Bayesian Agents

Bayesian Reasoning and Deep Learning in Agent-based Systems

Shakir Mohamed

DeepMind

@shakir_za shakir@google.com

Abstract

Bayesian Agents:
Bayesian Reasoning and Deep Learning in Agent-based Systems

Bayesian deep learning allows us to combine two needed components for building intelligent and autonomous systems: Deep learning, which provides a powerful framework for model building, and Bayesian analysis, which provides tools for optimal inference in these models. The outcome of this convergent thinking is our ability to develop and train a broad set of tools that are important components of systems that can reason and act in the real world. In this talk, we shall explore some of the ways in which Bayesian deep learning can be used in the tasks we expect from intelligent systems, such as scene understanding, concept formation, future-thinking, planning, and acting. These approaches remain far from perfect, and they allow us to unpack some of the challenges that remain for even wider application of Bayesian deep learning, and Bayesian reasoning more generally.

Deep Learning

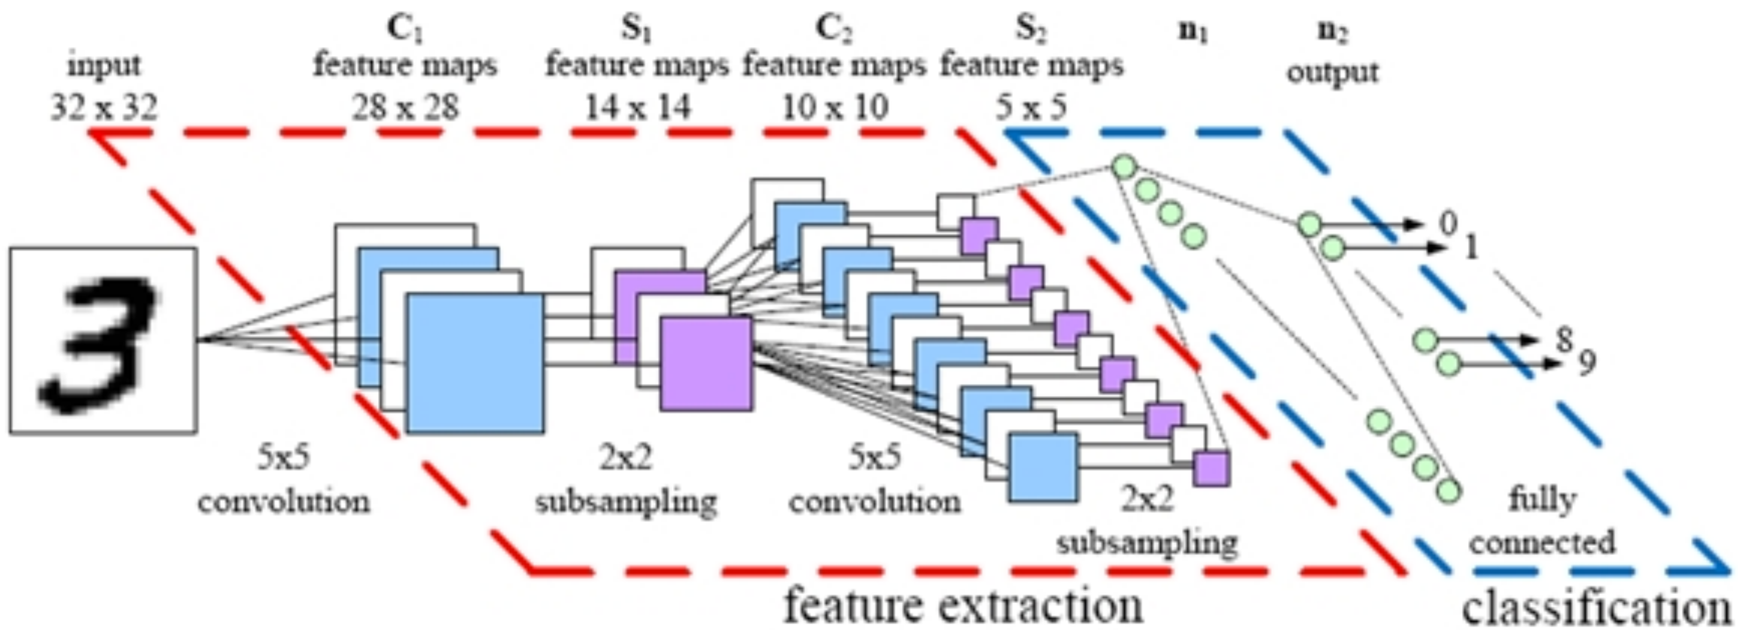

A framework for constructing flexible models

+ Rich non-linear models for classification and sequence prediction.

+ Scalable learning using stochastic approximations and conceptually simple.

+ Easily composable with other gradient-based methods

- Only point estimates

- Hard to score models, do model selection and complexity penalisation.

Bayesian Reasoning

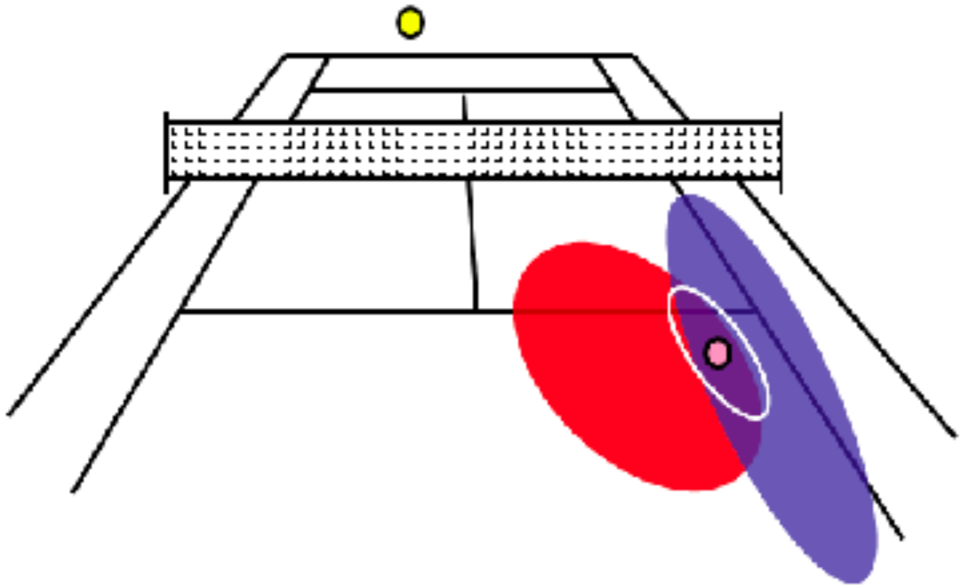

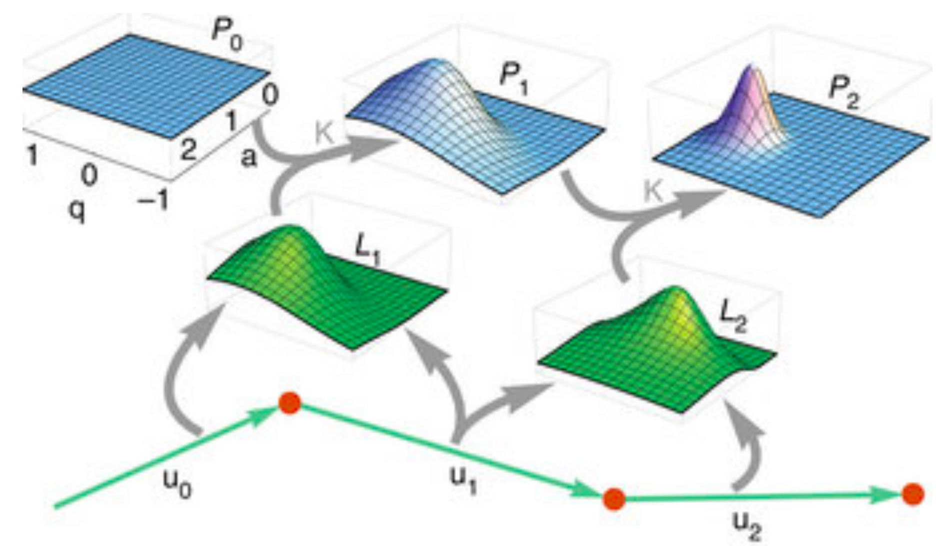

A framework for inference and decision making

+ Unified framework for model building, inference, prediction and decision making

+ Explicit accounting for uncertainty and variability of outcomes

+ Robust to overfitting; tools for model selection and composition.

- Mainly conjugate and linear models

- Intractable inference leading to expensive computation or long simulation times.

Bayesian Deep Learning

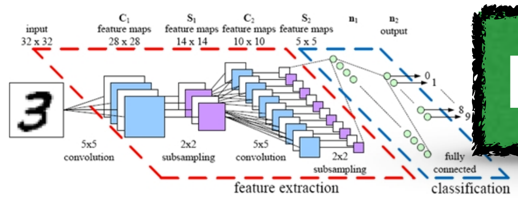

Deep Learning

+ Rich non-linear models for classification and sequence prediction.

+ Scalable learning using stochastic approximation and conceptually simple.

+ Easily composable with other gradient-based methods

- Only point estimates

- Hard to score models, do selection and complexity penalisation.

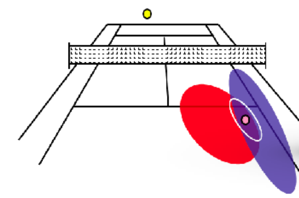

Bayesian Reasoning

- Many conjugate and linear models

- Potentially intractable inference, computationally expensive or long simulation time.

+ Unified framework for model building, inference, prediction and decision making

+ Explicit accounting for uncertainty and variability of outcomes

+ Robust to overfitting; tools for model selection and composition.

Natural to marry these approaches.

Inference and Decision-making

Two Distinct Processes:

What we can *know* about our data

What we can *do* with our data.

