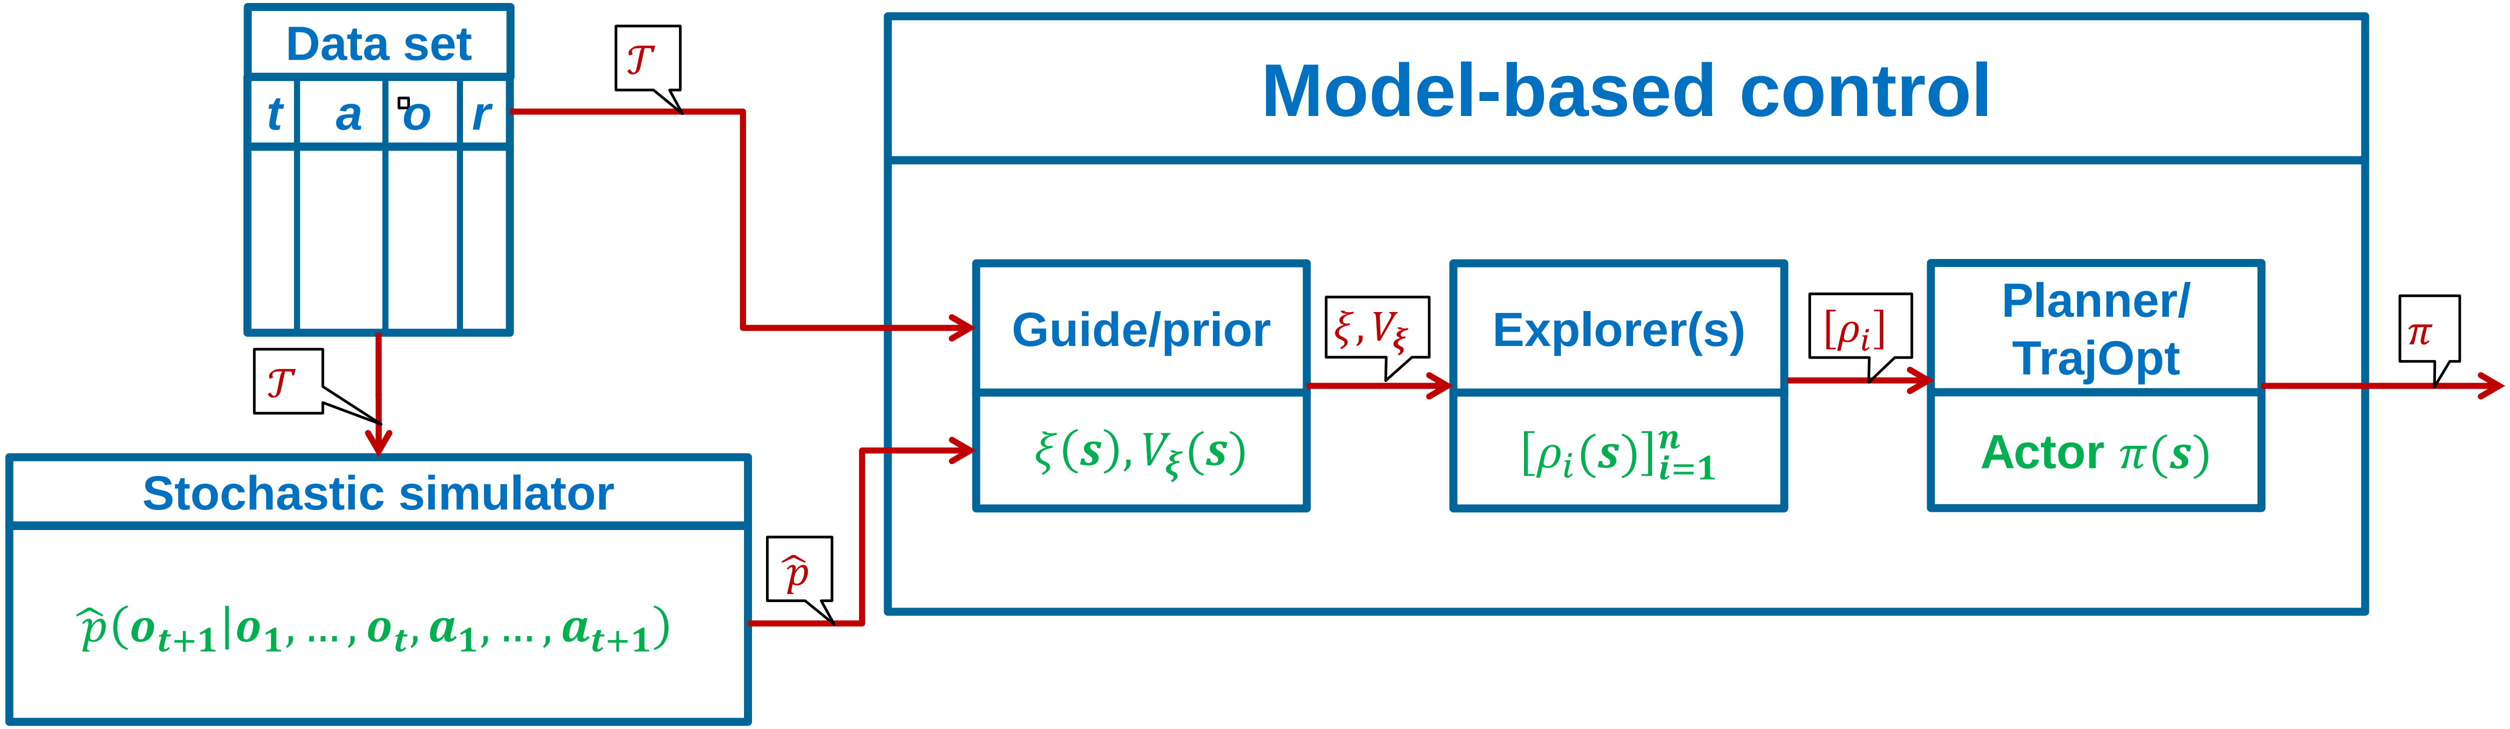

- **Learn a model-free policy** (pure Dyna-style)
- **Add exploration (*iterated* batch)**
- **Use it to guide planning**

Under review as a conference paper at ICLR 2023

THE GUIDE AND THE EXPLORER: SMART AGENTS FOR RESOURCE-LIMITED ITERATED BATCH REINFORCEMENT LEARNING

Anonymous authors
Paper under double-blind review

Why model-based batch RL?

