CS343
Artificial Intelligence

Profs: Peter Stone and Yuke Zhu

Department of Computer Science
The University of Texas at Austin

Good Morning, Colleagues

Final Exam

- Wednesday May 12th, 2pm-5pm (or some 3-hour window close to then)

Final Exam

- Wednesday May 12th, 2pm-5pm (or some 3-hour window close to then)
- Same rules as midterm (open book and notes, but don't search internet for answers)

Final Exam

- Wednesday May 12th, 2pm-5pm (or some 3-hour window close to then)
- Same rules as midterm (open book and notes, but don't search internet for answers)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning

Final Exam

- Wednesday May 12th, 2pm-5pm (or some 3-hour window close to then)
- Same rules as midterm (open book and notes, but don't search internet for answers)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm

Final Exam

- Wednesday May 12th, 2pm-5pm (or some 3-hour window close to then)
- Same rules as midterm (open book and notes, but don't search internet for answers)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm
- 3 hours rather than 1.5 hours

Final Exam

- Wednesday May 12th, 2pm-5pm (or some 3-hour window close to then)
- Same rules as midterm (open book and notes, but don't search internet for answers)
- Covers the whole semester
  - Slightly heavier emphasis on material since midterm
  - Certainly a question on planning
- Striving for similar difficulty to midterm
- 3 hours rather than 1.5 hours
- Samples - Berkeley exams

Planning Recap

- Graphplan

Is AI all about the data?

- Which programming assignments used data files?

Your Questions

- Vijay Vuyyuru: Is it useless to study the brain separately from the body since both interact?

Your Questions

- Vijay Vuyyuru: Is it useless to study the brain separately from the body since both interact?

- Theodore Venter: Is there anyone taking the idea of perceptrons and deep learning to the next level by using biological components? And if not, has anyone tried designing separate neuron-like circuits to create a real artificial brain?

Questions

- Nitin Jain: Is AI automation a good thing or bad thing and how heavily should this be considered when working on further AI advancements?

Questions

- Nitin Jain: Is AI automation a good thing or bad thing and how heavily should this be considered when working on further AI advancements?
- Are there some types of research we shouldn't do?

Questions

- Nitin Jain: Is AI automation a good thing or bad thing and how heavily should this be considered when working on further AI advancements?
- Are there some types of research we shouldn't do?
- If an AI technology runs amok, who is responsible?

Questions

- Ethan Lao: What are the newest changes in AI that were discovered after the textbook?

Questions

- Ethan Lao: What are the newest changes in AI that were discovered after the textbook?
- Joseph Jiang: If humans have made strides in AI that would be considered impossible just a decade ago, is there real merit as to the definite limitations of AI?

Questions

- Ethan Lao: What are the newest changes in AI that were discovered after the textbook?
- Joseph Jiang: If humans have made strides in AI that would be considered impossible just a decade ago, is there real merit as to the definite limitations of AI?
- Can computers perfectly simulate a human's decision-making (weak AI)?
- When will AI reach human-level intelligence?