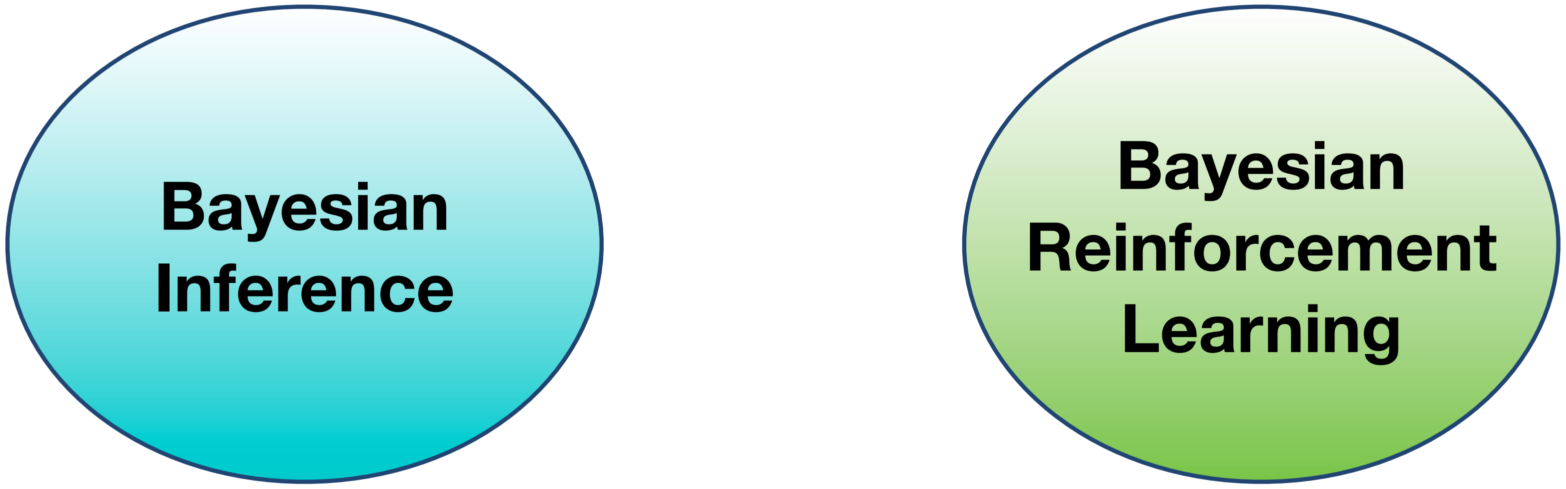

Have the core tools to build reasoning systems:

1. Flexible ways of building models
2. Ability to learn and make consistent inferences and maintain beliefs
3. Reason about potential outcomes and take actions.

Probabilistic Models

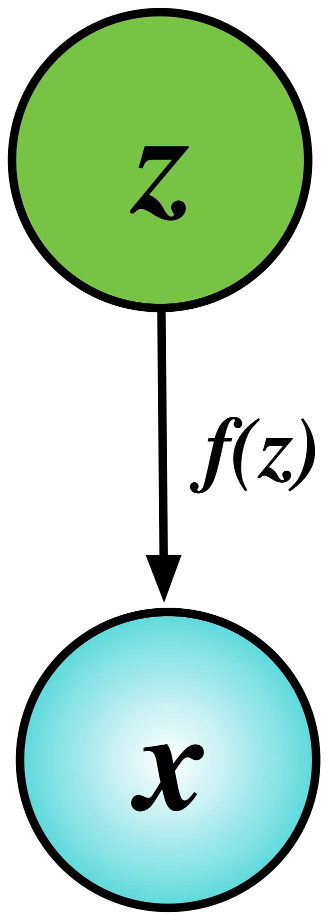

Prescribed Probabilistic Models:
Likelihood helps avoid pathologies of support and lots of different approaches.

Variational inference

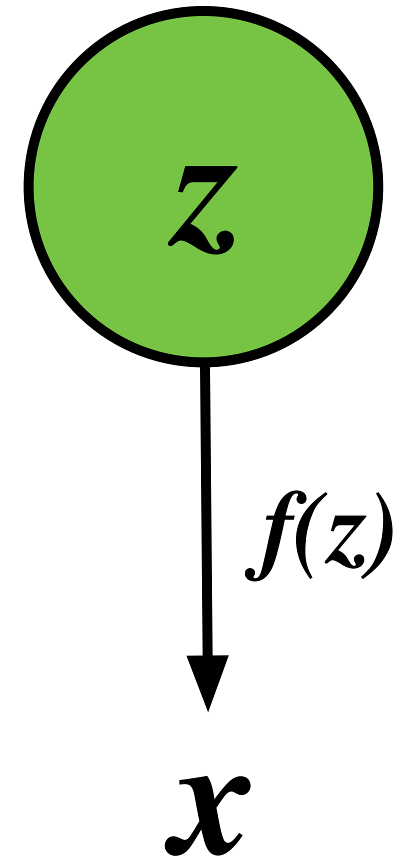

Implicit Probabilistic Models:
A more fundamental way of building models

Hypothesis test-driven Learning

A Pragmatic Bayesian Approach

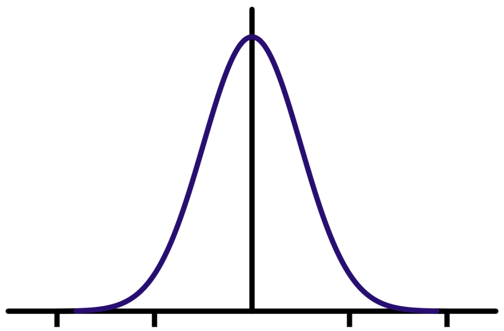

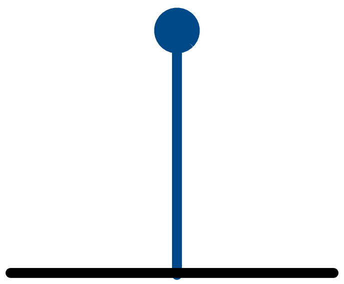

Full Posterior ⟷ ***Point Estimation***

$p(z|x)$ $\delta(z^*)$

Bayesian analysis thinks of probability as a belief.
We situate ourselves in this space subject to resource constraints, e.g., computation, memory.

Default Statistical Inference

Use variational inference as a default approach for statistical inference.