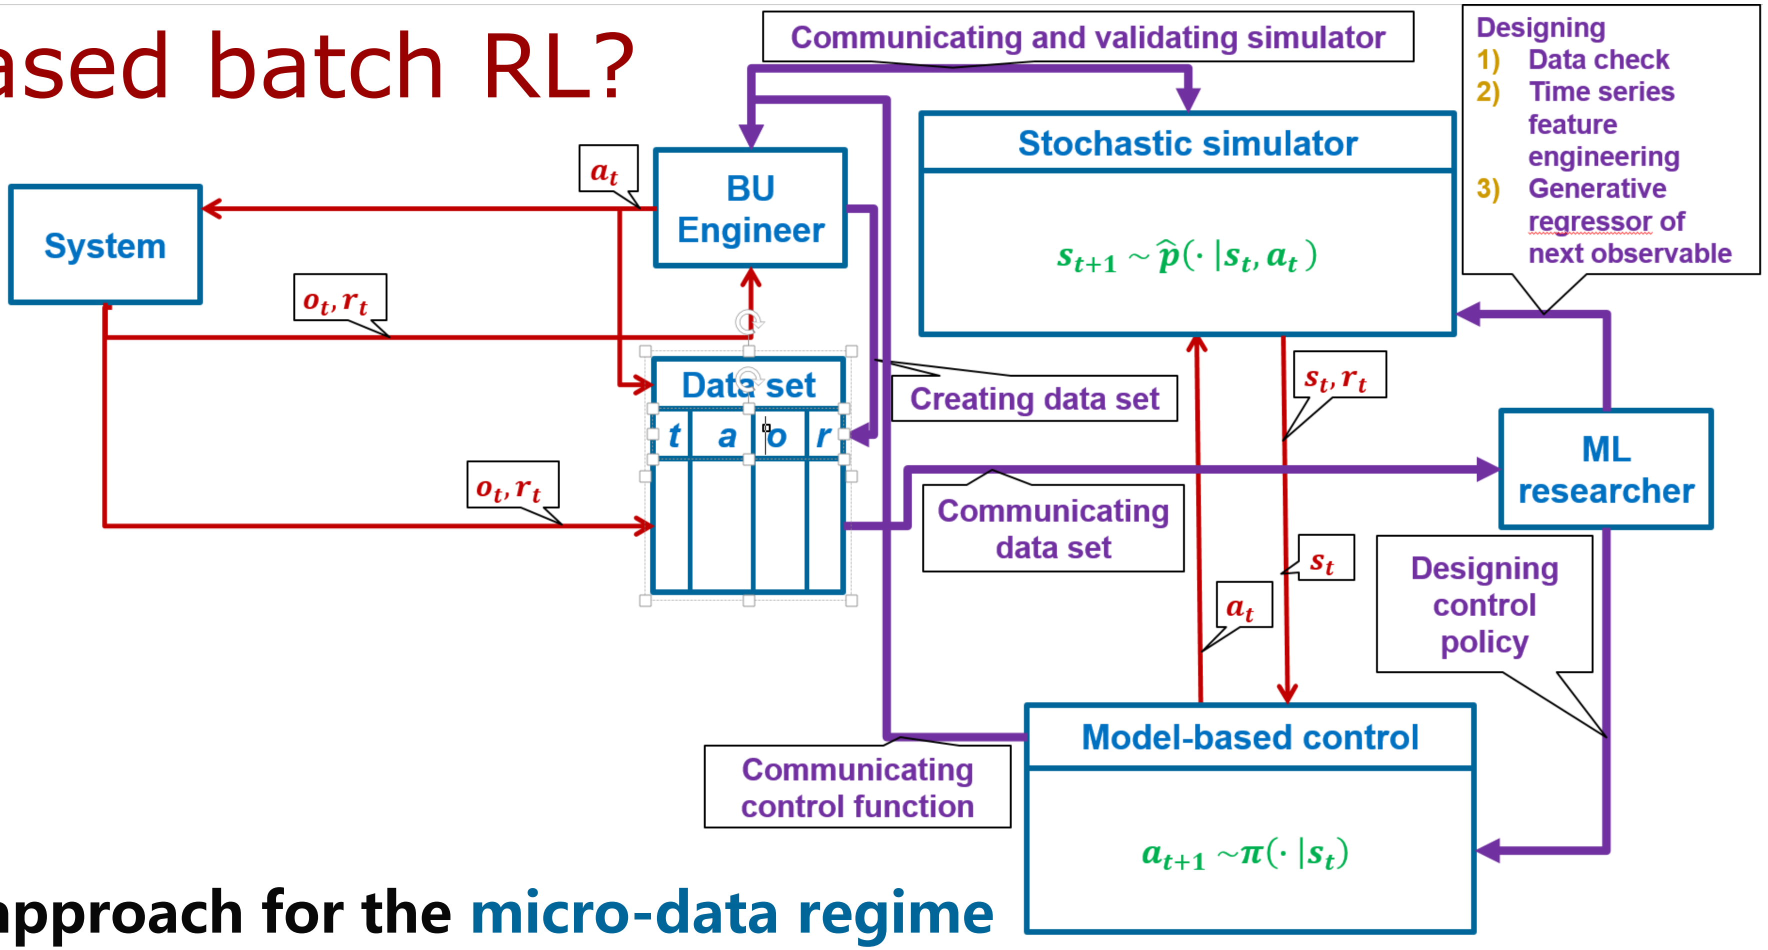

- **Considered the best approach for the micro-data regime**

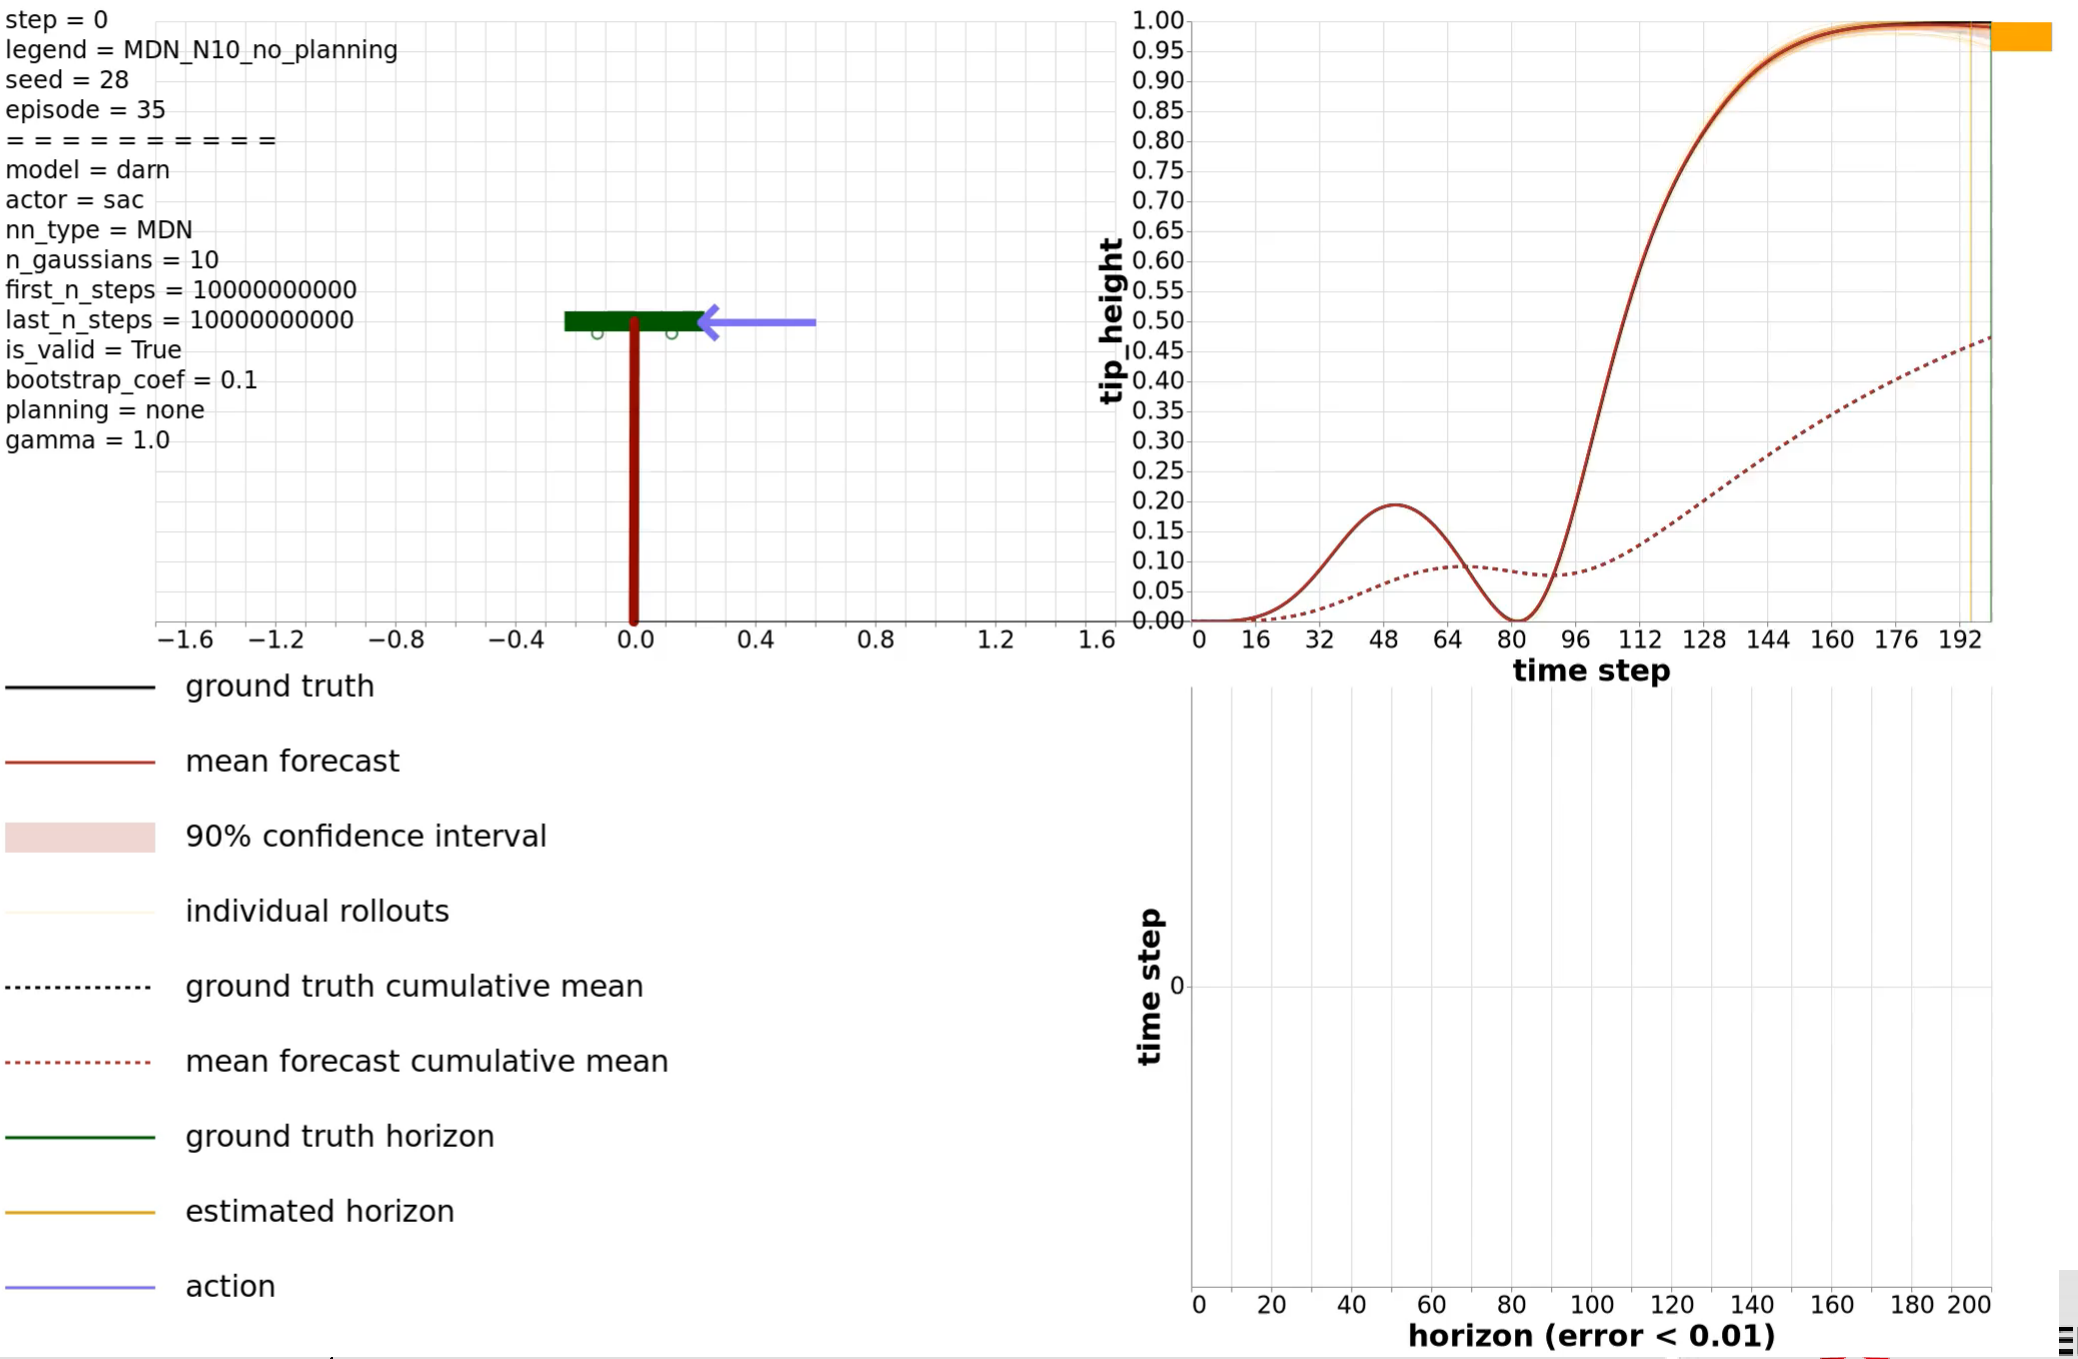
step = 0
legend = MDN_N10_no_planning
seed = 28
episode = 35
= = = = = = = = = =
model = darn
actor = sac
nn_type = MDN
n_gaussians = 10
first_n_steps = 10000000000
last_n_steps = 10000000000
is_valid = True
bootstrap_coef = 0.1
planning = none
gamma = 1.0
−1.6 −1.2 −0.8 −0.4 0.0 0.4 0.8 1.2 1.6
tip_height
1.00 0.95 0.90 0.85 0.80 0.75 0.70 0.65 0.60 0.55 0.50 0.45 0.40 0.35 0.30 0.25 0.20 0.15 0.10 0.05 0.00
0 16 32 48 64 80 96 112 128 144 160 176 192
time step
ground truth
mean forecast
90% confidence interval
individual rollouts
ground truth cumulative mean
mean forecast cumulative mean
ground truth horizon
estimated horizon
action
time step
0
0 20 40 60 80 100 120 140 160 180 200
horizon (error < 0.01)