Some Current AI Research

Building General-Purpose Robot Autonomy

Yuke Zhu

May 4, 2021

UT-Austin Robot Perception and Learning Lab

Traditional Form of Robot Automation

Structured Environments

Fixed Set of Tasks

Pre-programmed Procedures

General-Purpose Robot Autonomy: Our North Star Goal

Natural Environments

Ever-Changing Tasks

Human Involvement

Robot Autonomy as a “Pixels-to-Torques” Problem

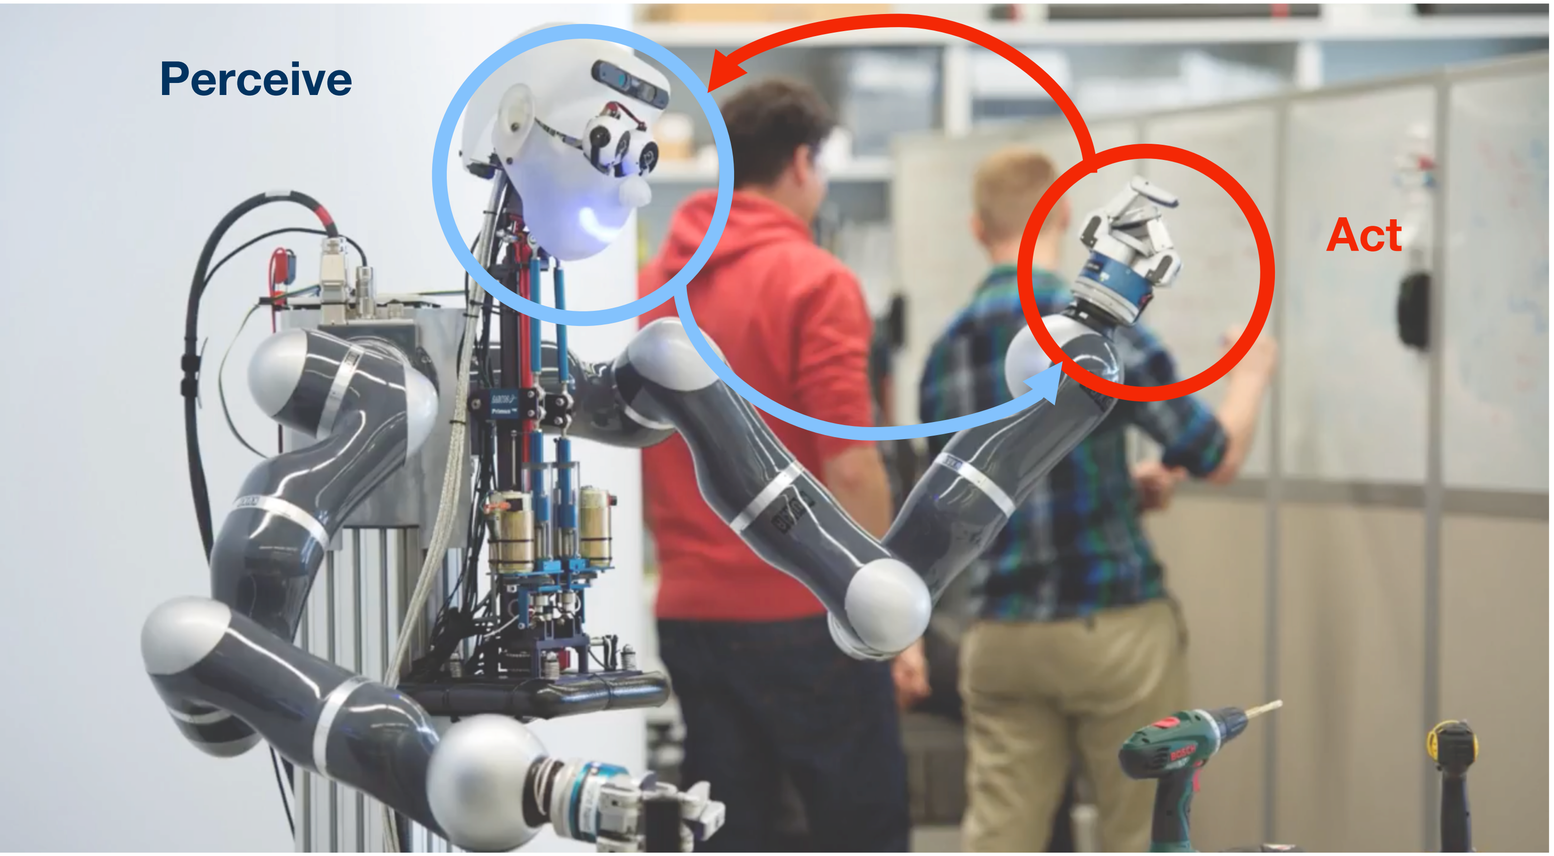

Policy $\pi : \mathcal{O} \rightarrow \mathcal{A}$

pixel

torque

Robot Autonomy as a “Pixels-to-Torques” Problem

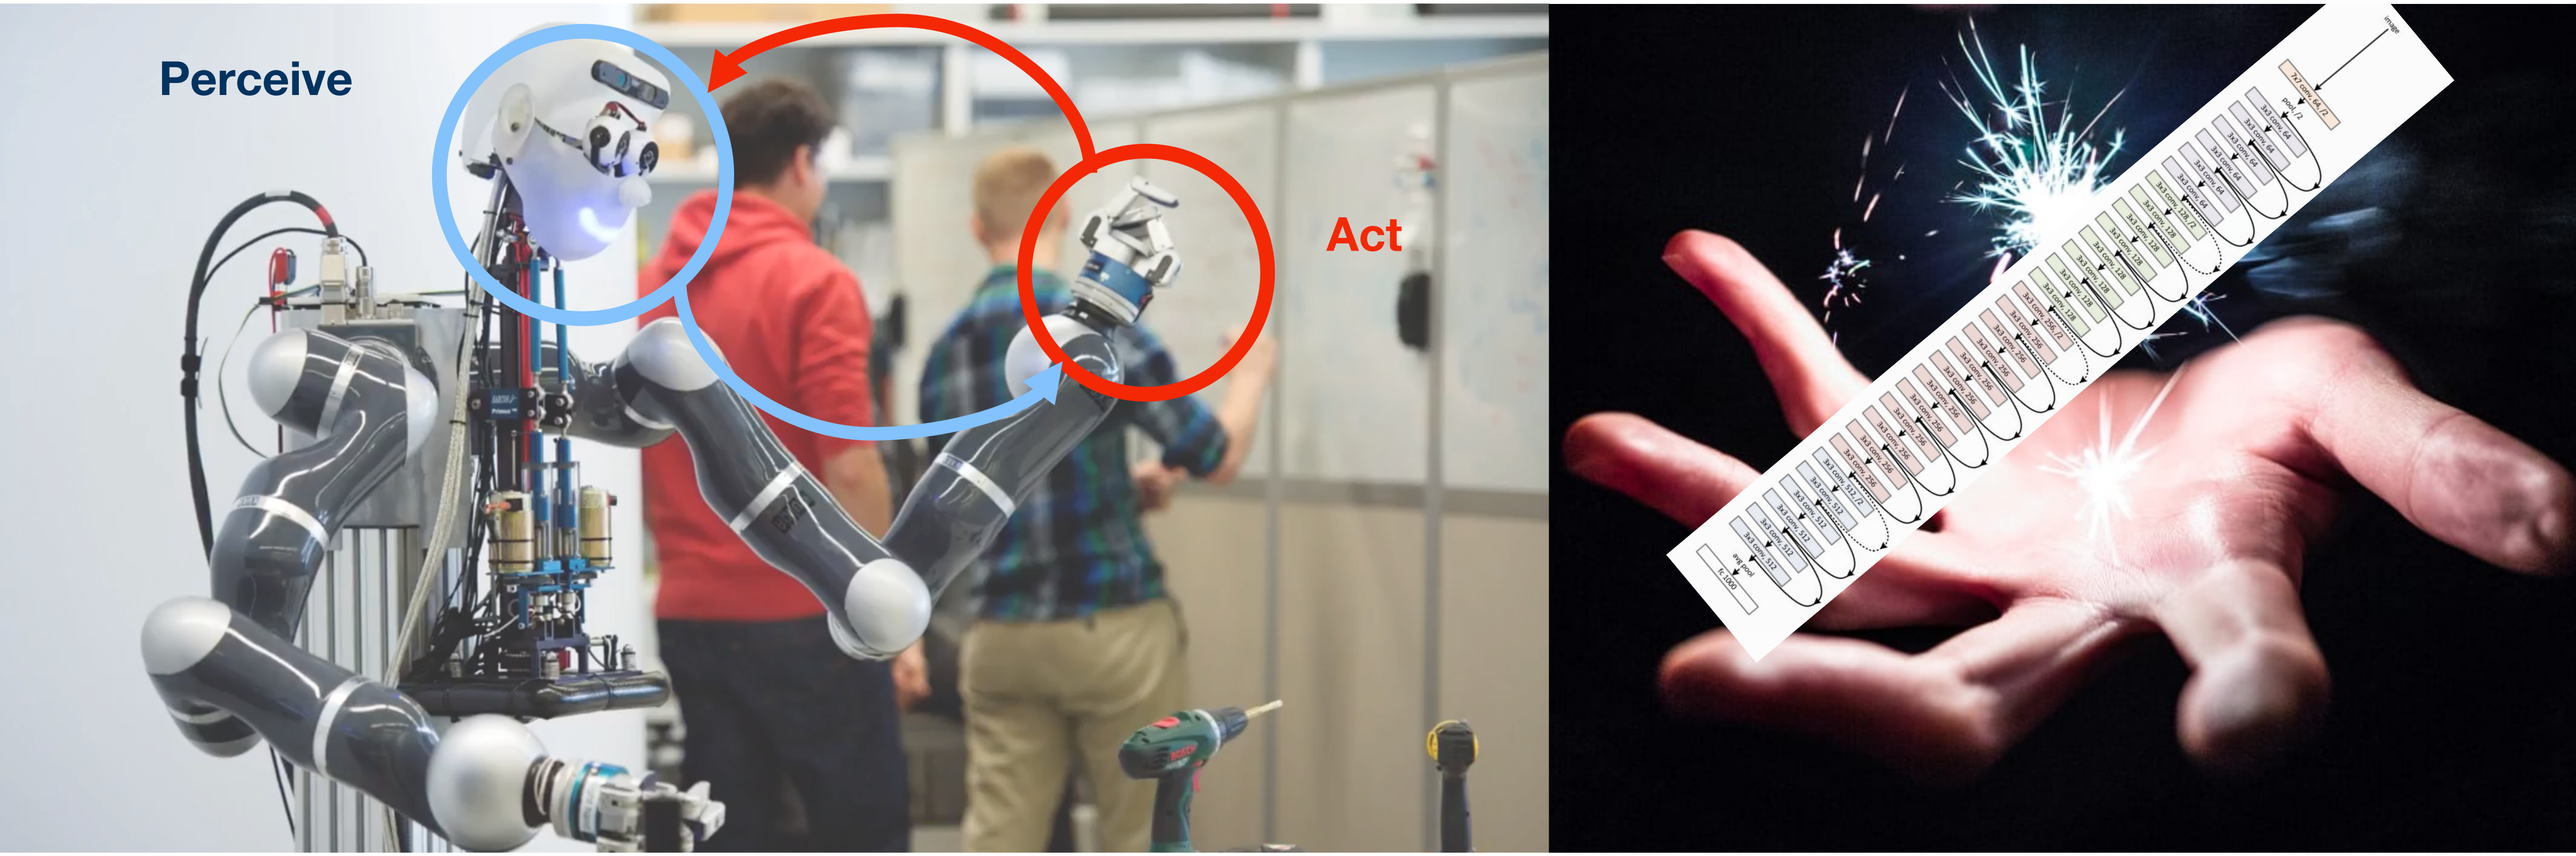

“End-to-End Deep Learning Magic”

Robot Autonomy as a “Pixels-to-Torques” Problem

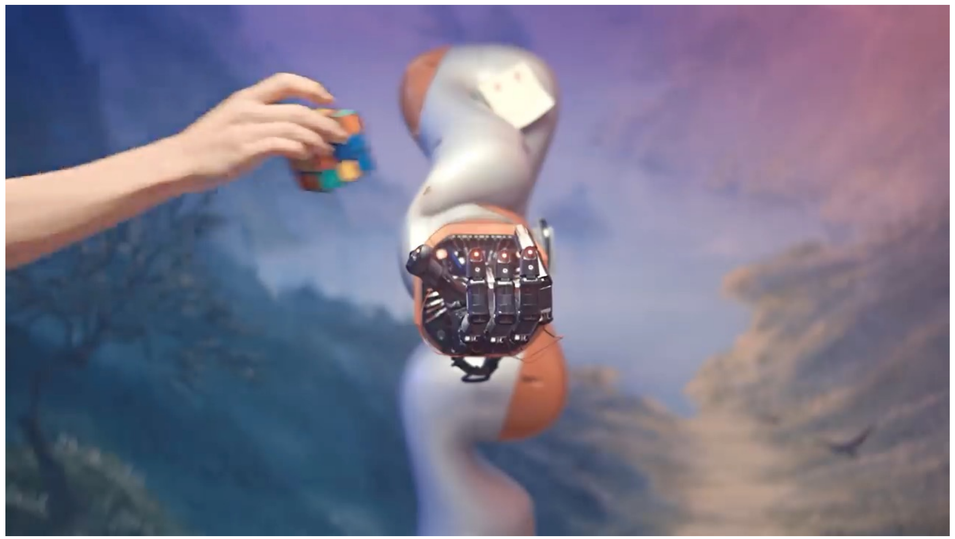

Rubik’s Cube Manipulation (OpenAI; 2019)

Sensorimotor behaviors **emerge** from end-to-end deep reinforcement learning (60% success rate).

Robot played with the Rubik’s Cube for an equivalent of **10,000 years** in simulated training.