$$p(\mathbf{y}|\mathcal{M}) = \int p(\mathbf{y}|\mathbf{z})p(\mathbf{z})dz$$

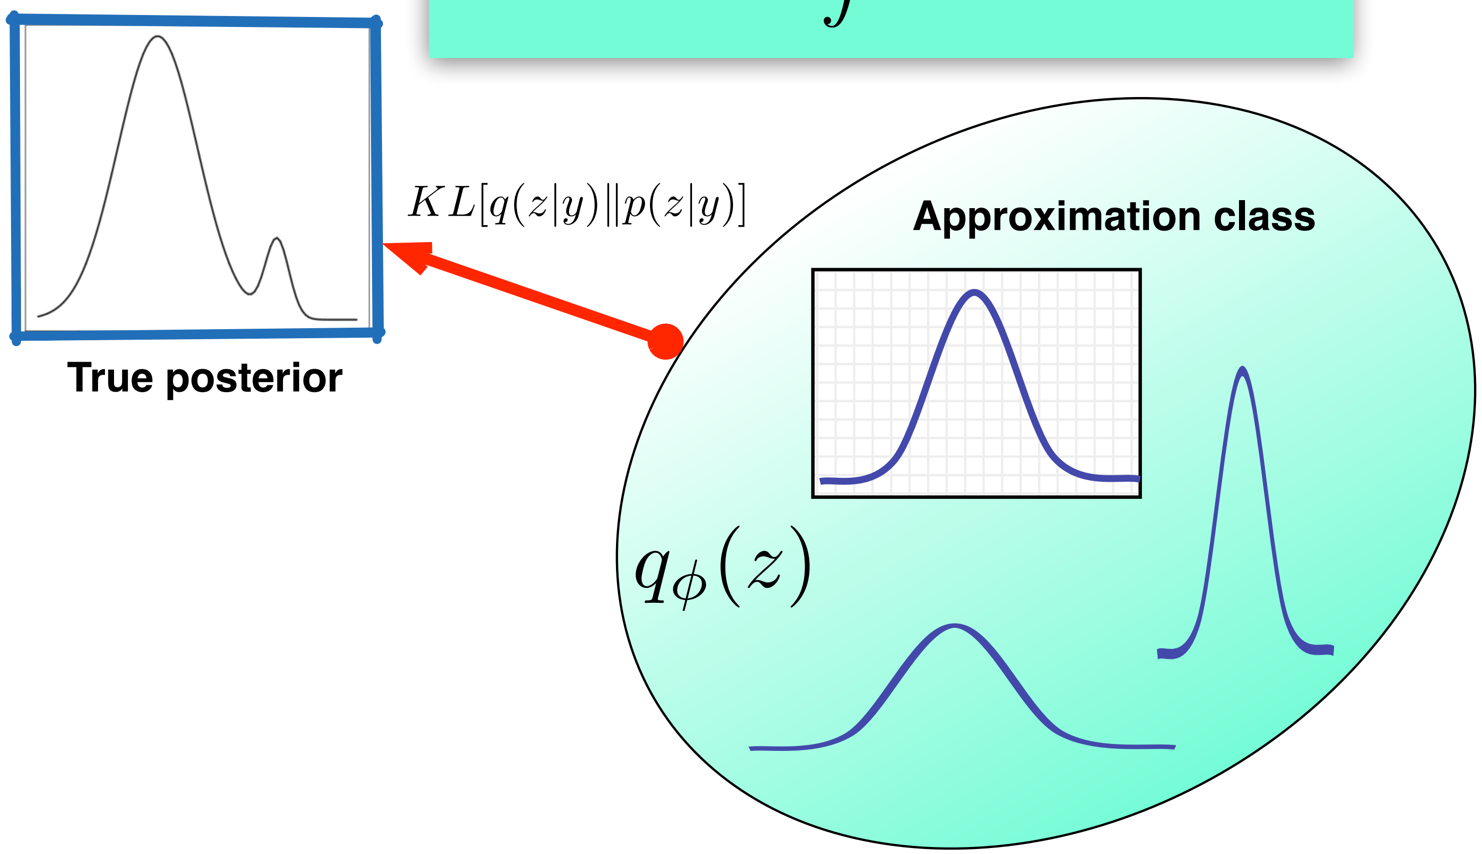

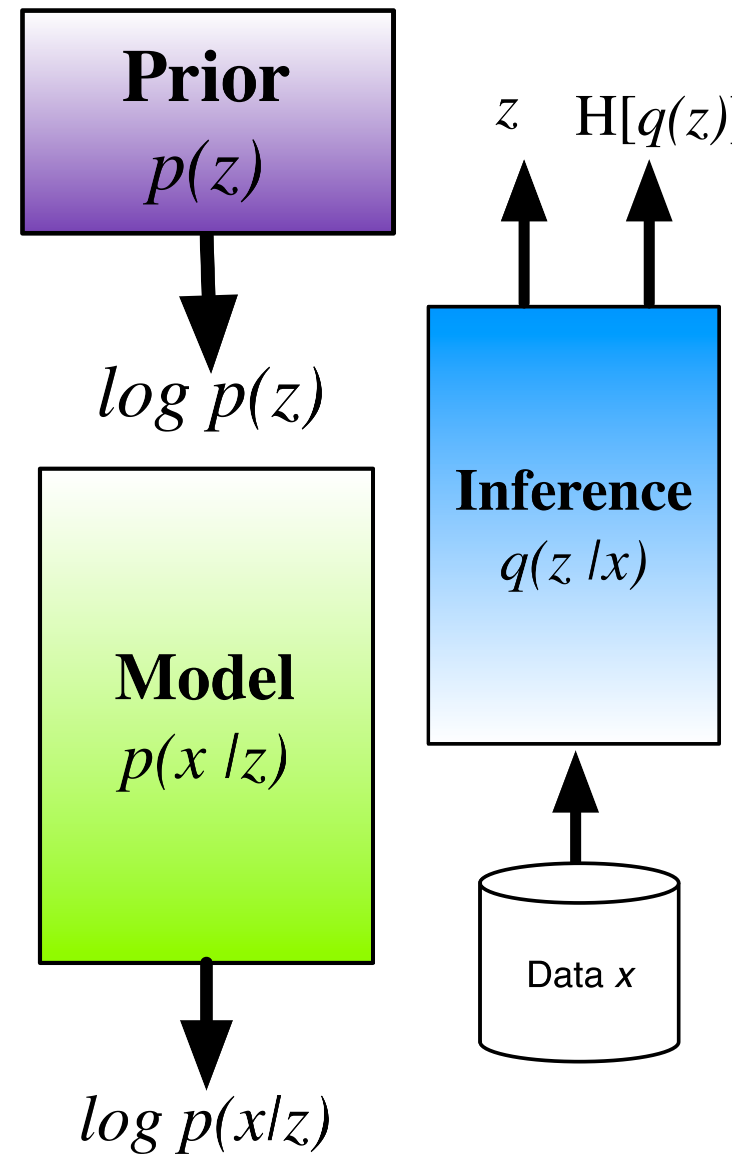

$$\mathcal{F}(y,q) = \underbrace{\boxed{\mathbb{E}_{q(z)}[\log p(y|z)]}}_{\text{Reconstruction}} - \underbrace{\boxed{KL[q(z)\|p(z)]}}_{\text{Penalty}}$$

Highly rich approaches for designing the distribution q available.

Hypothesis Test-driven Learning

Bayesian Inference

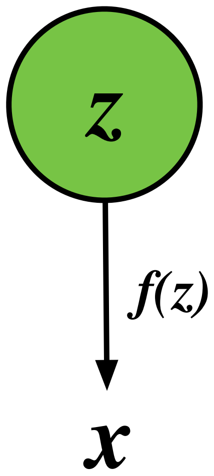

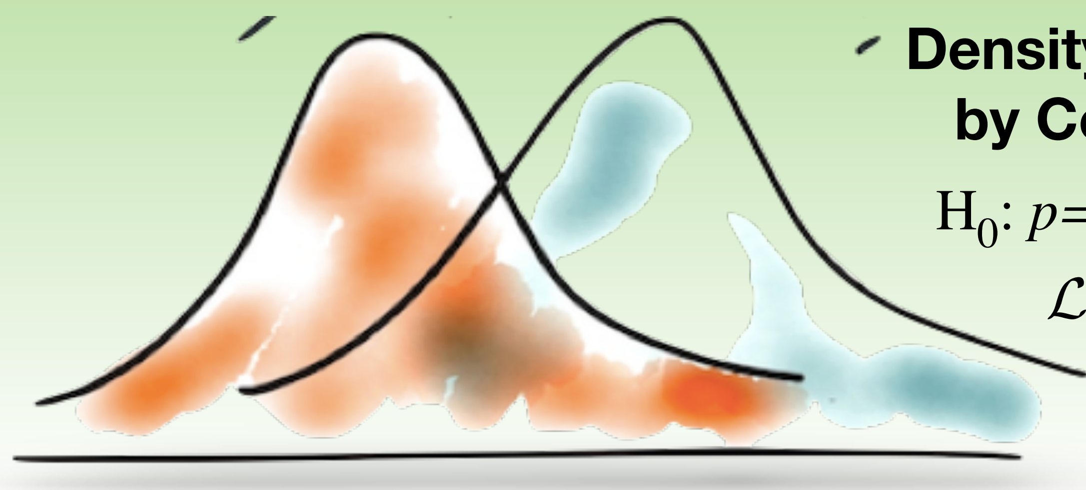

Density Estimation by Comparison

H_0: $p=q$ vs H_1: $p\neq q$

$\mathcal{L}(\theta, \phi)$

Density Difference

$r = p - q$

Density Ratio

$r = \frac{p}{q}$

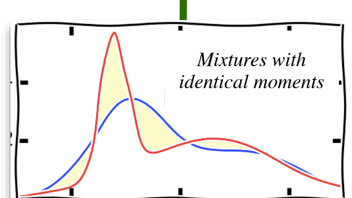

Moment matching

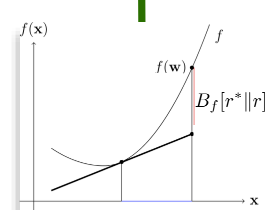

Bregman Divergence

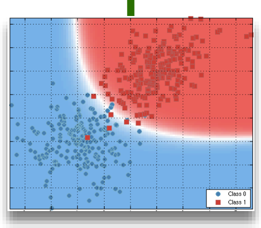

Class Probability Estimation

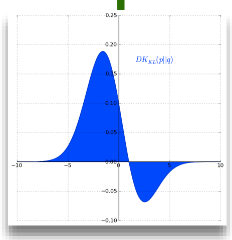

f-Divergence

$f(u) = u \log u - (u+1)\log(u+1)$

Bayesian: Classifier ABC, ABC-MCMC. *Point-estimation*: GAN, GSMM

Bayesian Policy Search

Bayesian Reinforcement Learning

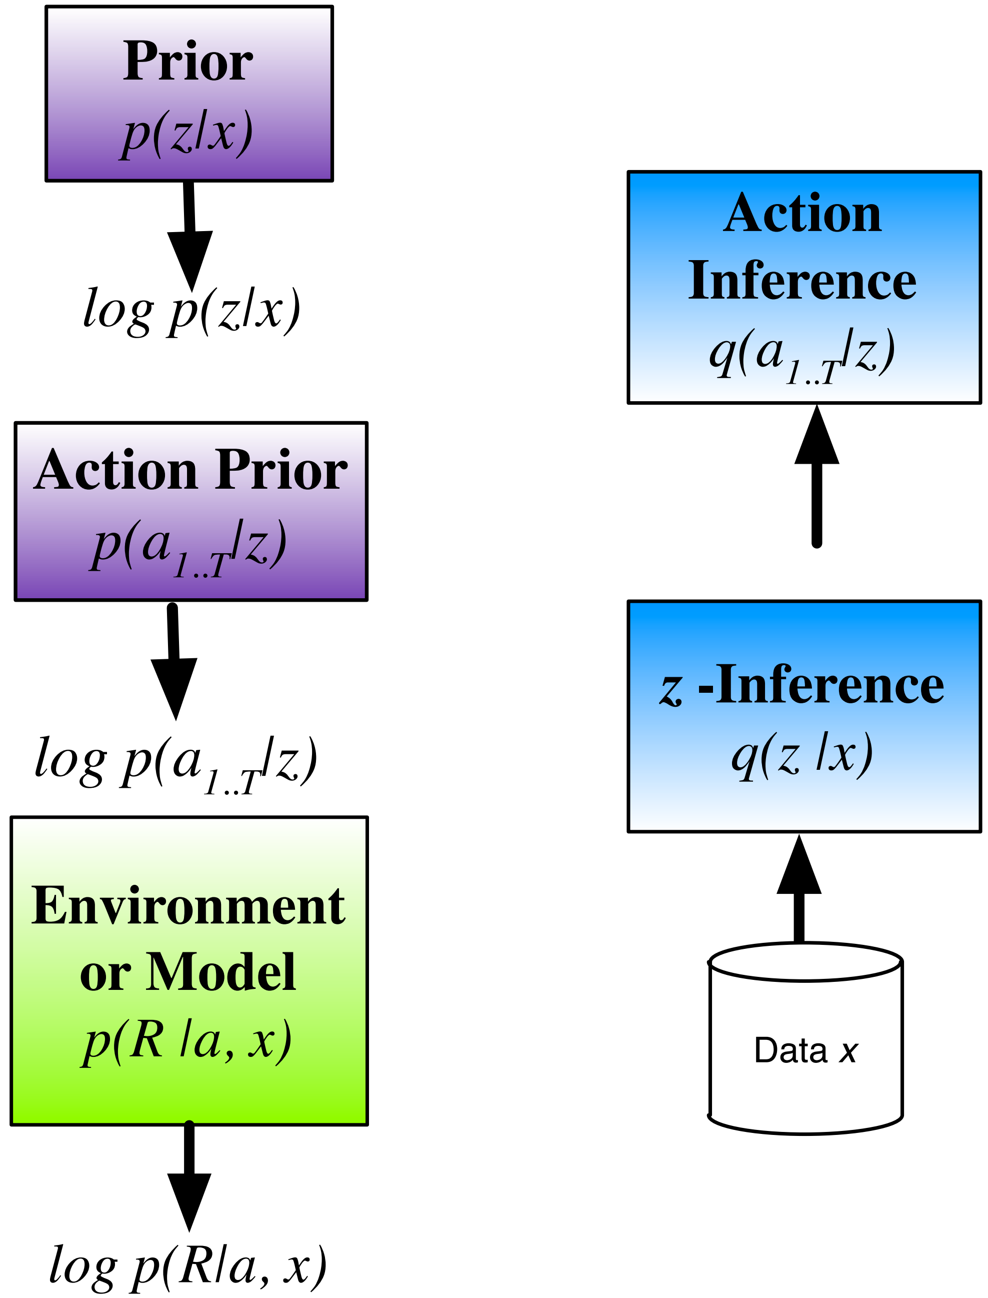

Variational MDP

$$\mathcal{F}^{\pi}(\theta) = \mathbb{E}_{q(a,z|x)}[R(a|x)] - \alpha KL[q_{\theta}(\mathbf{z}|\mathbf{x})\|p(\mathbf{z}|\mathbf{x})] + \alpha\mathbb{H}[\pi_{\theta}(\mathbf{a}|\mathbf{z})]$$

Agent Reasoning

Concrete tasks of a self-reliant agent:

- Reason about physics, objects and scenes.
- Form conceptual understanding of environments.
- Engage future thinking by making predictions about future outcomes and counterfactuals.
- Make decisions and take actions.