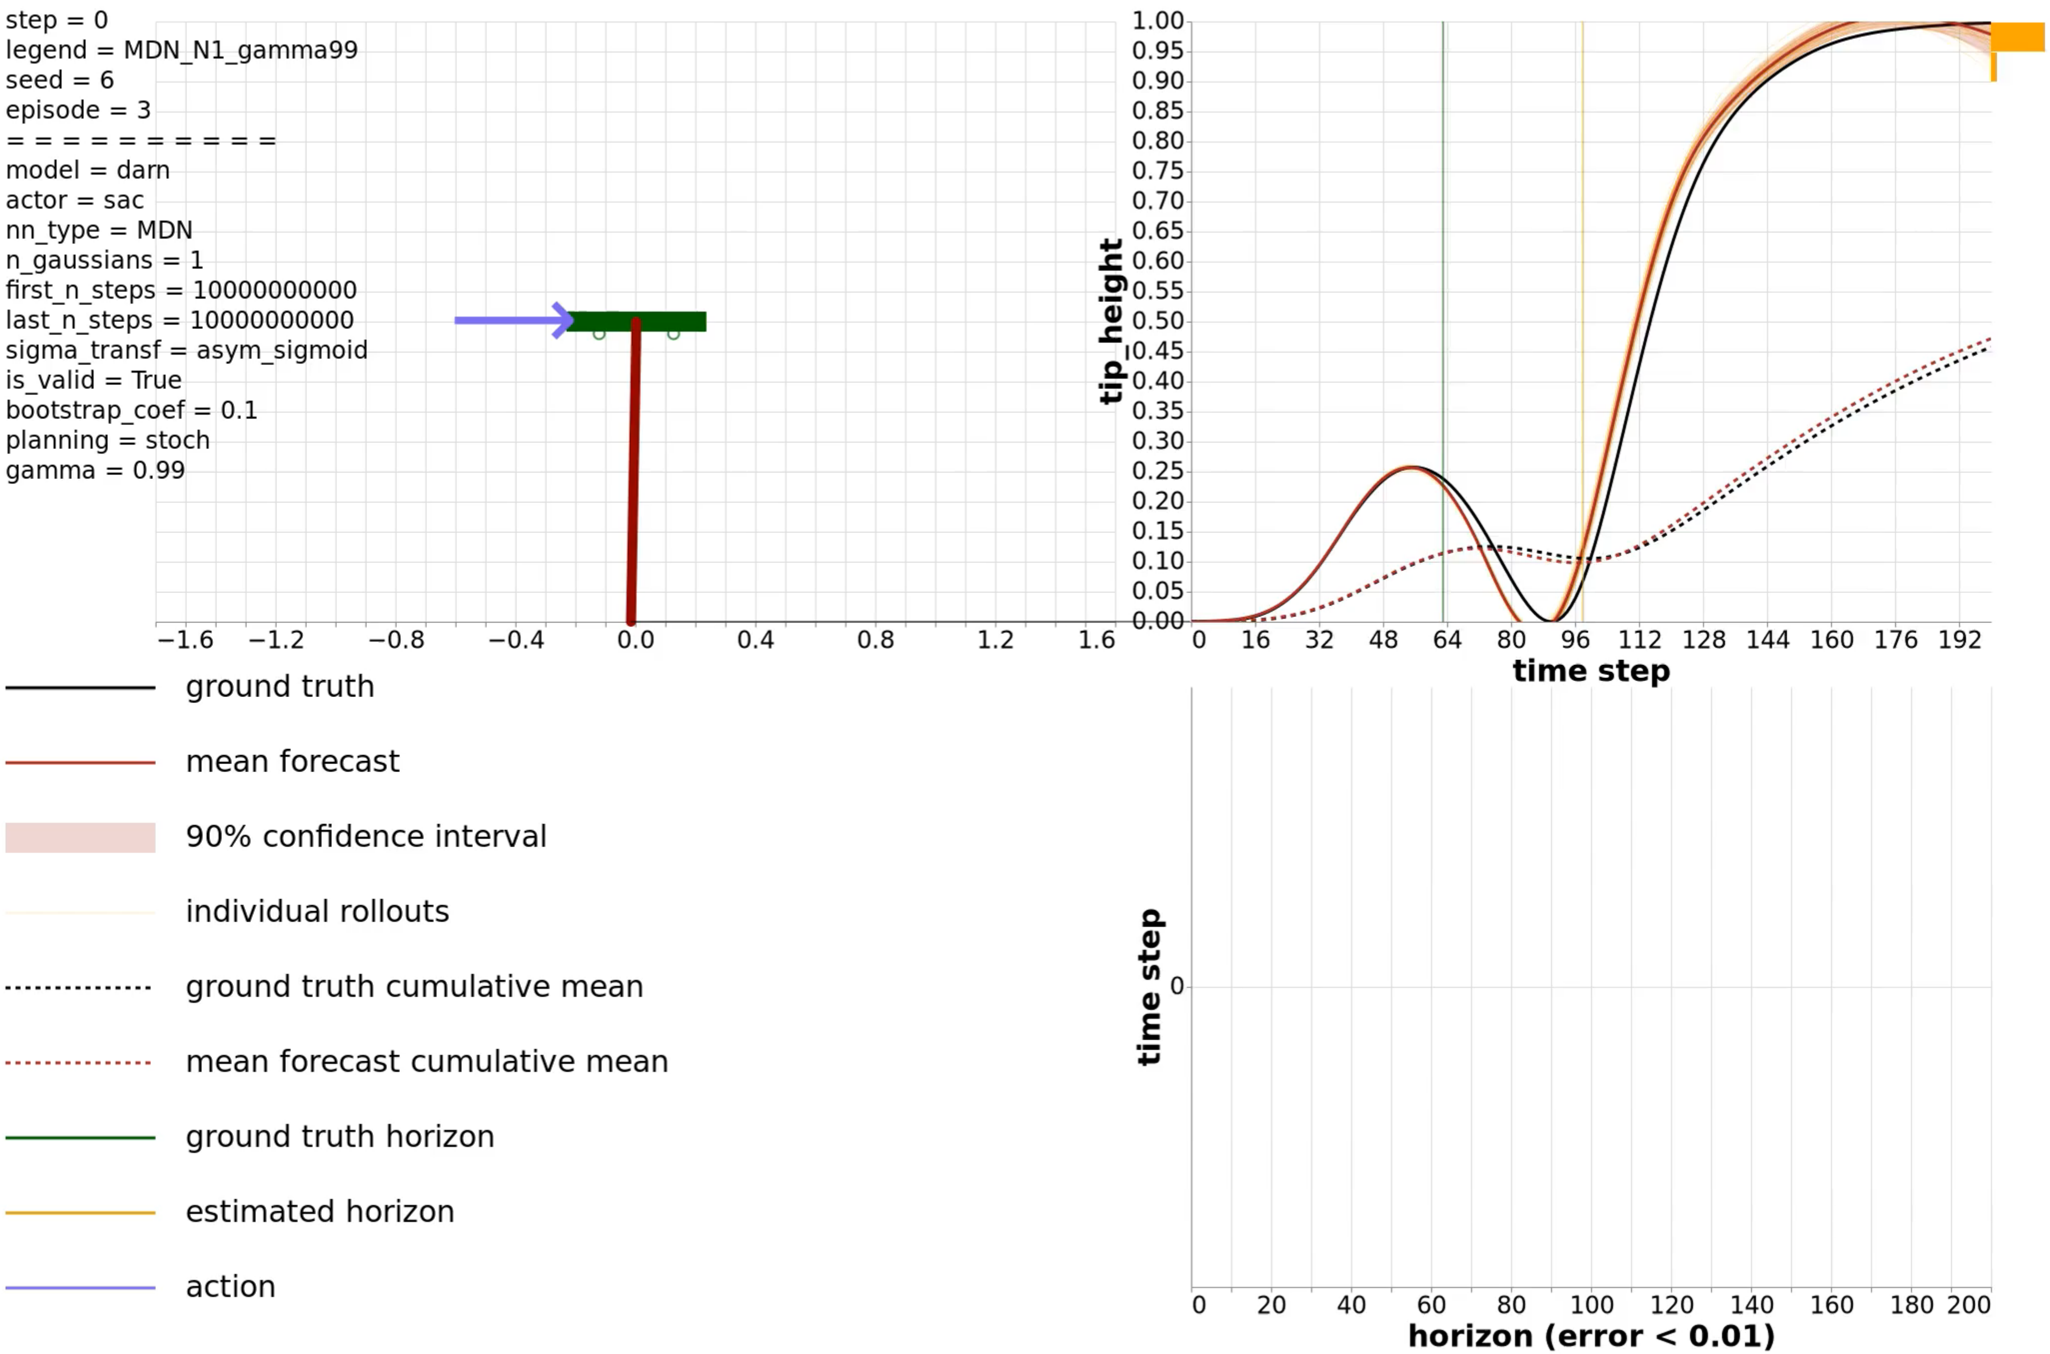
step = 0
legend = MDN_N1_gamma99
seed = 6
episode = 3
= = = = = = = = = = =
model = darn
actor = sac
nn_type = MDN
n_gaussians = 1
first_n_steps = 10000000000
last_n_steps = 10000000000
sigma_transf = asym_sigmoid
is_valid = True
bootstrap_coef = 0.1
planning = stoch
gamma = 0.99
−1.6 −1.2 −0.8 −0.4 0.0 0.4 0.8 1.2 1.6
tip_height
1.00 0.95 0.90 0.85 0.80 0.75 0.70 0.65 0.60 0.55 0.50 0.45 0.40 0.35 0.30 0.25 0.20 0.15 0.10 0.05 0.00
0 16 32 48 64 80 96 112 128 144 160 176 192
time step
ground truth
mean forecast
90% confidence interval
individual rollouts
ground truth cumulative mean
mean forecast cumulative mean
ground truth horizon
estimated horizon
action
time step
0
0 20 40 60 80 100 120 140 160 180 200
horizon (error < 0.01)