[Price & Boutilier 2003; Bentivegna et al. 2004; Latzke et al. 2007; Conn & Peters 2007; Ross & Bagnell 2014; Kumar et al. 2016; Sun et al. 2017; Vecerik et al. 2007; Zhu et al. 2018; Rajeswaran et al. 2018; Cheng et al. 2018; Nair et al. 2018; Pfeiffer et al. 2018; Codevilla et al. 2018]

“There is no ladder to the moon”

π :

pixel

$\mathcal{O}$

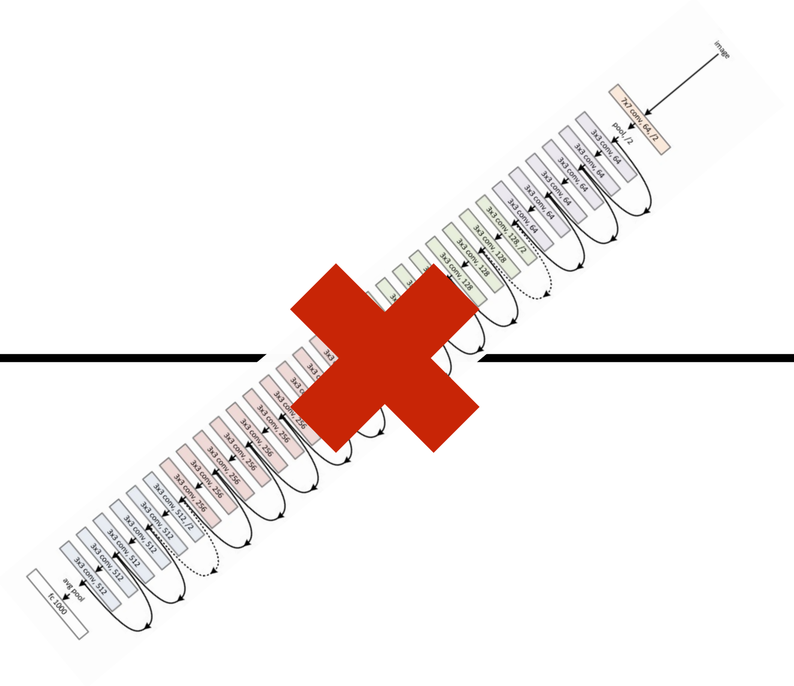

$\mathcal{A}$

torque

My take on robot autonomy: The master algorithm

requires a grand unification of AI methodologies

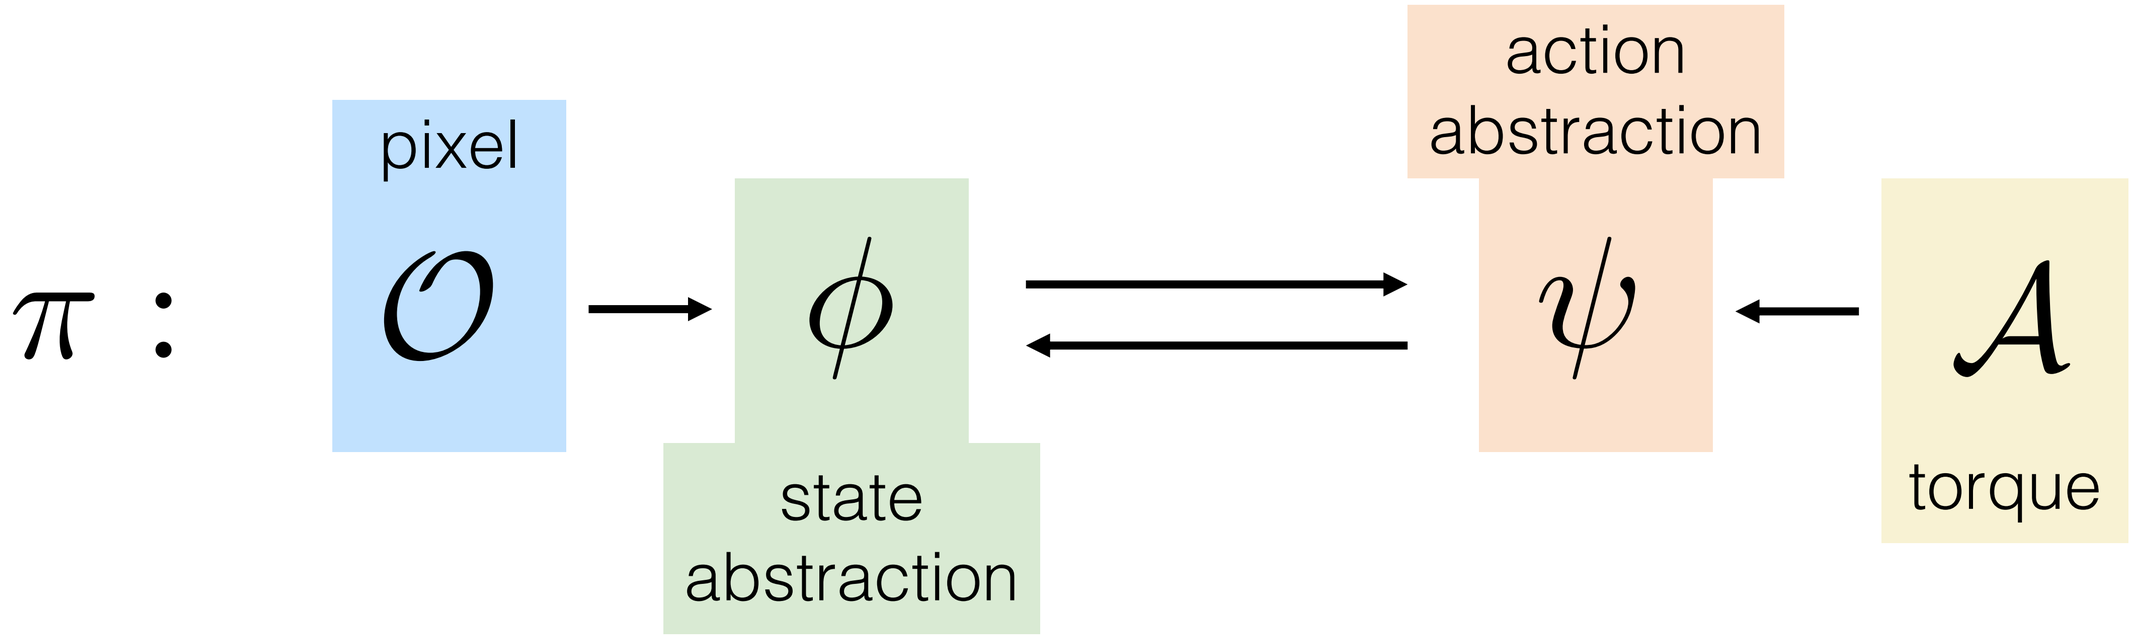

AI methodologies you have learned in CS 343	probabilistic inference	symbolic reasoning	planning	reinforcement learning

Deep Affordance Foresight:

Planning Through What Can Be Done in the Future

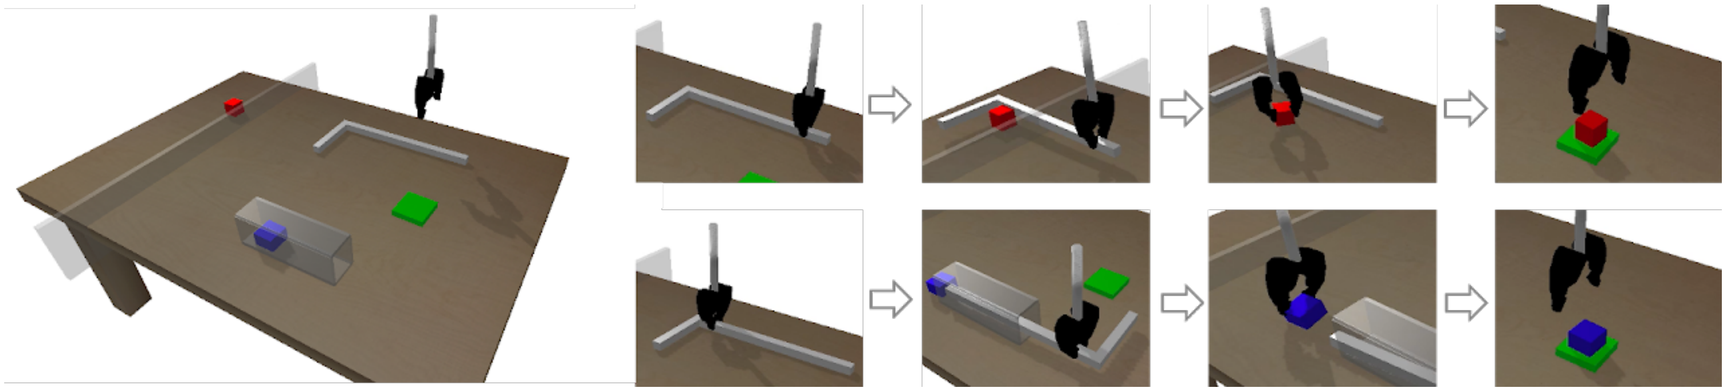

Xu et al. ICRA 2021

What is Affordance?