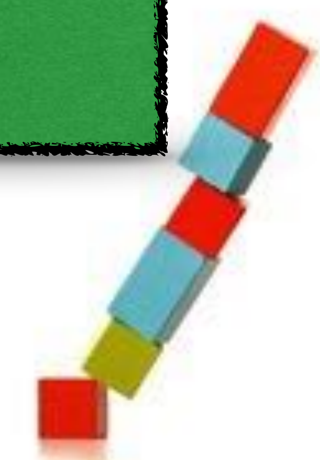

Scene Understanding

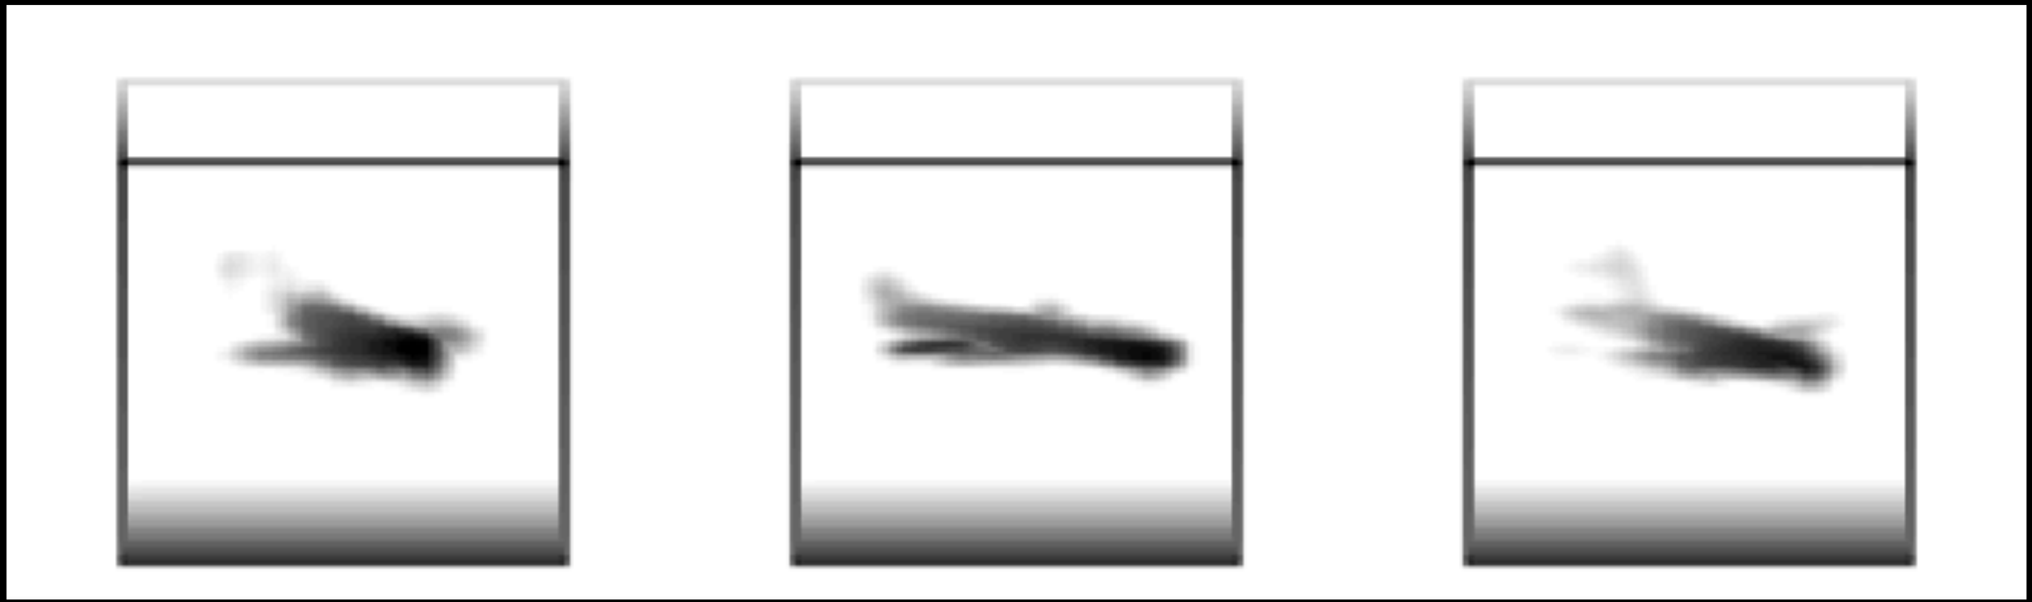