Cartpole

Model-free (SAC)

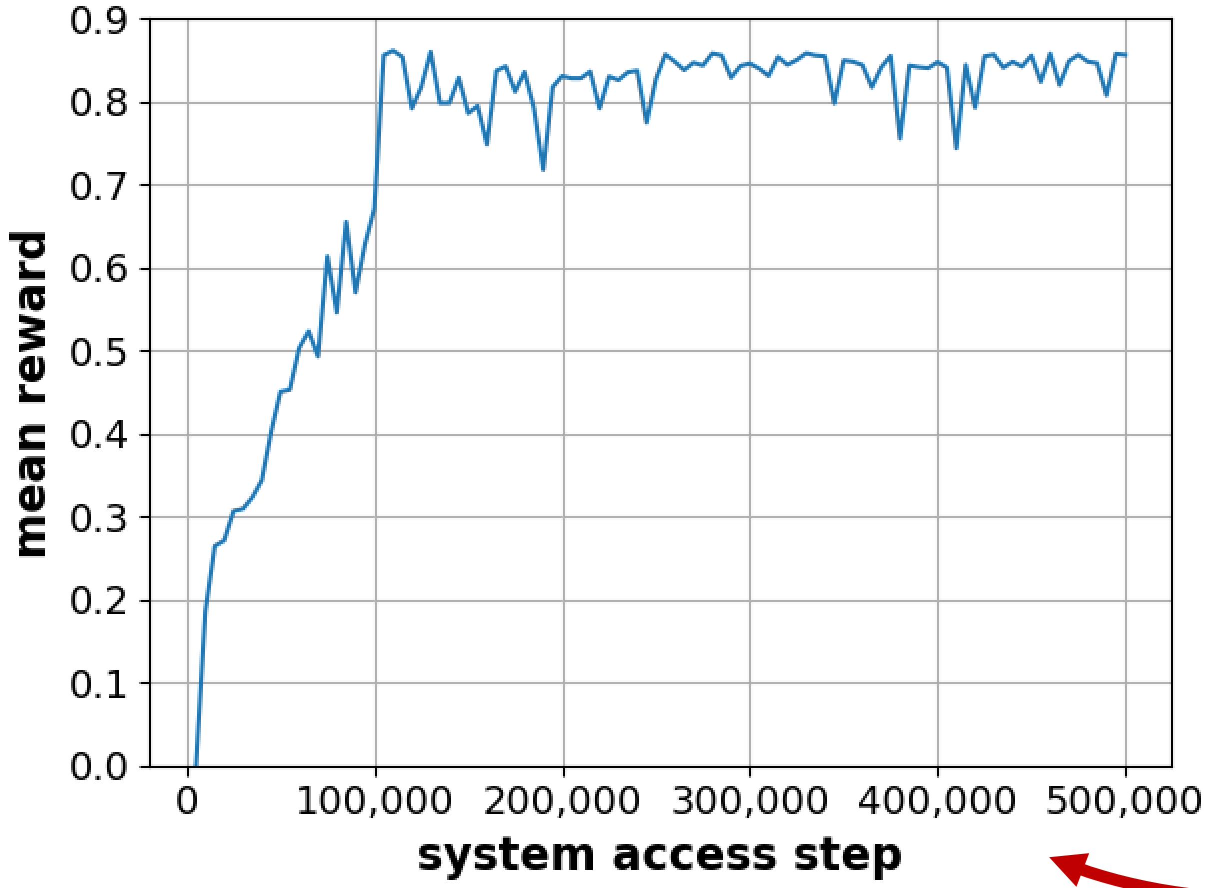

Model-based (Dyna + planning)