Affordances are *potential for actions* *successes* that the environment *affords* to the agent. (Gibson, 1977)

[Gibson 1979; Kirklik 1993; Zaff 1995; Stoytchev 2015; Amant 1999; Bousmalis 2018; Detry 2011; Zhu et al. 2017; Mahler 2017; Nagarajan 2020; Ugur 2007; Dang 2020; Fang 2018; Song 2010; Zeng 2018; Abel 2014; Abel 2015; Cruz 2016; Khetarpal 2020; Ardon 2020; Mandikal 2020]

Classical definition of **affordance** is **not** suitable for planning

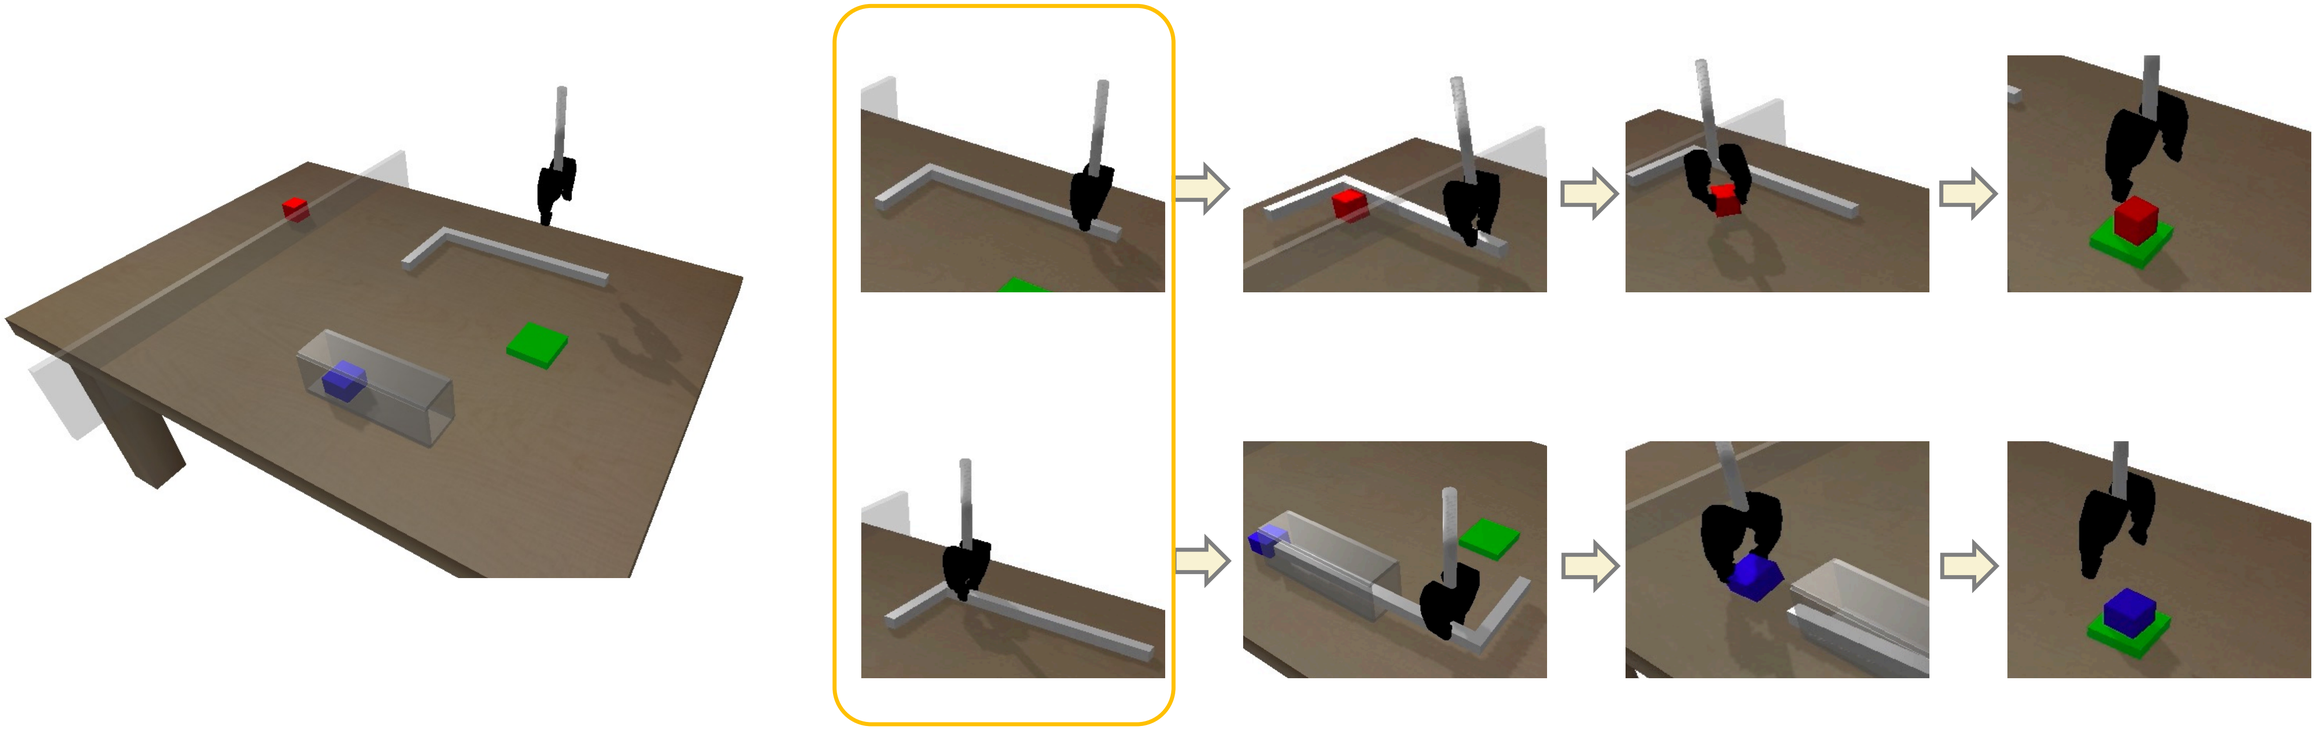

Xu et al. “Deep Affordance Foresight” ICRA 2021

Classical definition of **affordance** is **not** suitable for planning

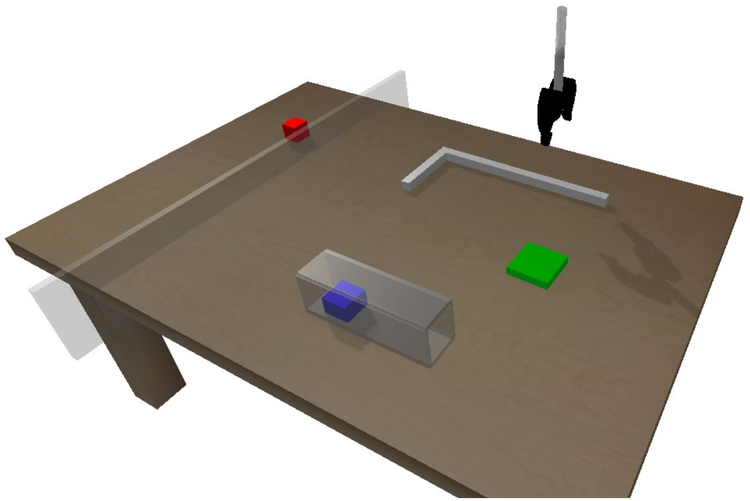

feasible? **YES**

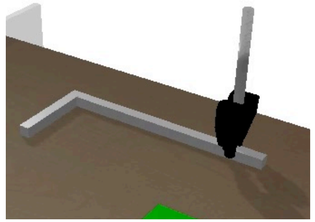

feasible? **YES**

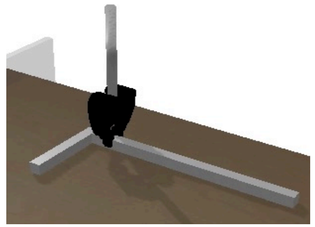

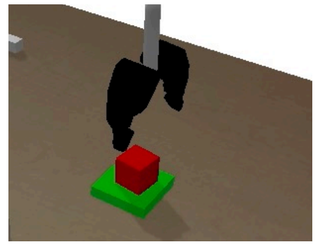

Classical affordance: whether an action is *feasible*

No way to choose actions with respect to a long-horizon task goal

Xu et al. “Deep Affordance Foresight” ICRA 2021

Classical definition of **affordance** is **not** suitable for planning

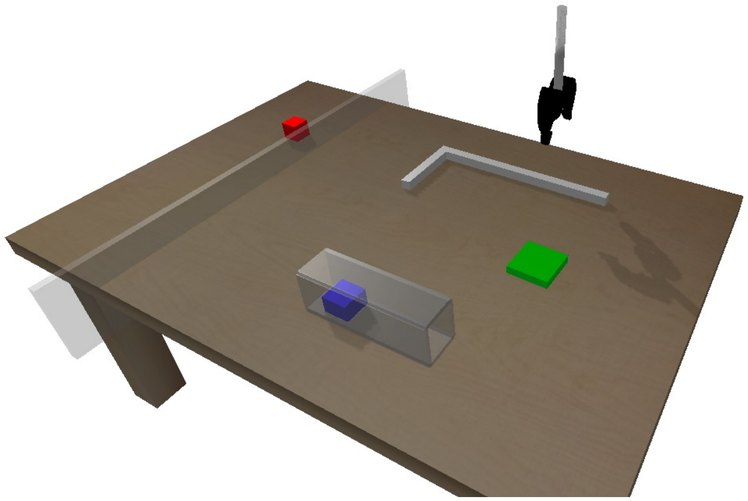

enables hooking red cube? **YES**

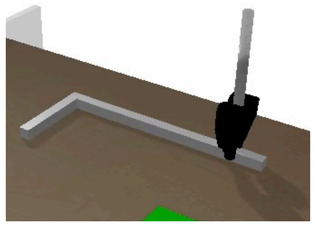

enables grasping red cube? **YES**

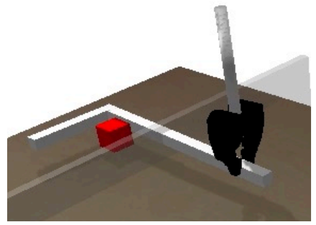