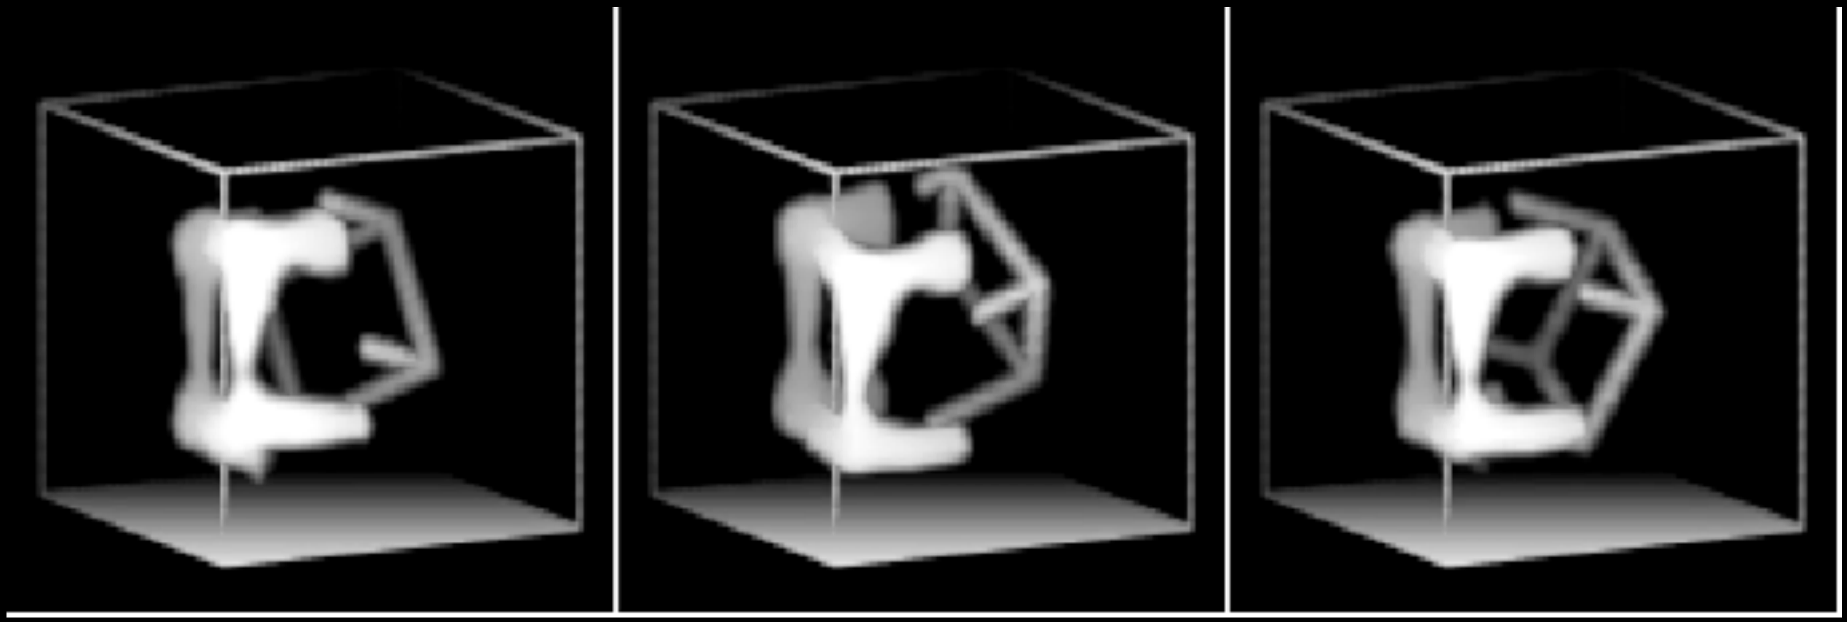

Scene Understanding

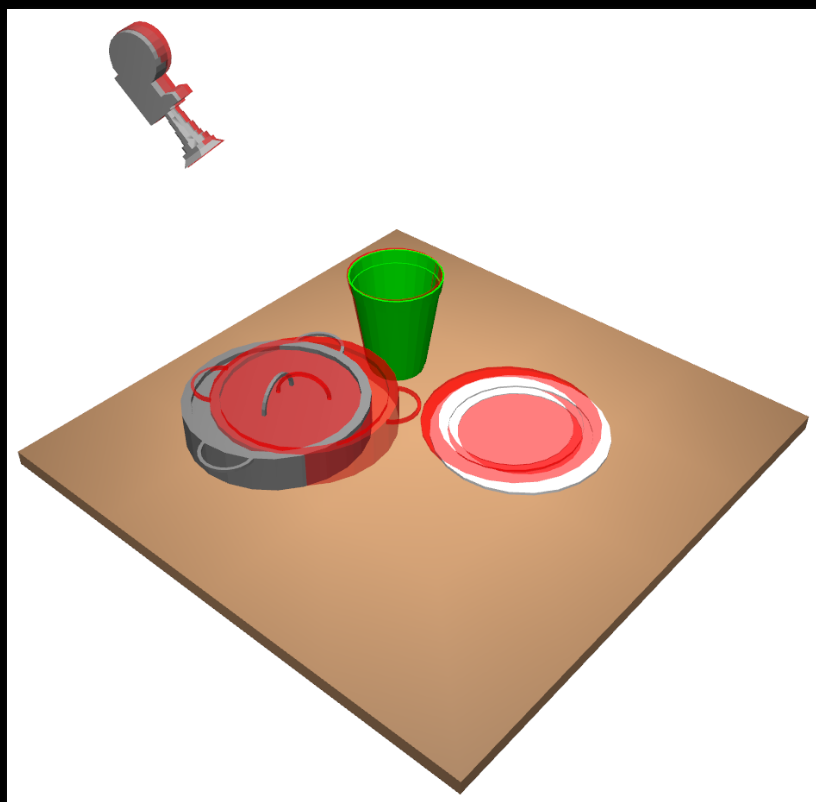