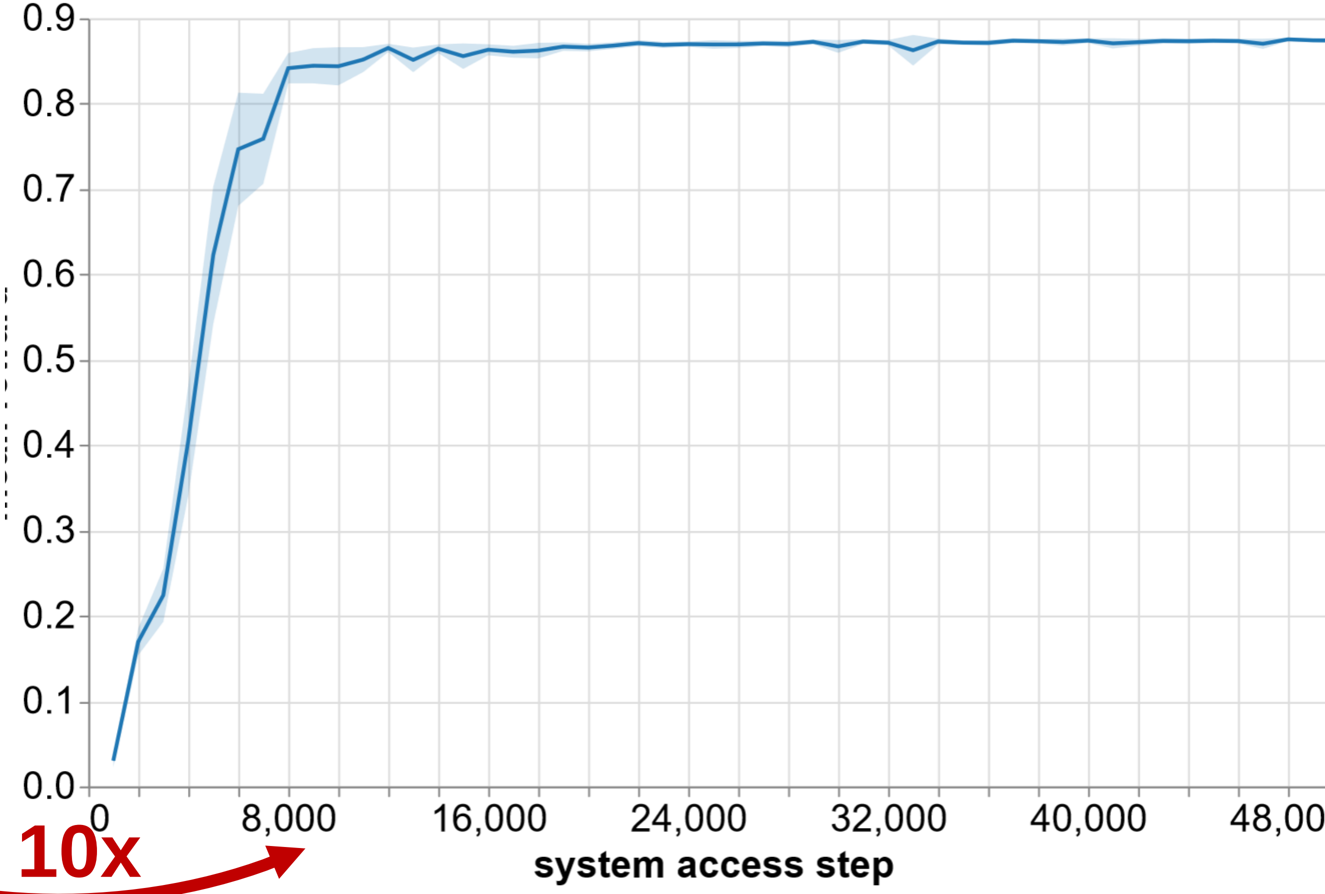

10x

Acrobot

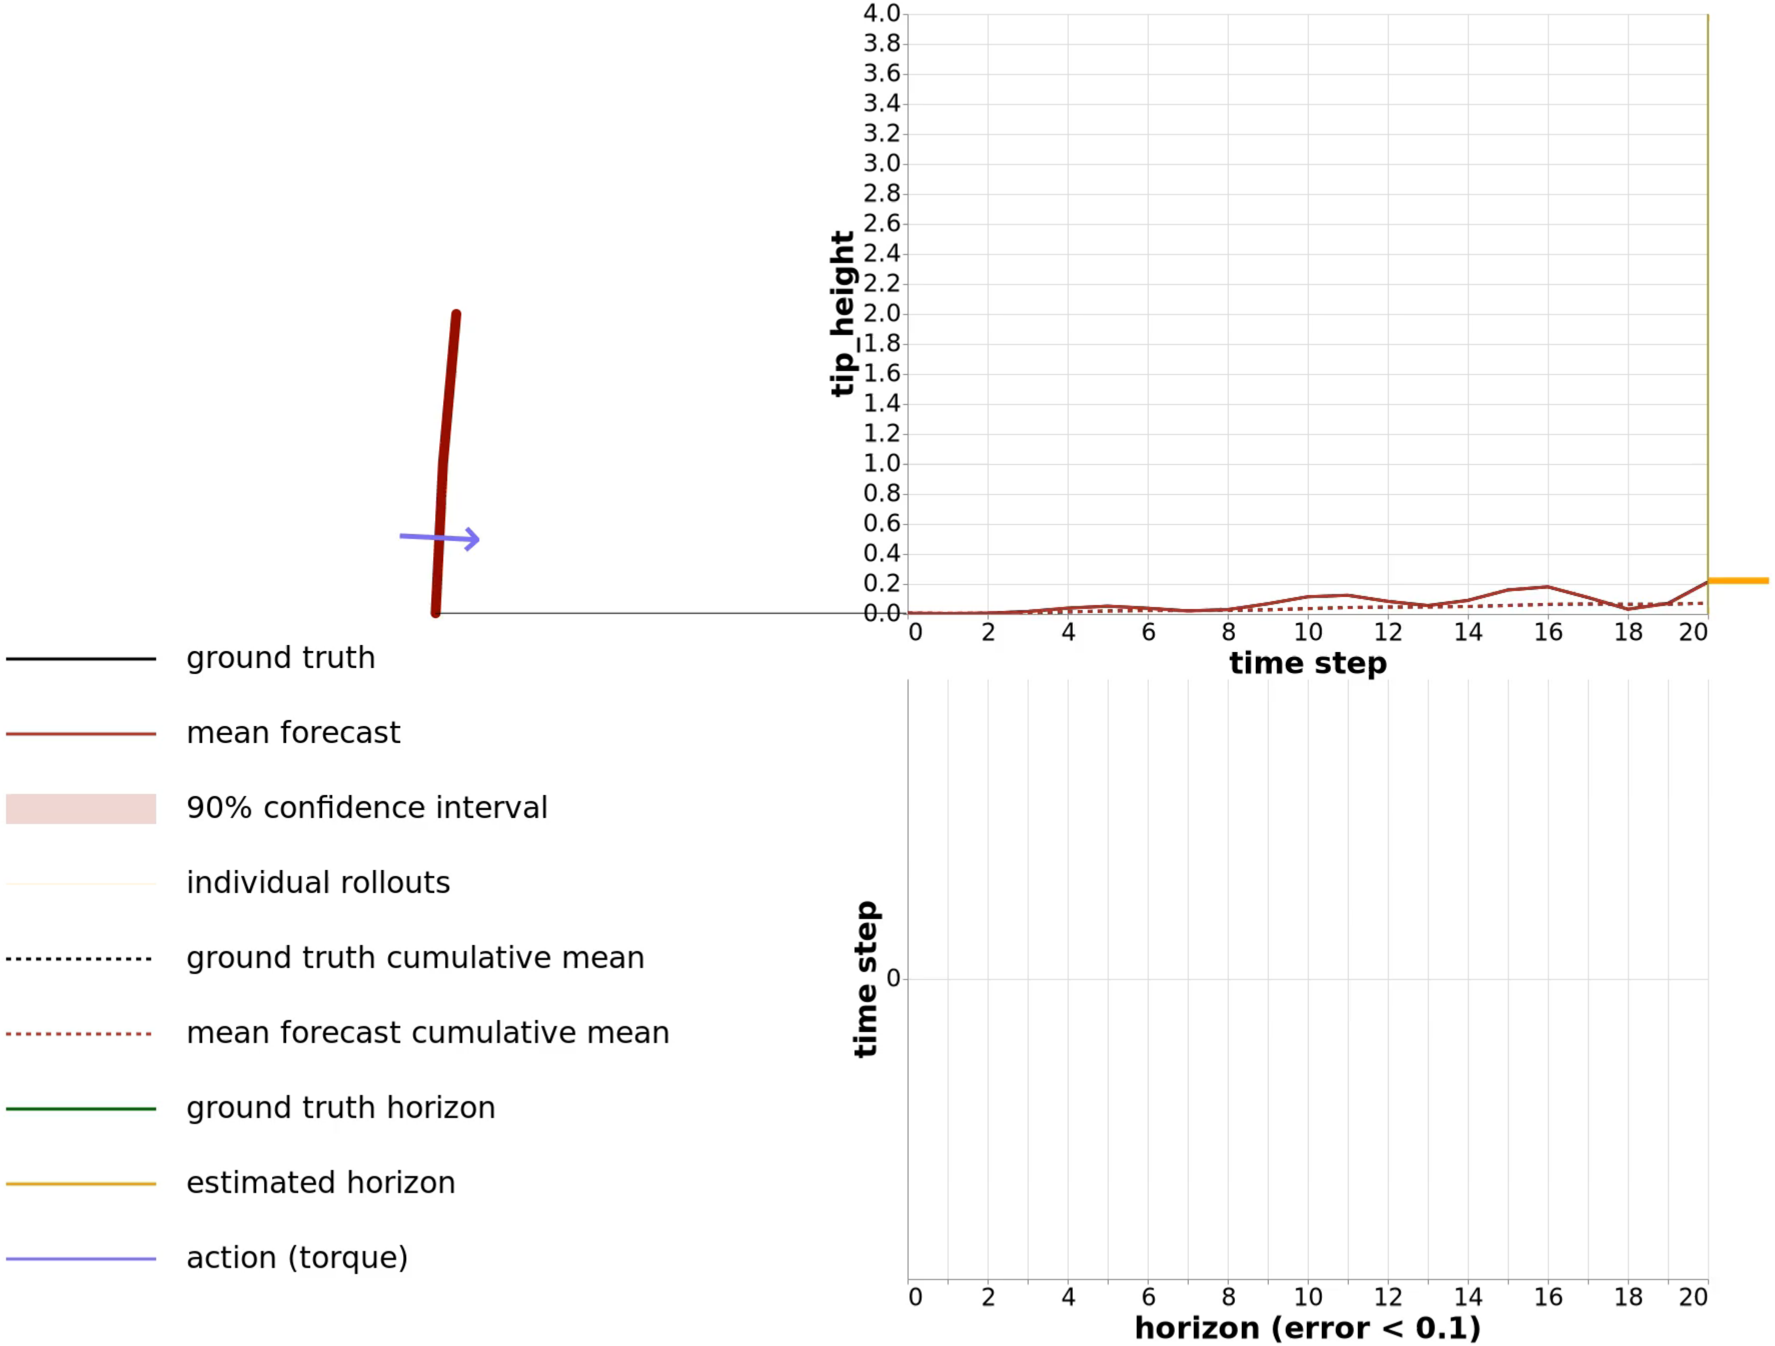
tip_height
time step
time step
horizon (error < 0.1)
ground truth
mean forecast
90% confidence interval
individual rollouts
ground truth cumulative mean
mean forecast cumulative mean
ground truth horizon
estimated horizon
action (torque)

Acrobot

Model-free (DQN)