enables placing on target? **YES**

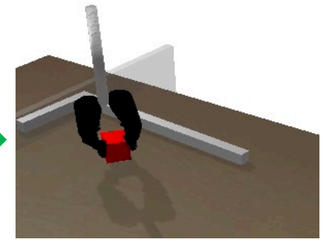

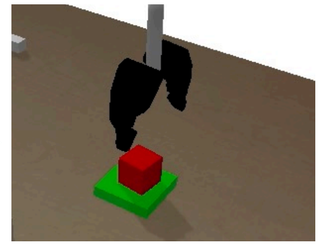

enables fetching red cube? **NO**

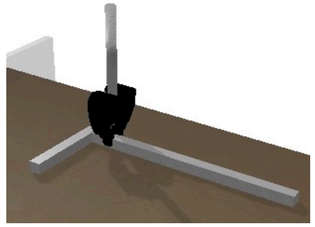

Classical affordance: whether an action is *feasible*

Our new affordance: will an action **make future actions feasible**?

Choose actions that enable the subsequent steps in the task plan

Skill Affordances

Parameterized skill plan grasp(θ) ⟶ hook(θ) ⟶ grasp(θ) ⟶ $\cdots$

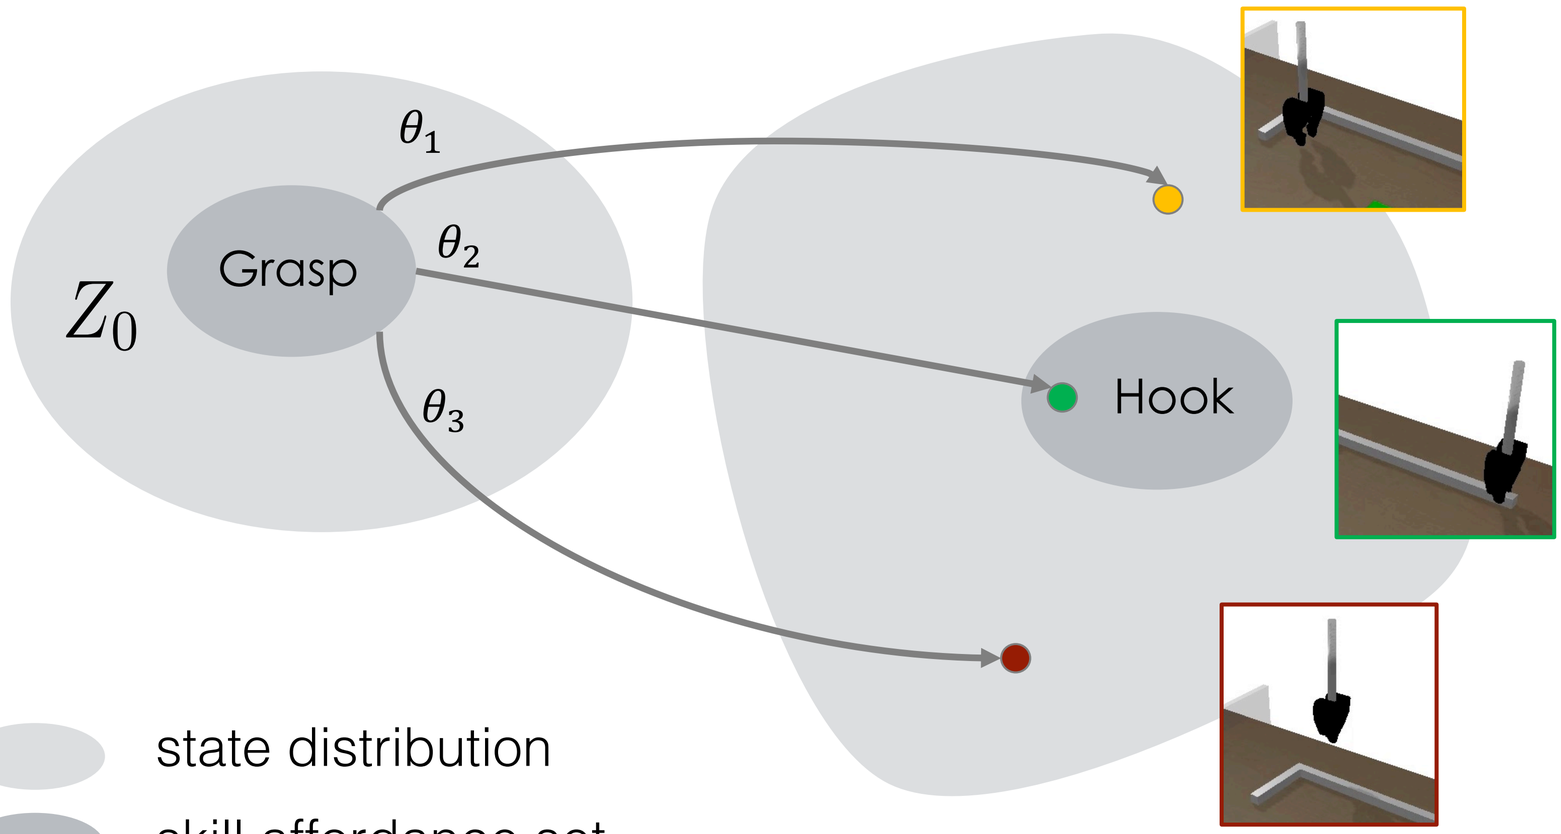

Affordance function:

$$\mathcal{A}_{\pi,\theta}(s)$$

Whether state s is in the affordance set of (π, θ)

Xu et al. “Deep Affordance Foresight” ICRA 2021

Verifying that a plan is executable with affordances

Parameterized skill plan grasp(θ) ⟶ hook(θ) ⟶ grasp(θ) ⟶ ···

How likely is this plan executable by the robot?

The plan is executable if every skill in the plan is afforded.

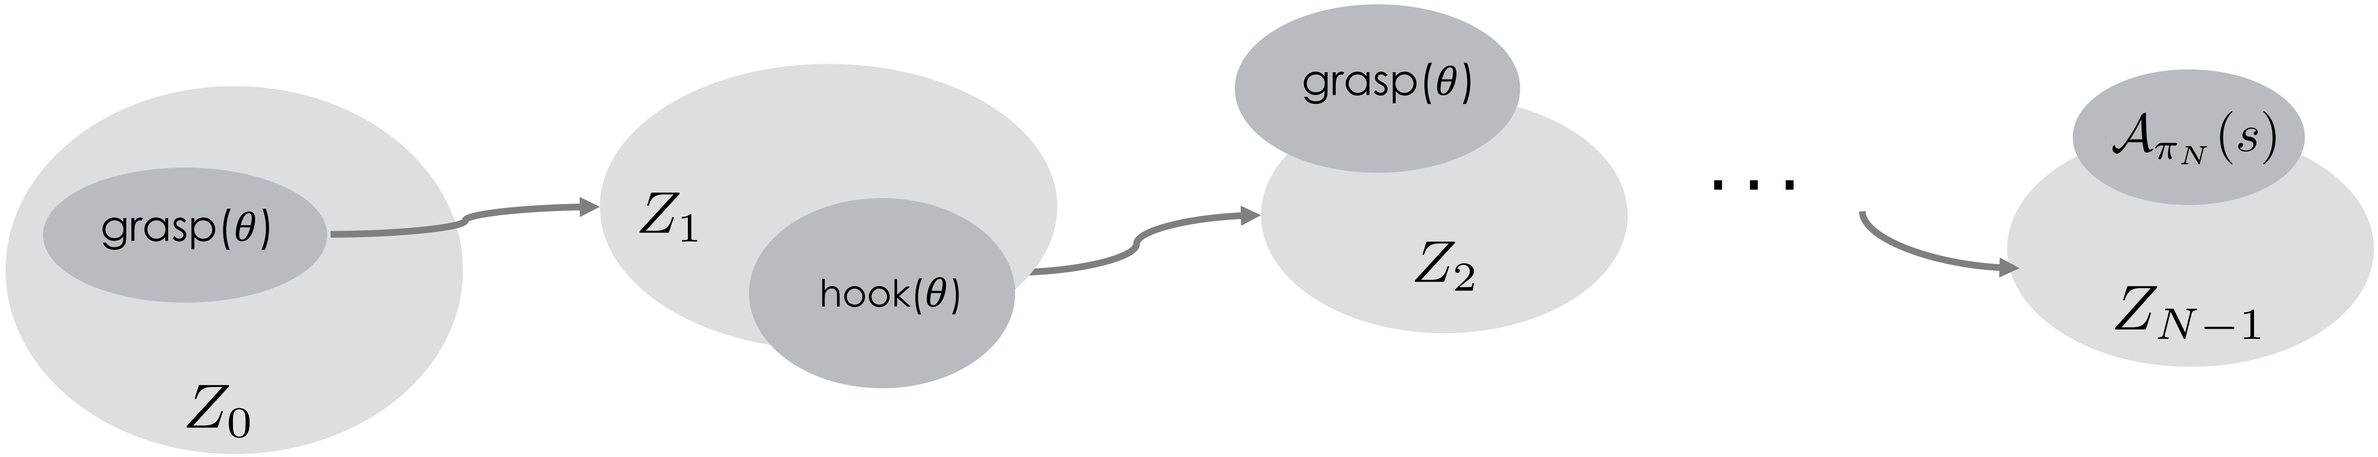

state distribution

skill affordance set

transition function

Probability that a plan is executable:

$$C_{plan}(\{(\pi_i, \theta_i)\}_{i=1}^{N}) = \sum_{s \in S} Z_{N-1}(s) \mathcal{A}_{\pi_N, \theta_N}(s)$$

state distribution

affordance likelihood

Learning to Plan with Deep Affordance Foresight (DAF)