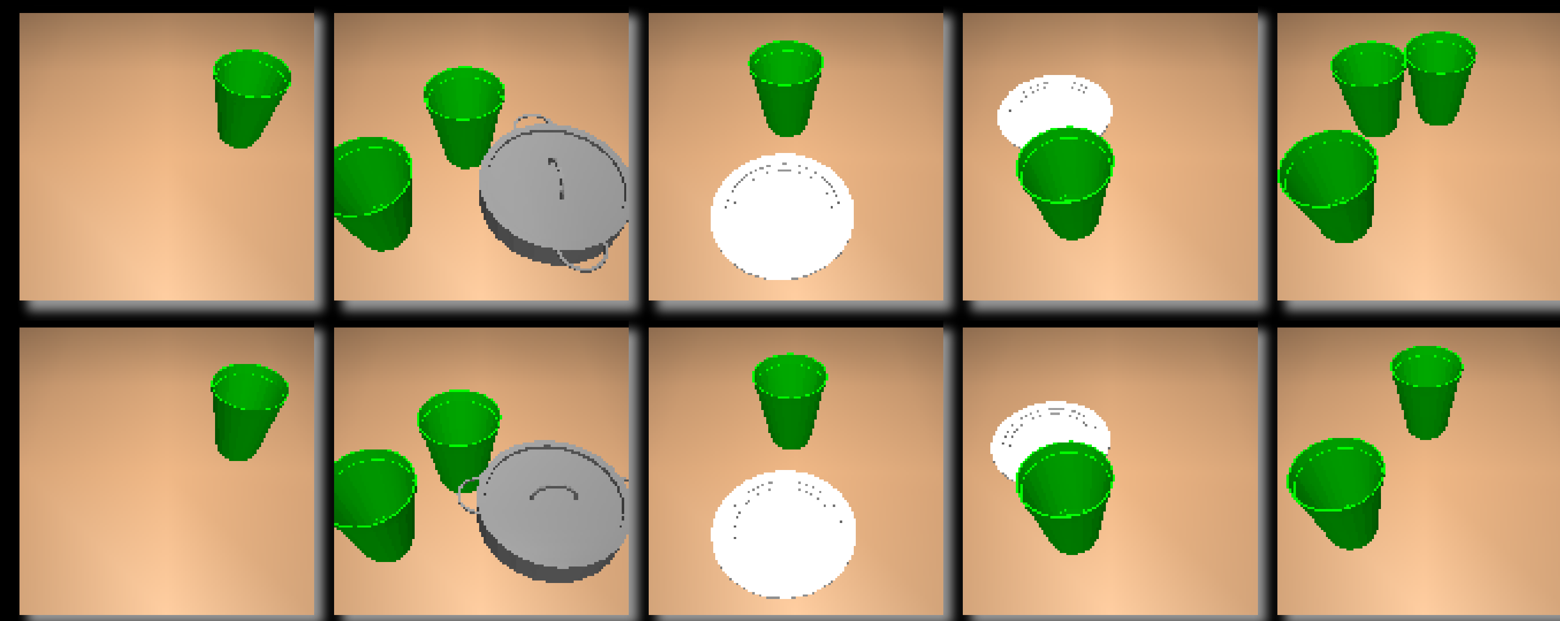

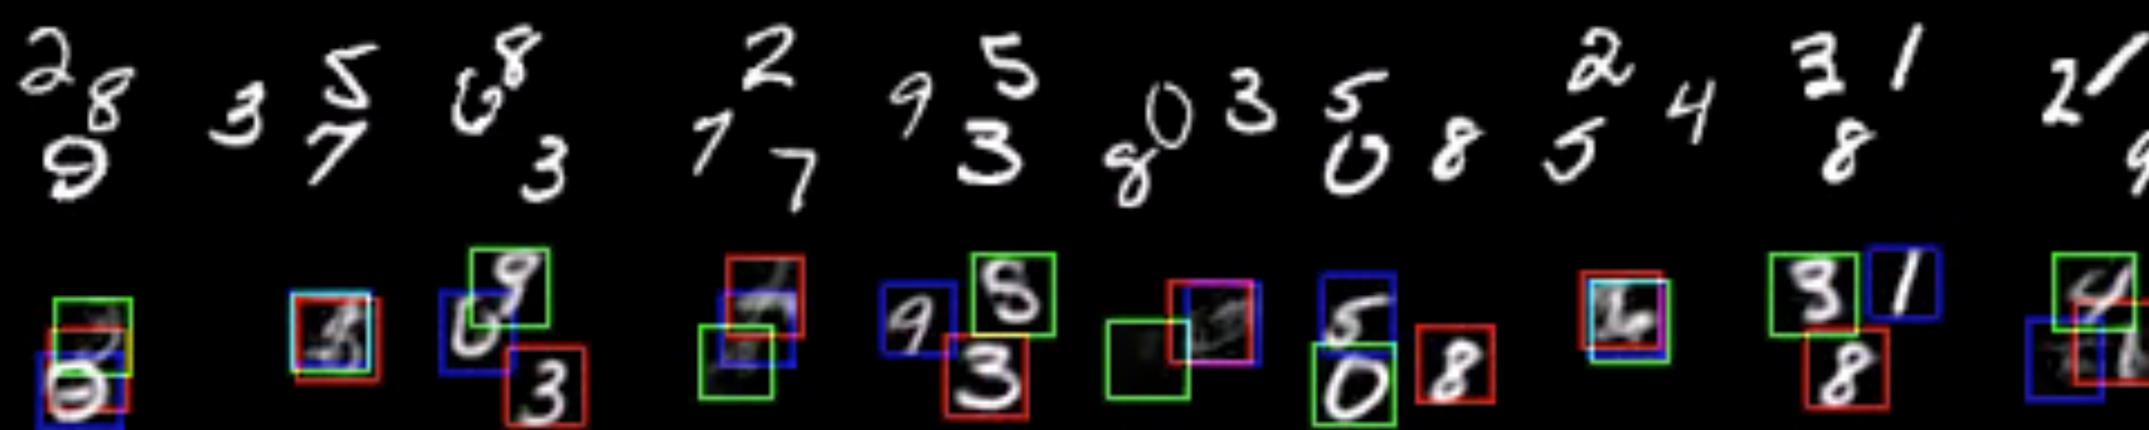