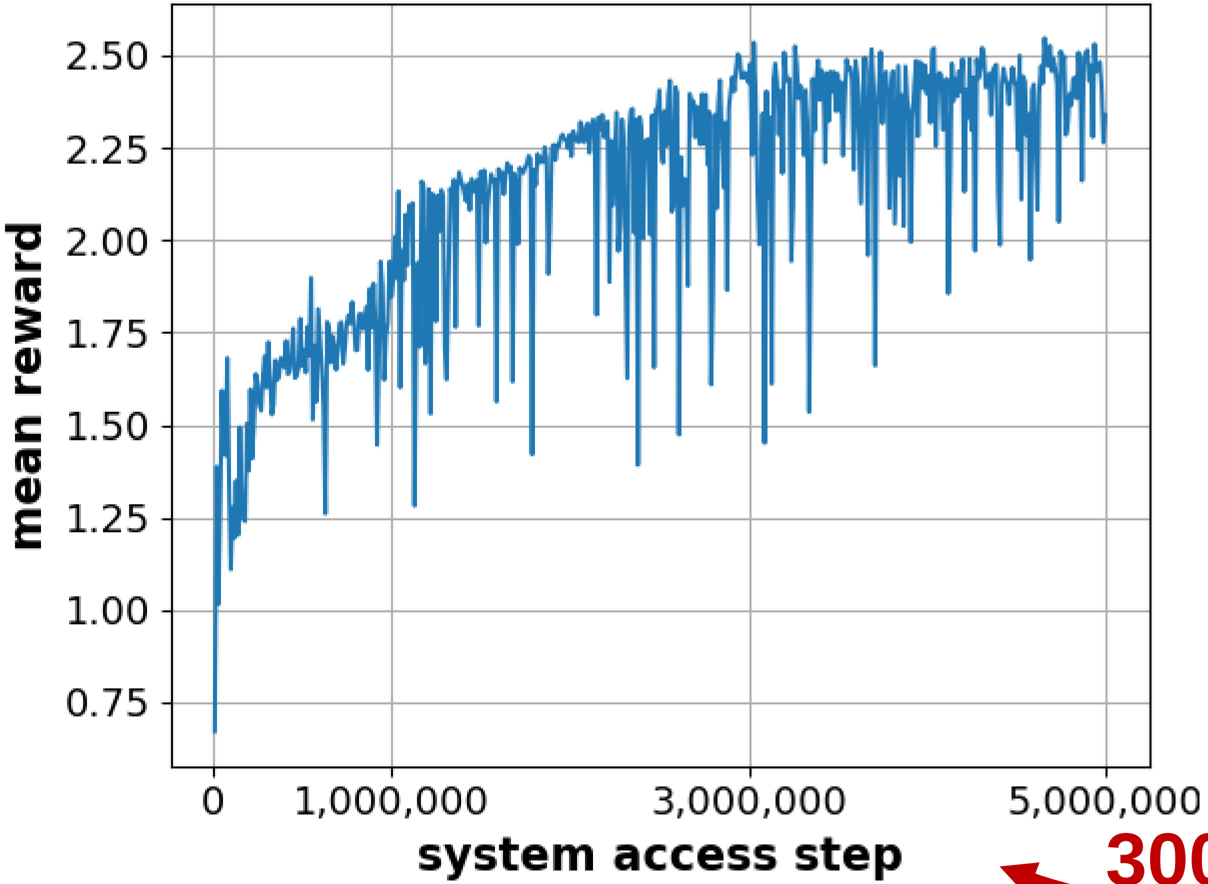

Model-based (Dyna + planning)

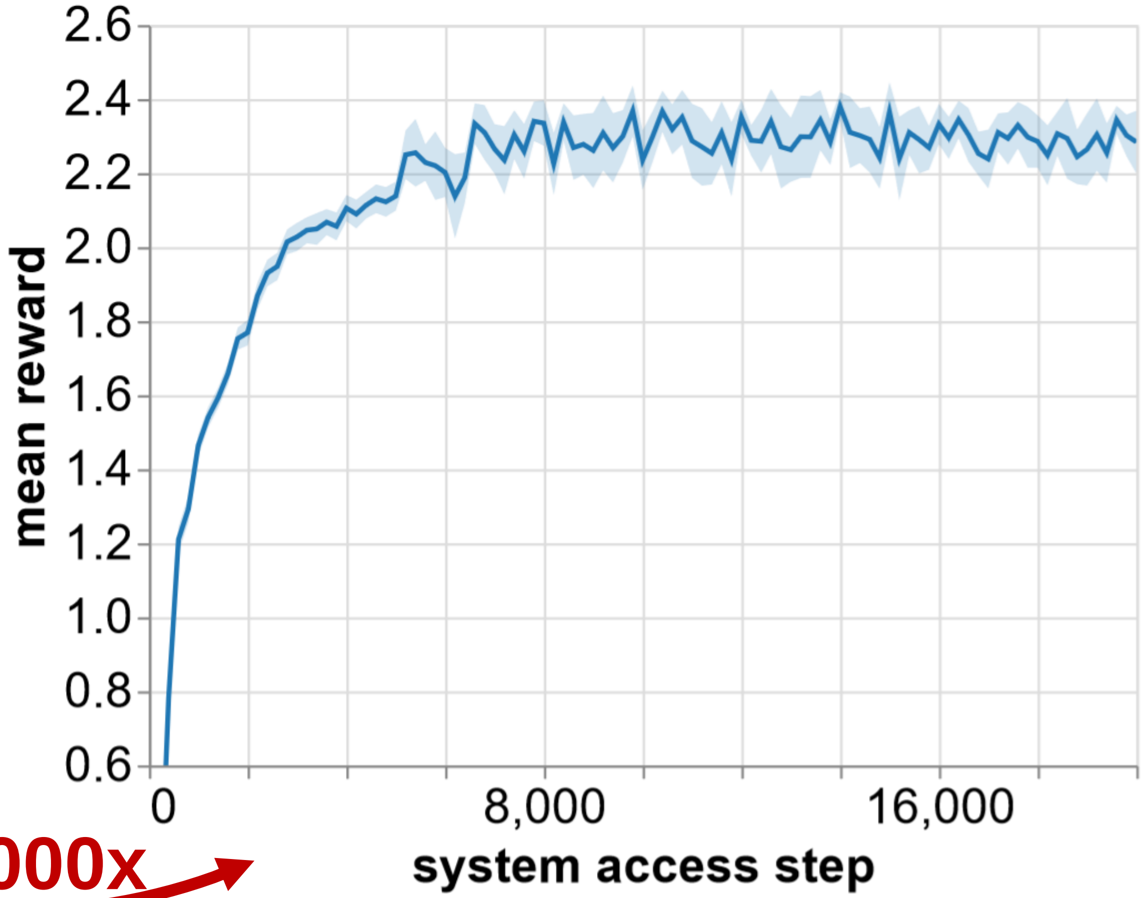

30000x

Why model-based batch RL?