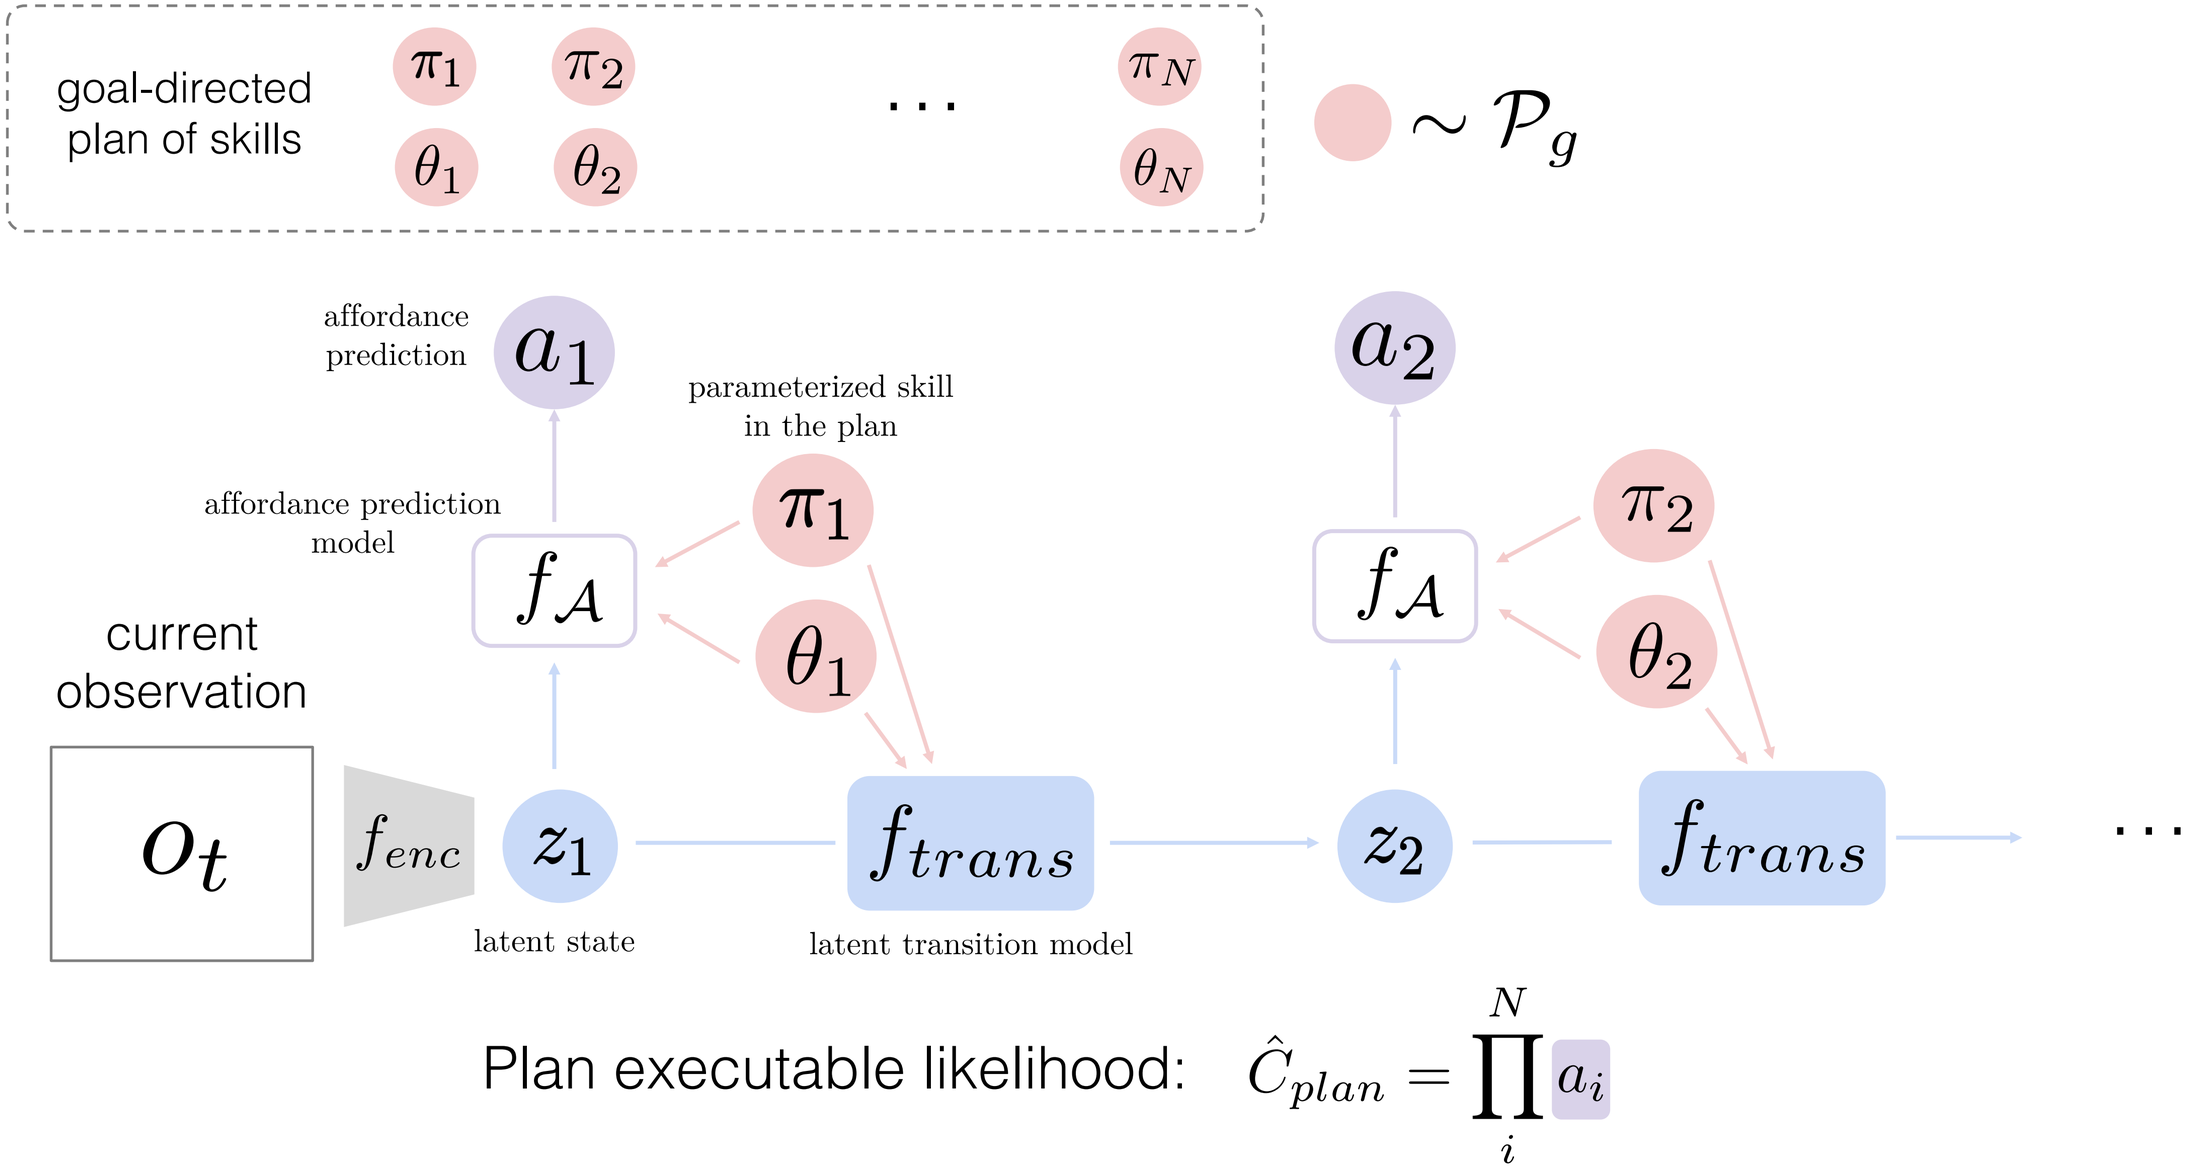

Plan executable likelihood: $\hat{C}_{plan} = \prod_{i}^{N} a_i$

Xu et al. “Deep Affordance Foresight” ICRA 2021

Results: Tool-Use

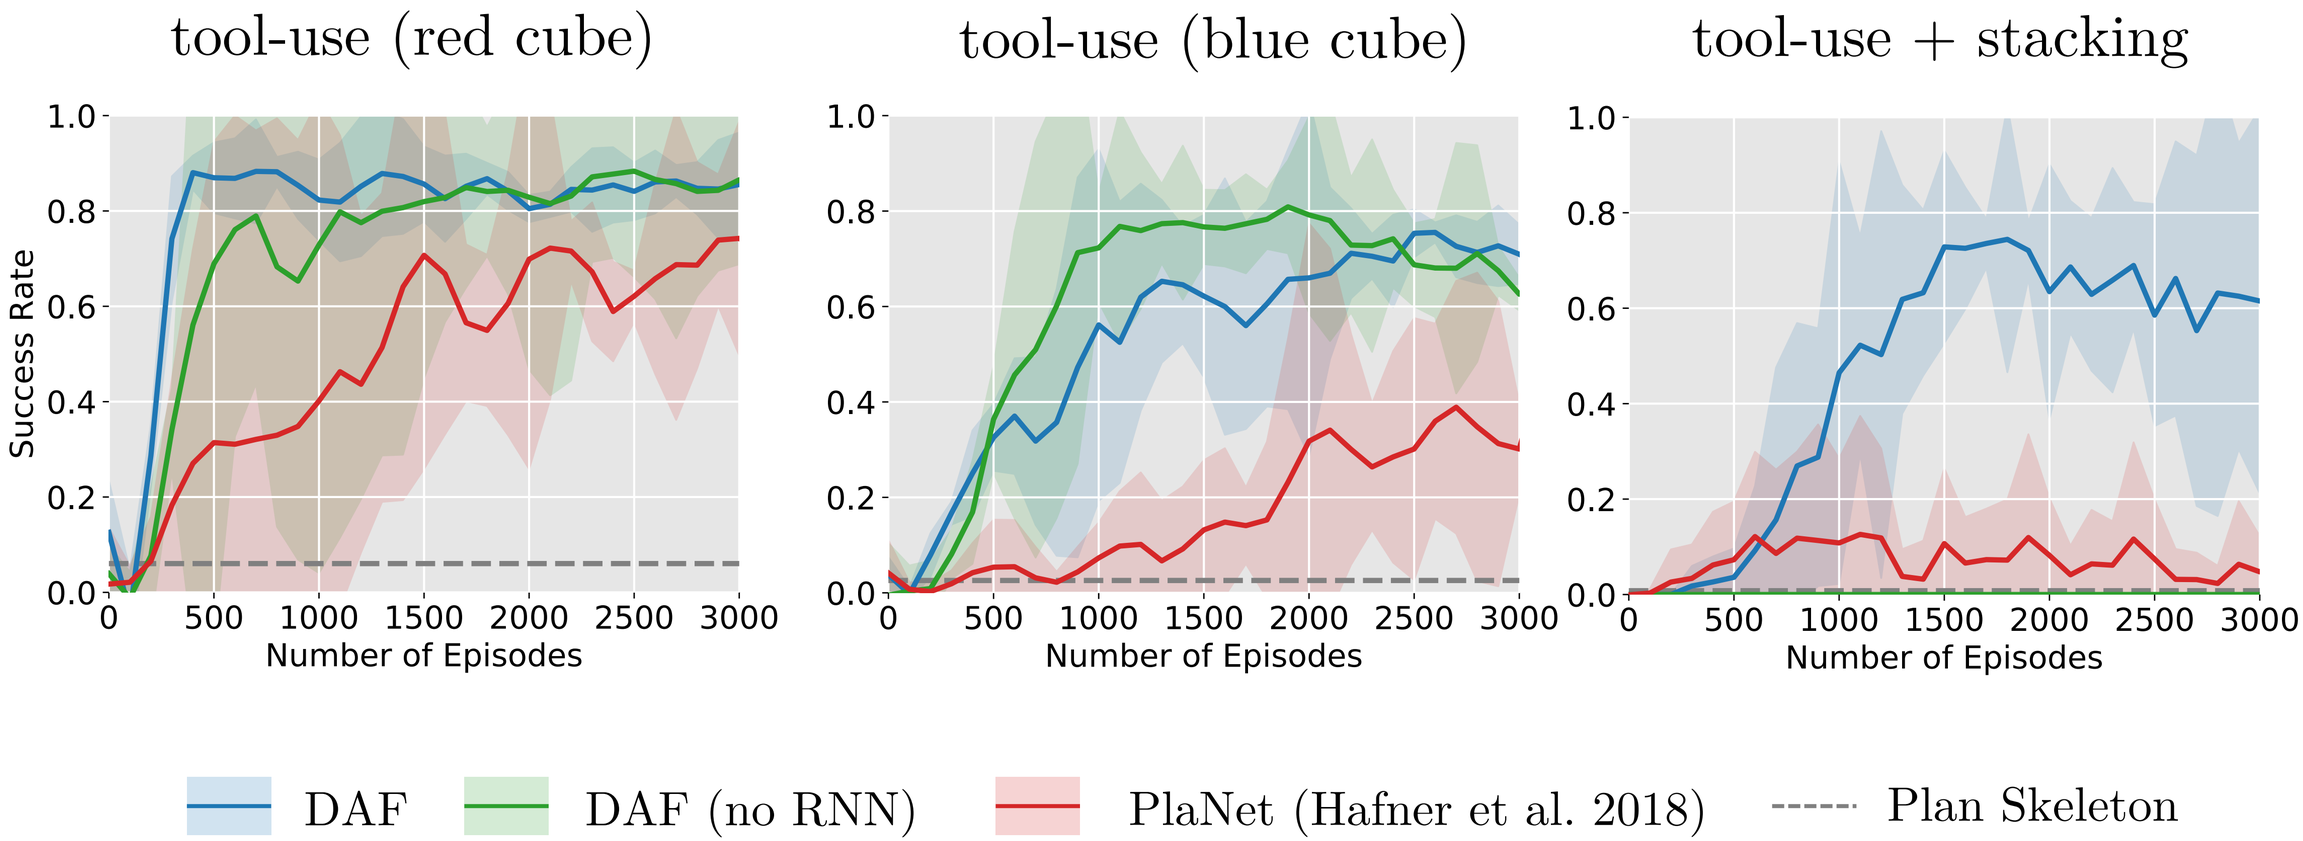

Xu et al. “Deep Affordance Foresight” ICRA 2021

Sample Rollouts: Tool-Use + Stacking

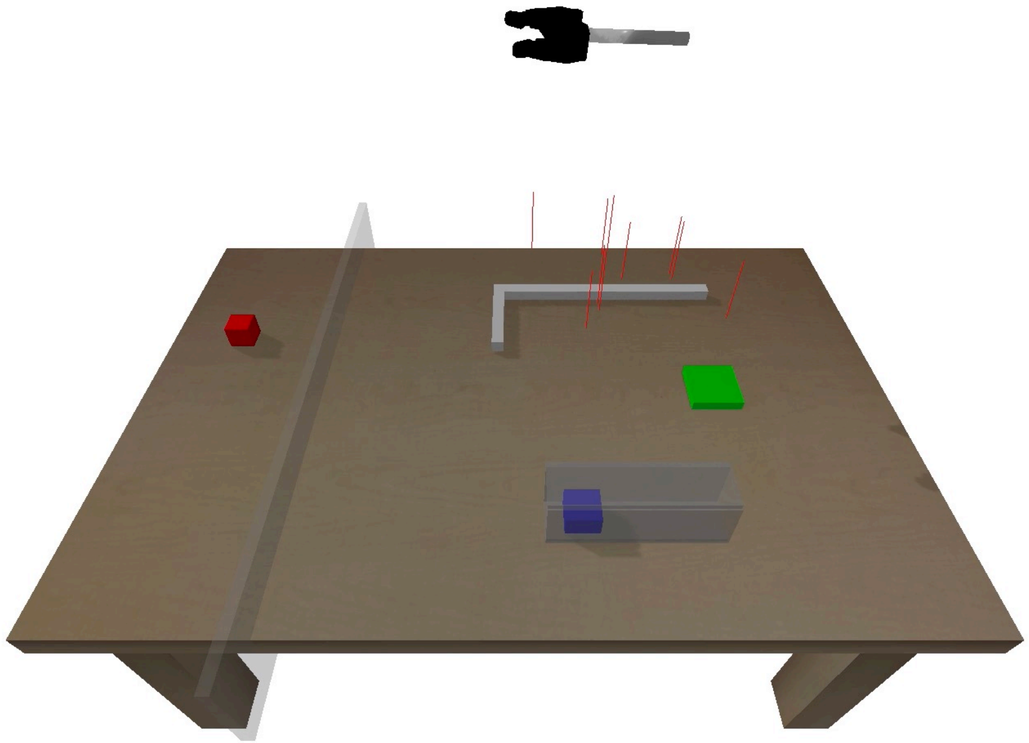