Concept Learning

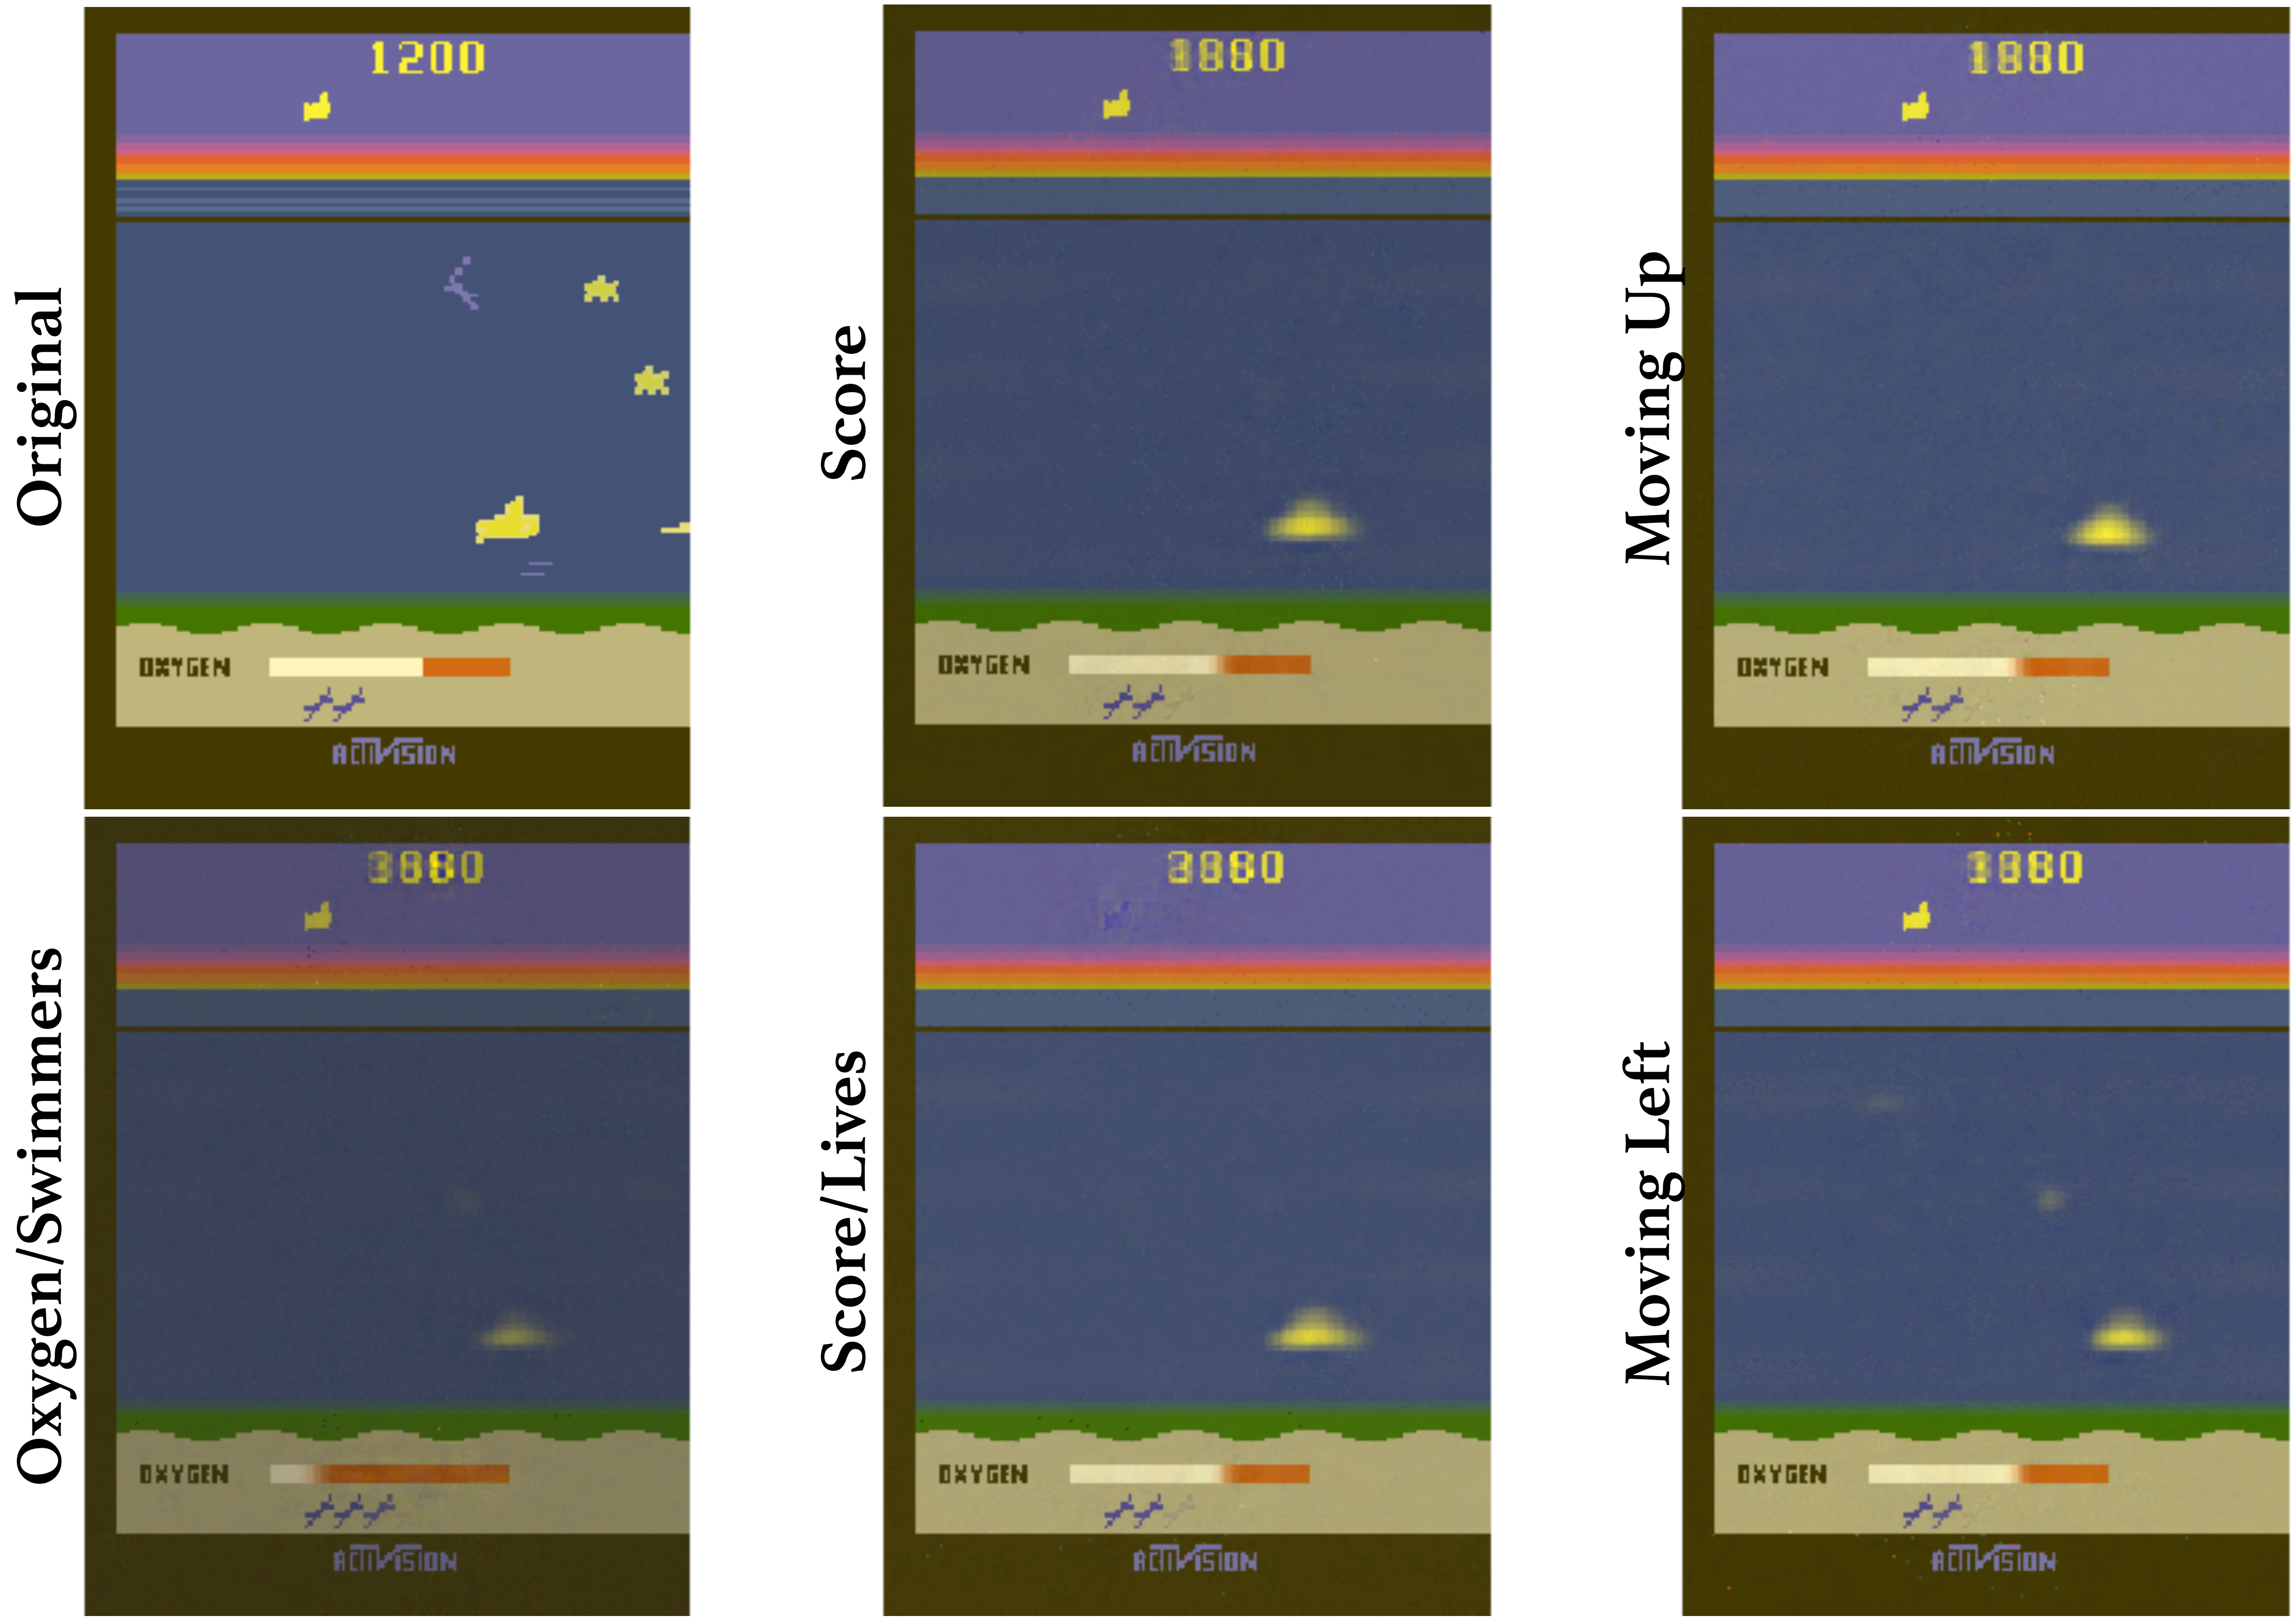

Concept Learning

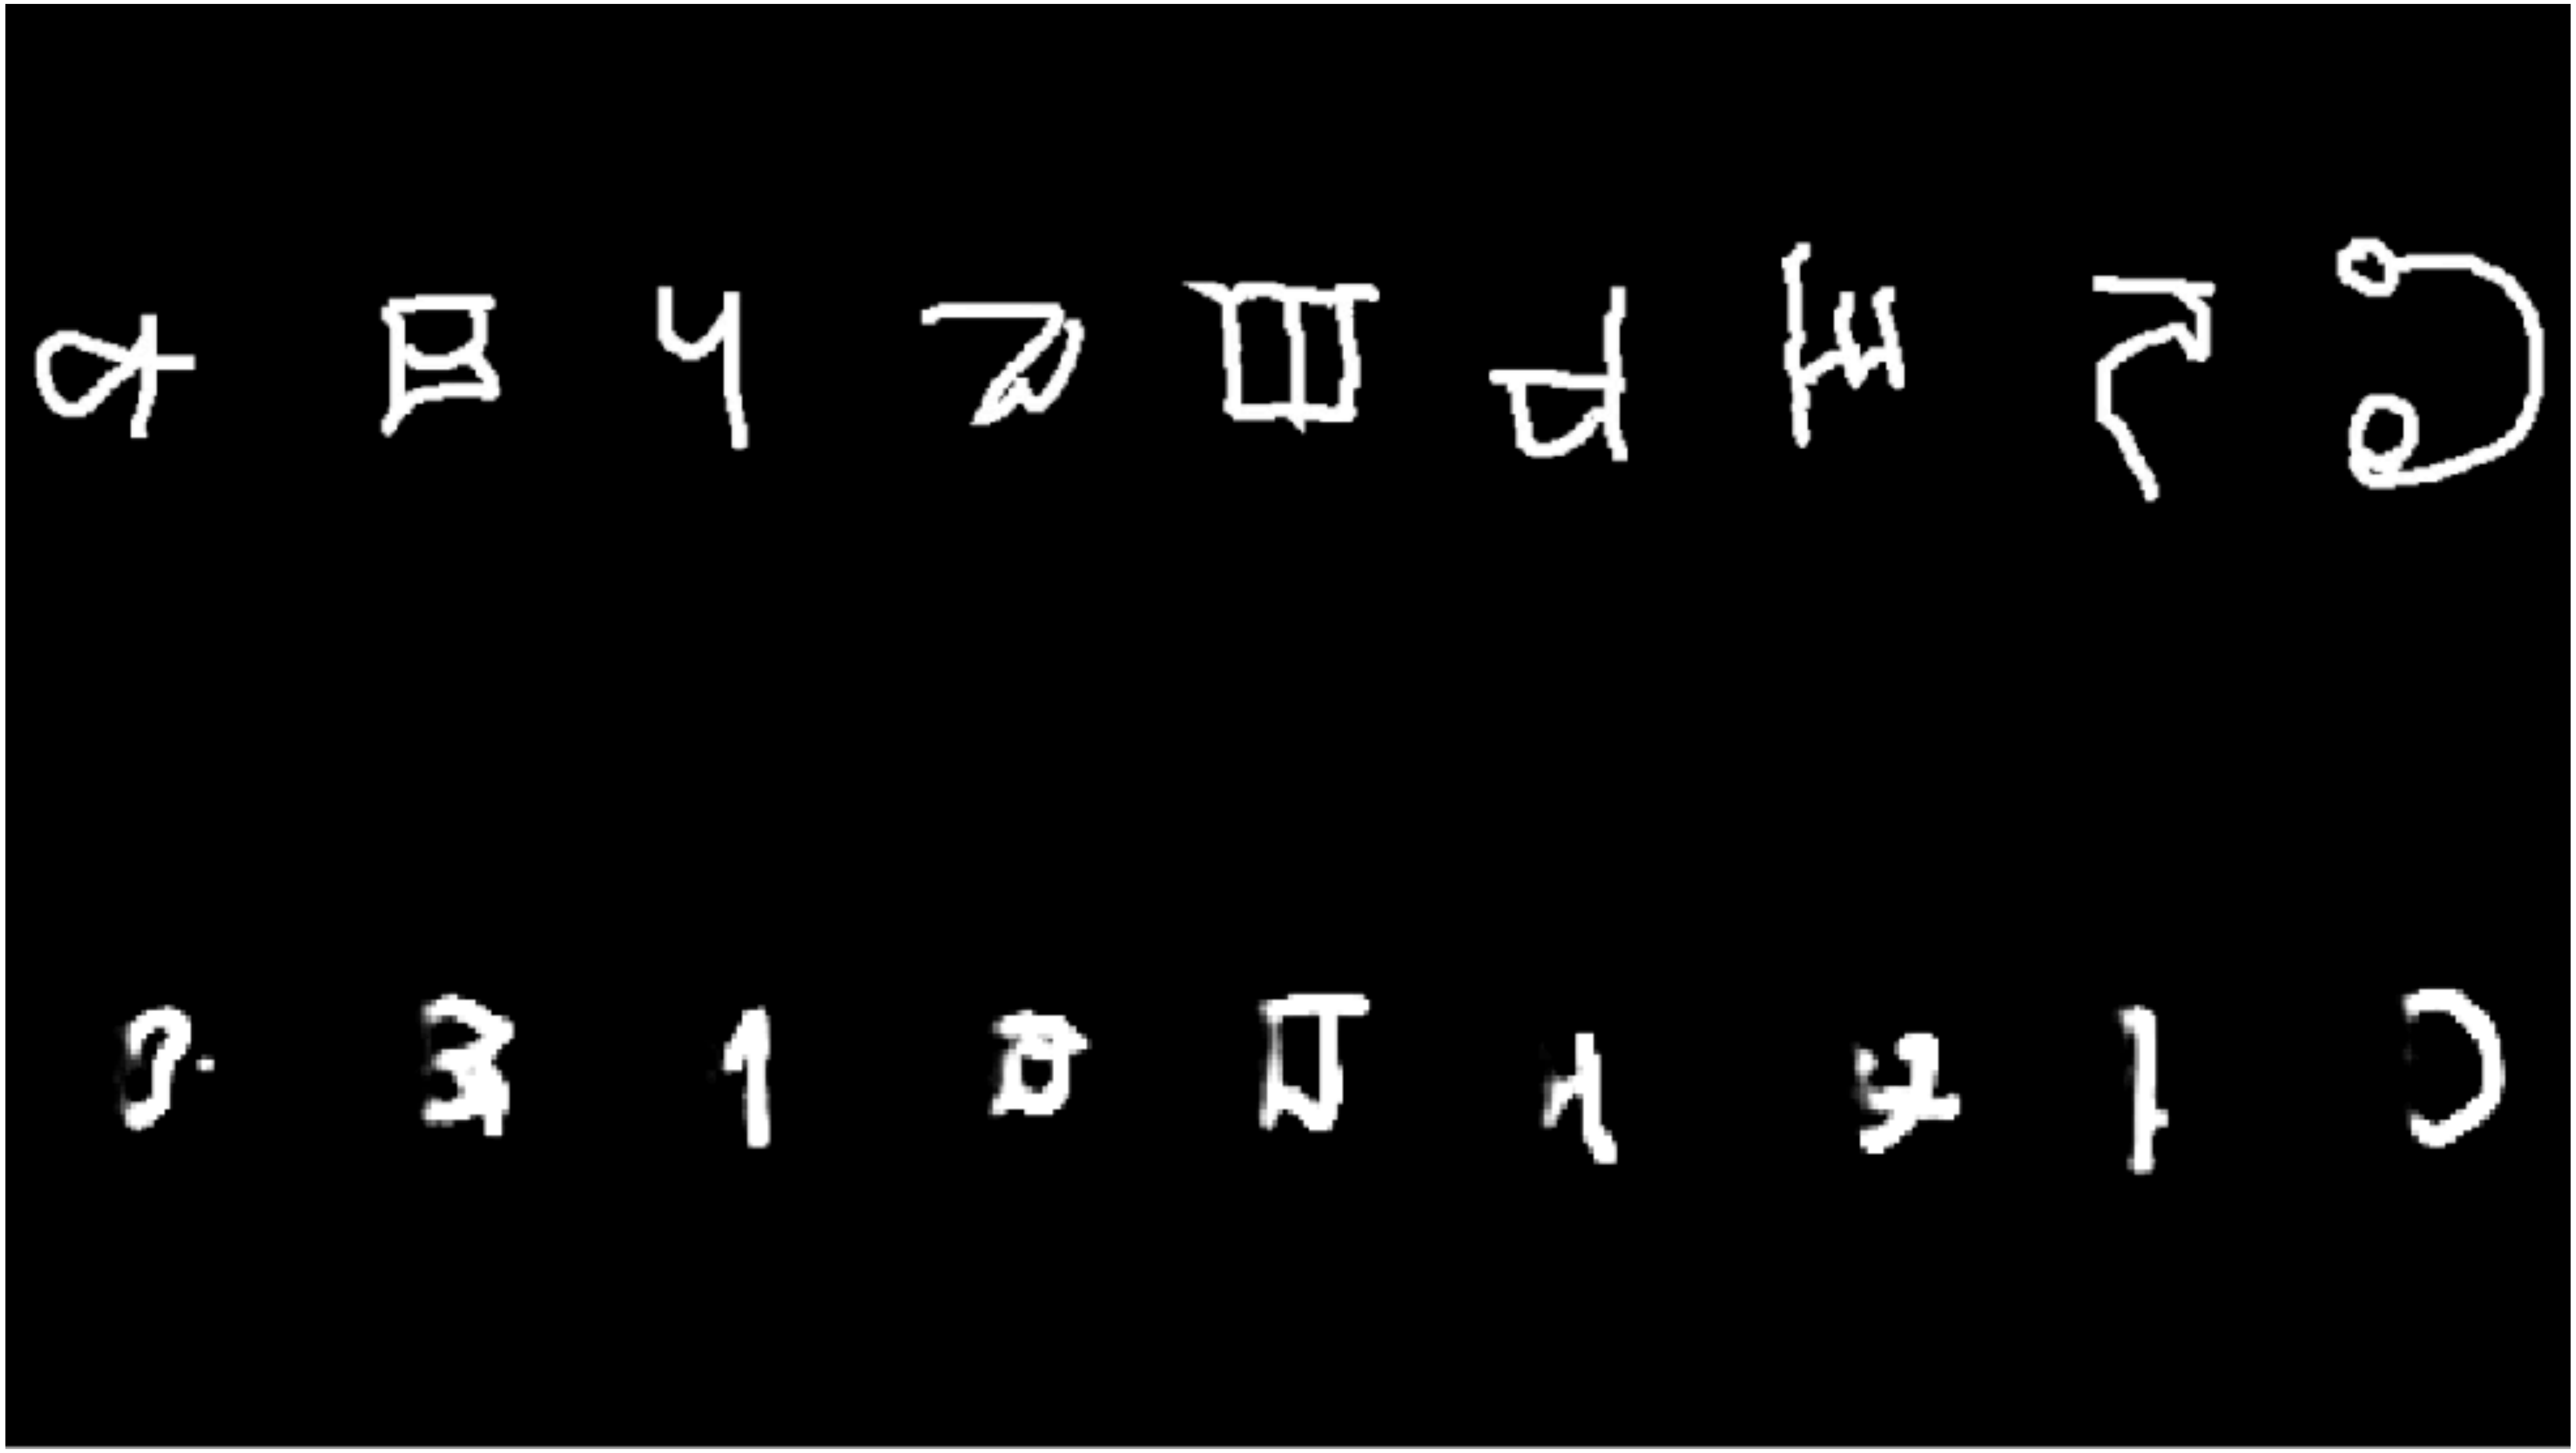