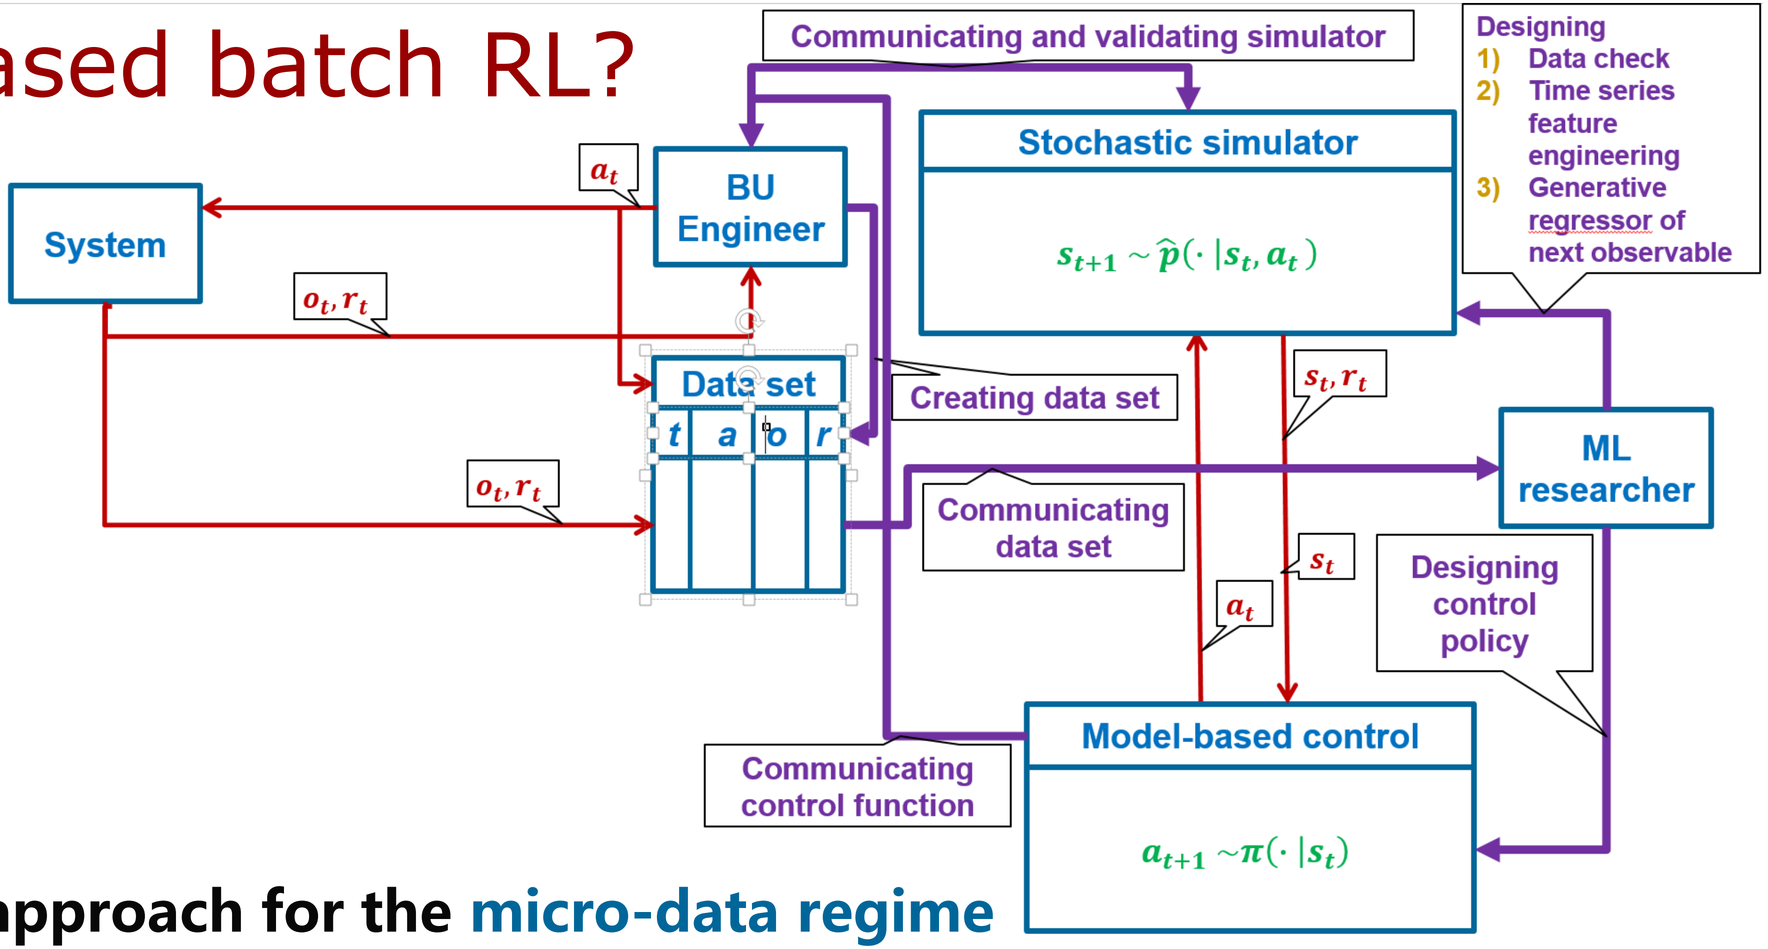

- Considered the best approach for the micro-data regime
- We do not waste predictive power (unlike, e.g., on images)
- System models (simulators or digital twins) are useful on their own
- Self-supervision (learning without human labeling) in RL

Part III

Mind Everywhere

Review of throughput for video streaming

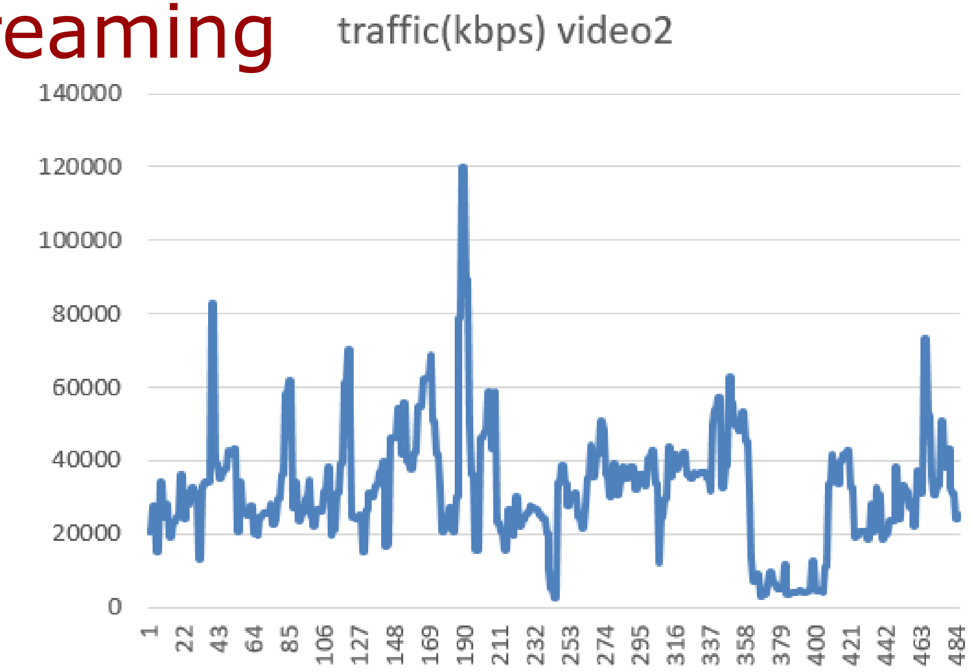

- Unstable throughput in the experiments
- From a Netflix blog post, similar behavior

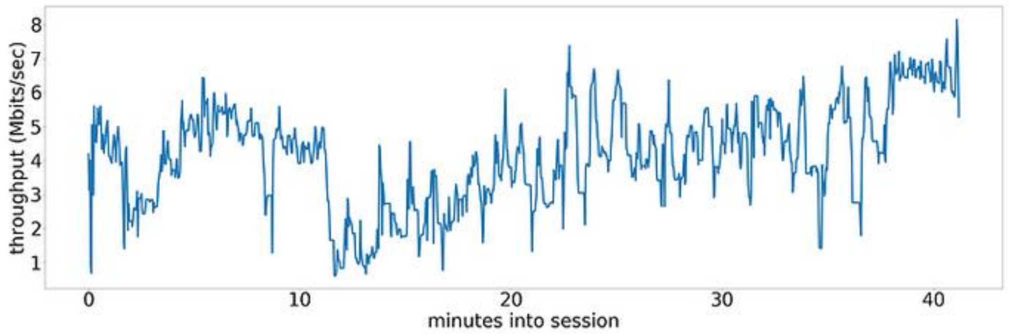

Examples of network throughput traces measured from real viewing sessions.

- Typical behavior (Netflix, Youtube): initial buffering then on-off cycles triggered by thresholds on the buffer size

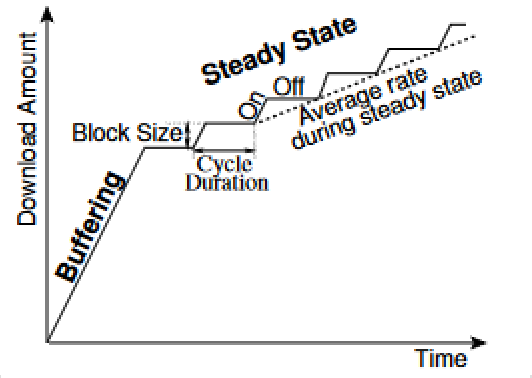

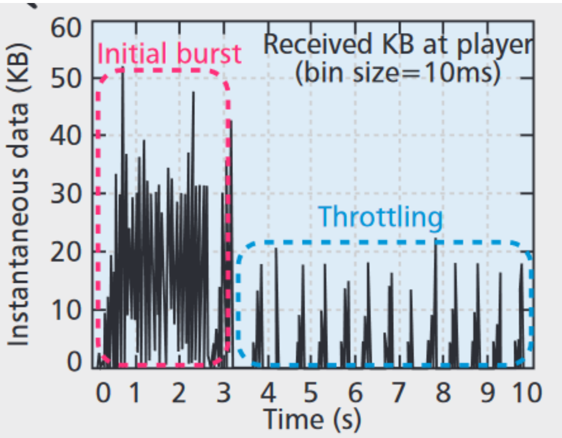

Mind everywhere

- **Phones are agential, not passive machines**
  - Intelligent buffering: a (future) video app uses AI to predict the video, the behavior of its owner and the available throughput from the base station
  - They learn and self-update
  - They cannot be fully controlled by the base station
- **Base stations are not fully in control**
  - They can vary the response to the phone requests
  - Create a landscape of states and rewards
  - Want to avoid arms races and deadlocks, for the common good
- **The body analogy**
  - Phones are cells with their own monocellular goals
  - Base station – set of phones organism is a higher level multi-cellular organism
  - Emergence and emanation

Mind everywhere

- **Thesis**
  - Future intelligent systems will look more like **biological bodies** than classical control systems
  - So their design should be inspired by the latest research in **distributed cognition** and **multi-layer developmental system biology**
  - Understanding and creating these embodied systems will complement language-based narrow AI on the **road to AGI**
- **Inspirations**
  - Michael Levin: Technological Approach to Mind Everywhere
  - Mark Solms: The Hidden Spring, affect-based dual-aspect monistic view of consciousness
  - Lina Bariah and Merouane Debbah: AI embodiment through 6G: shaping the future of AGI
  - John Vervaeke: Artificial Intelligence, The Meaning Crisis, & The Future of Humanity (YouTube)

Mind everywhere

- **Intelligence by William James:**

Solving a higher-level problem through *whatever* lower-level mechanism is available.