PlaNet

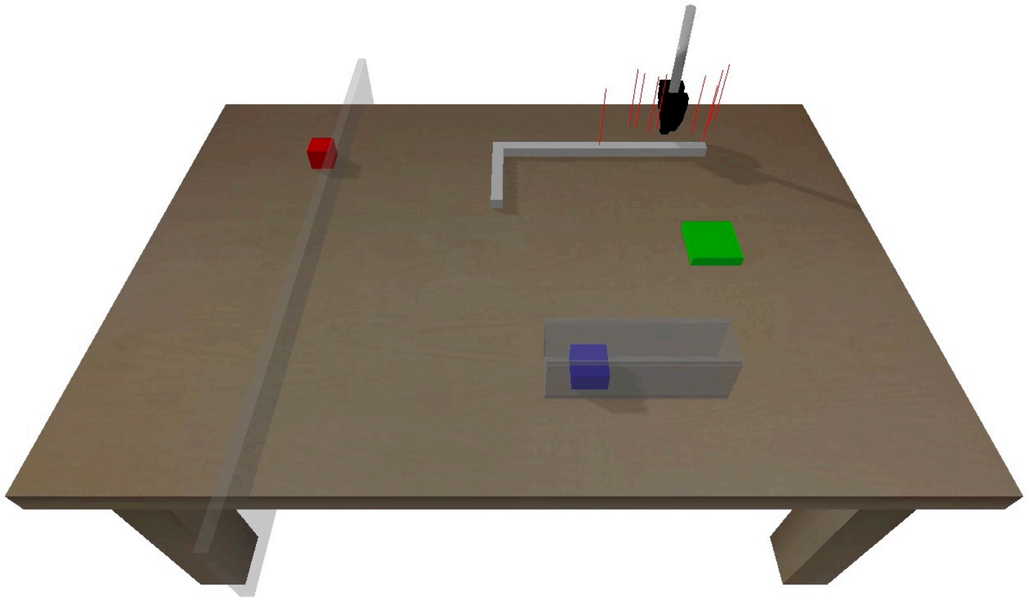

DAF

Xu et al. “Deep Affordance Foresight” ICRA 2021

Rollouts: Kitchen (Coffee)

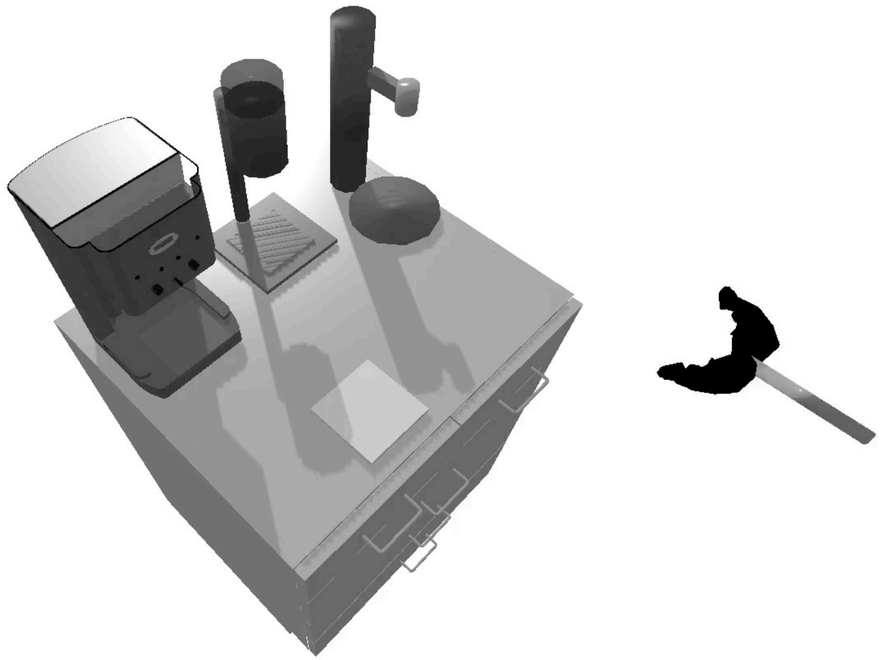

PlaNet

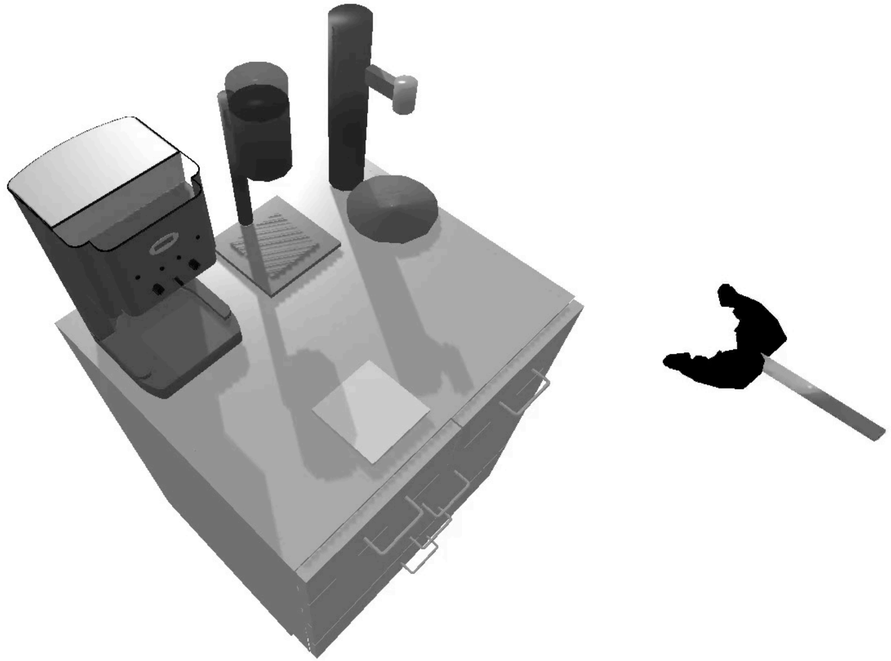

DAF

Xu et al. “Deep Affordance Foresight” ICRA 2021

Deep Affordance Foresight: Summary

- Mapping skills to **action symbols** (skill affordance) from raw observations (NN)
- Reason about skill symbolic plans with **probabilistic methods** of success likelihood (PI)
- Model-based **reinforcement learning** learned from trial and error (RL)

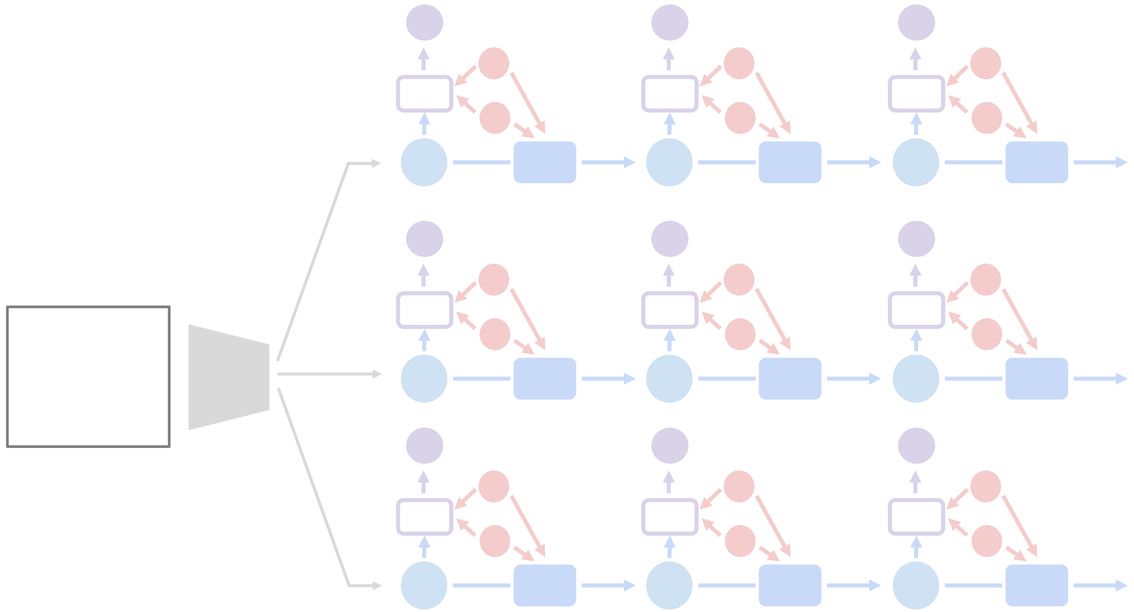

"My advice to young people" Donald Knuth

https://youtu.be/75Ju0eM5T2c