- **Top-down causation**

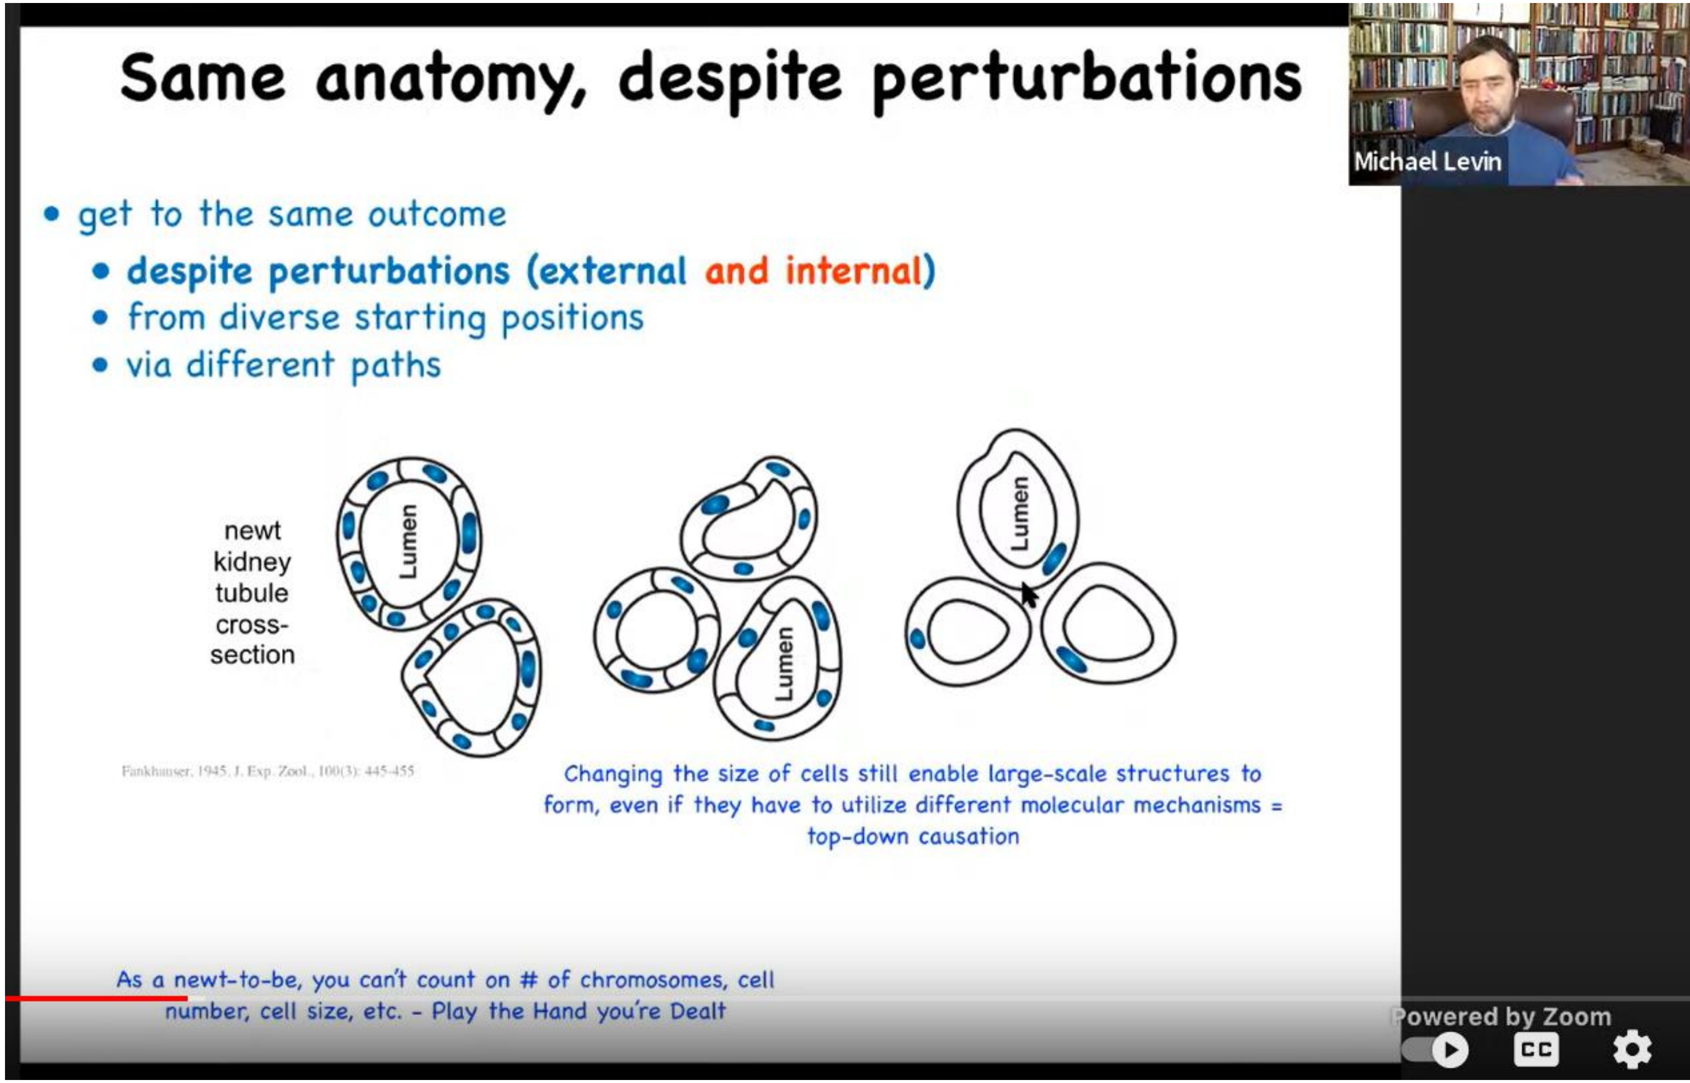

Conditions for consciousness

- **Autonomous, self-organizing, auto-poietic systems**
- **Separated from their environment by a Markov blanket: sensors for extero and interoception**
- **Responsible for their persistence: maintaining and optimizing vital internal states**
  - Modelling future states and recording predictive errors
  - Feeling, affect = feedback signal: how well am I doing?
  - Actions: what to do to improve?
- **More than one vital dimensions which require arbitration at the organism level, cannot be automated because of unknown unknowns in the environment**

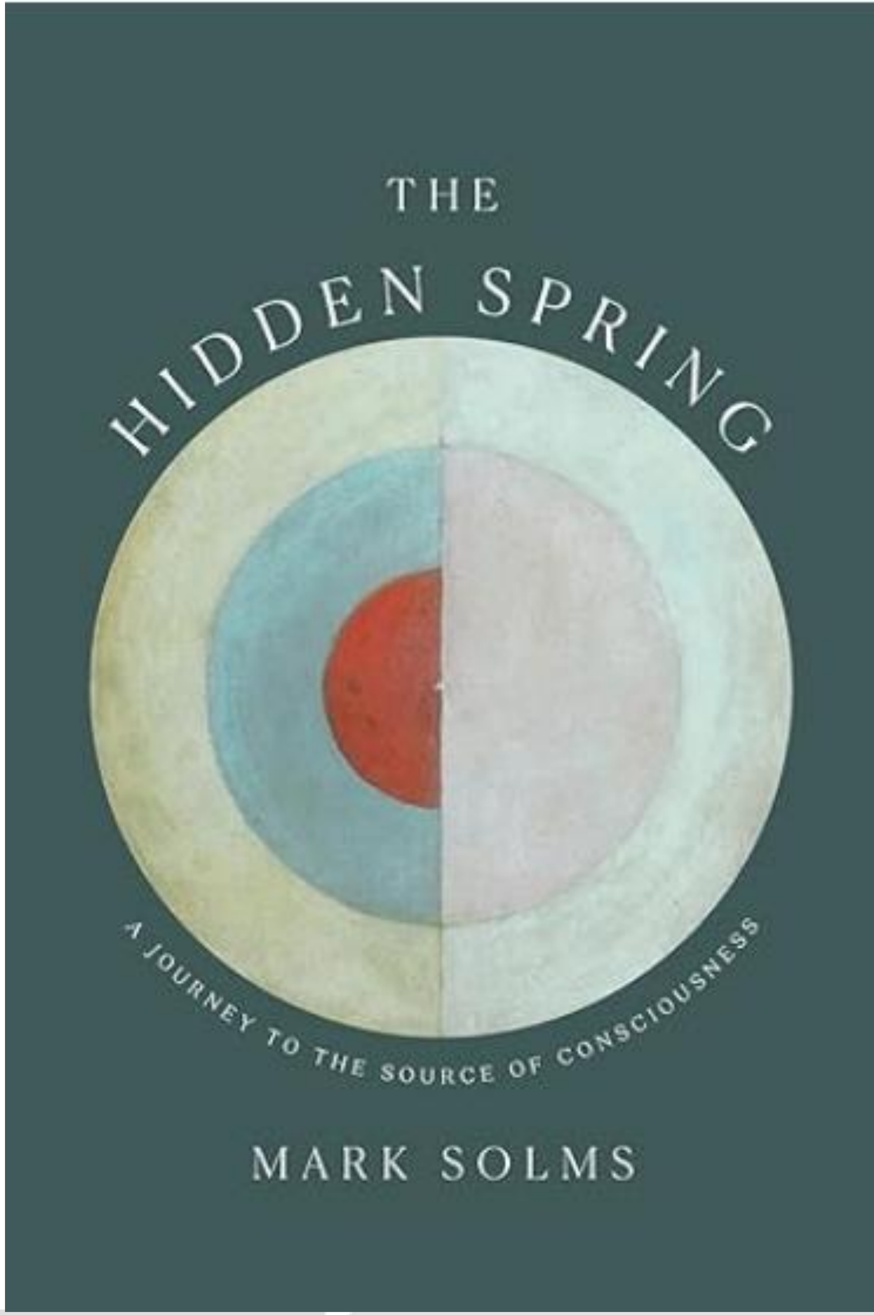

How will engineering systems look like?

- **They will consist of intelligent parts that all build predictive models**
- **They will act to satisfy and optimize multiple key performance indicators**
  - Which are often contradictory
- **They will act to safely explore, collect data, and learn about the world**
- **They will be hierarchical and multi-agent**
  - Horizontal and vertical channels
- **They will not be in full control of their parts**
  - Request, persuasion, shaping states and rewards, hierarchical adaptation
- **They will individuate**
  - Old hardware will need to teach new hardware
- **They will communicate, partially through language**
  - They will talk to us, and maybe to each other

The four pillars of embodied AIoT

- **World models and digital twins**
  - **Data efficient**
  - **Self-calibrated**: will know about their own uncertainties
  - Consistently **multi-timestep and multi-timescale**
  - Robust and **auto-tuning**
- **Model-based control**
  - **Data efficient** and **safe**
  - **Continual learning** and deployment
  - Combining model-based, model-free control and planning
- **Communication through LLMs**
  - **Episodic memory**: each device will remember its experience
- **Hierarchical multi-agent systems**
  - **Collective intelligence**, persuasion, reward and state-shaping, emergence and emanation

https://balazskegl.medium.com/mind-everywhere-embodied-ai-of-things-and-the-future-of-engineering-124e1b3a35e7

Twitter: https://twitter.com/balazs.kegl
Blog: https://balazskegl.substack.com
Podcast: https://www.youtube.com/@I.scientist.balazskegl