Future Thinking

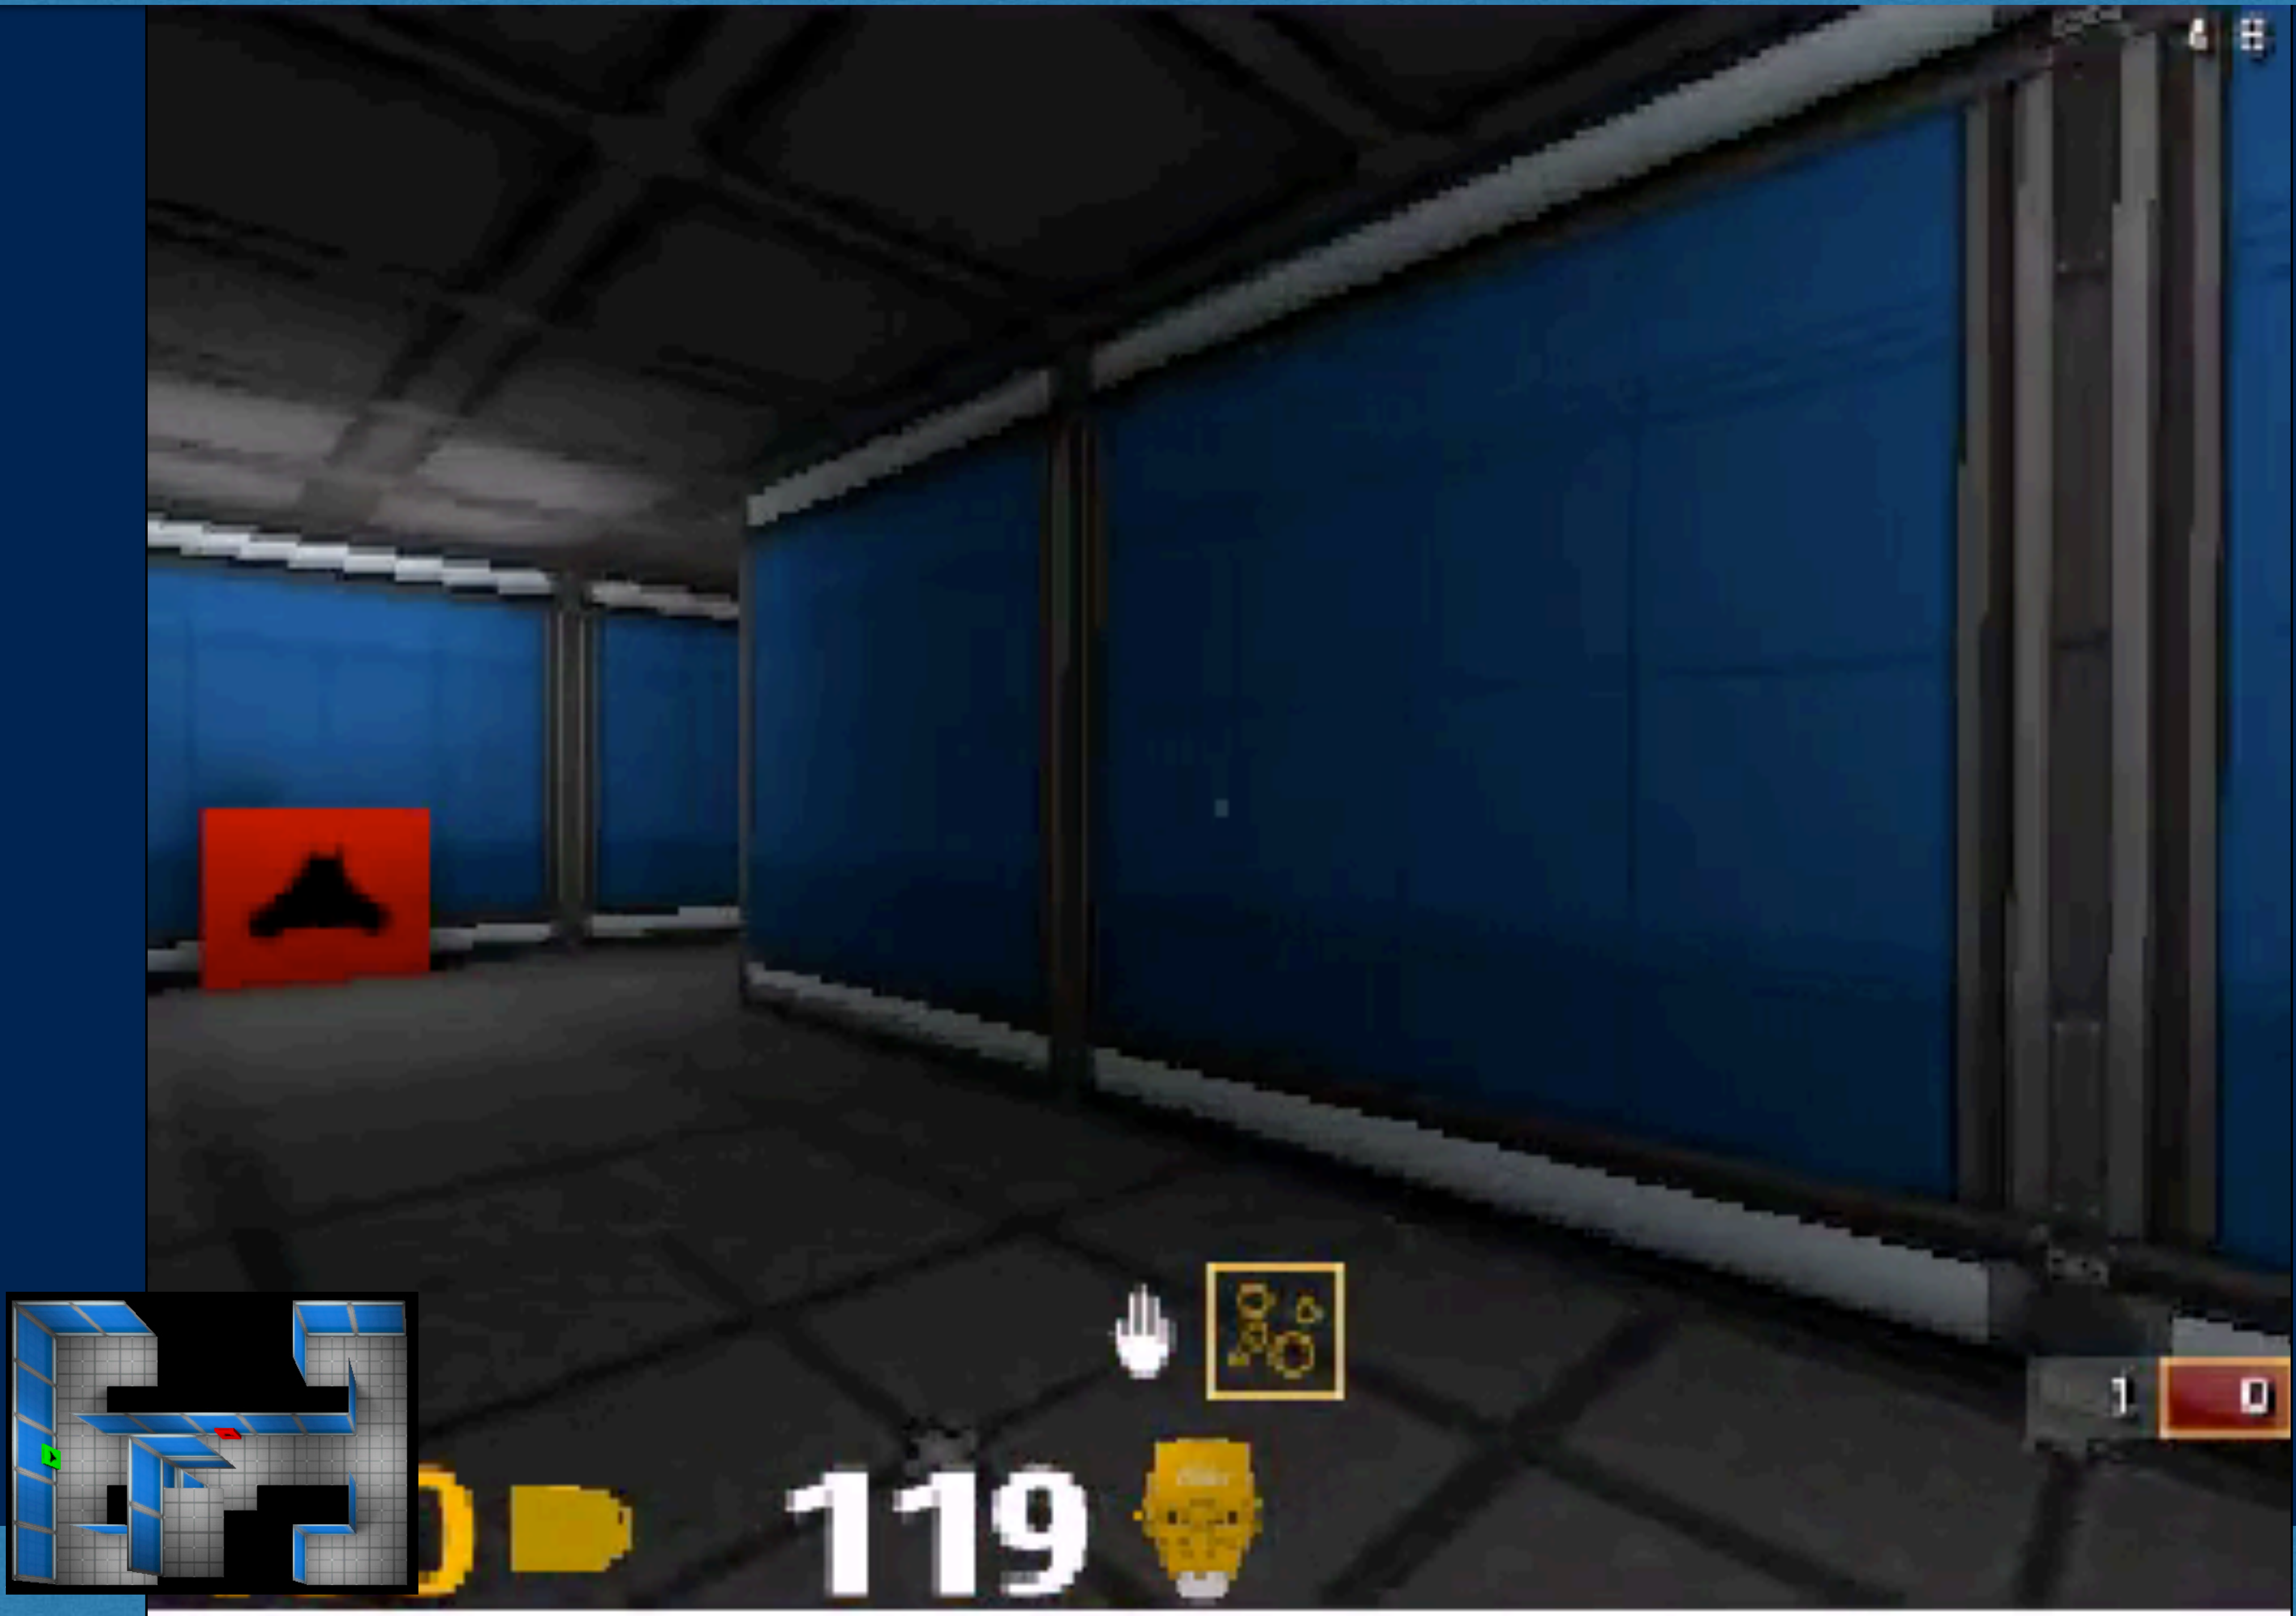

Planning and Acting

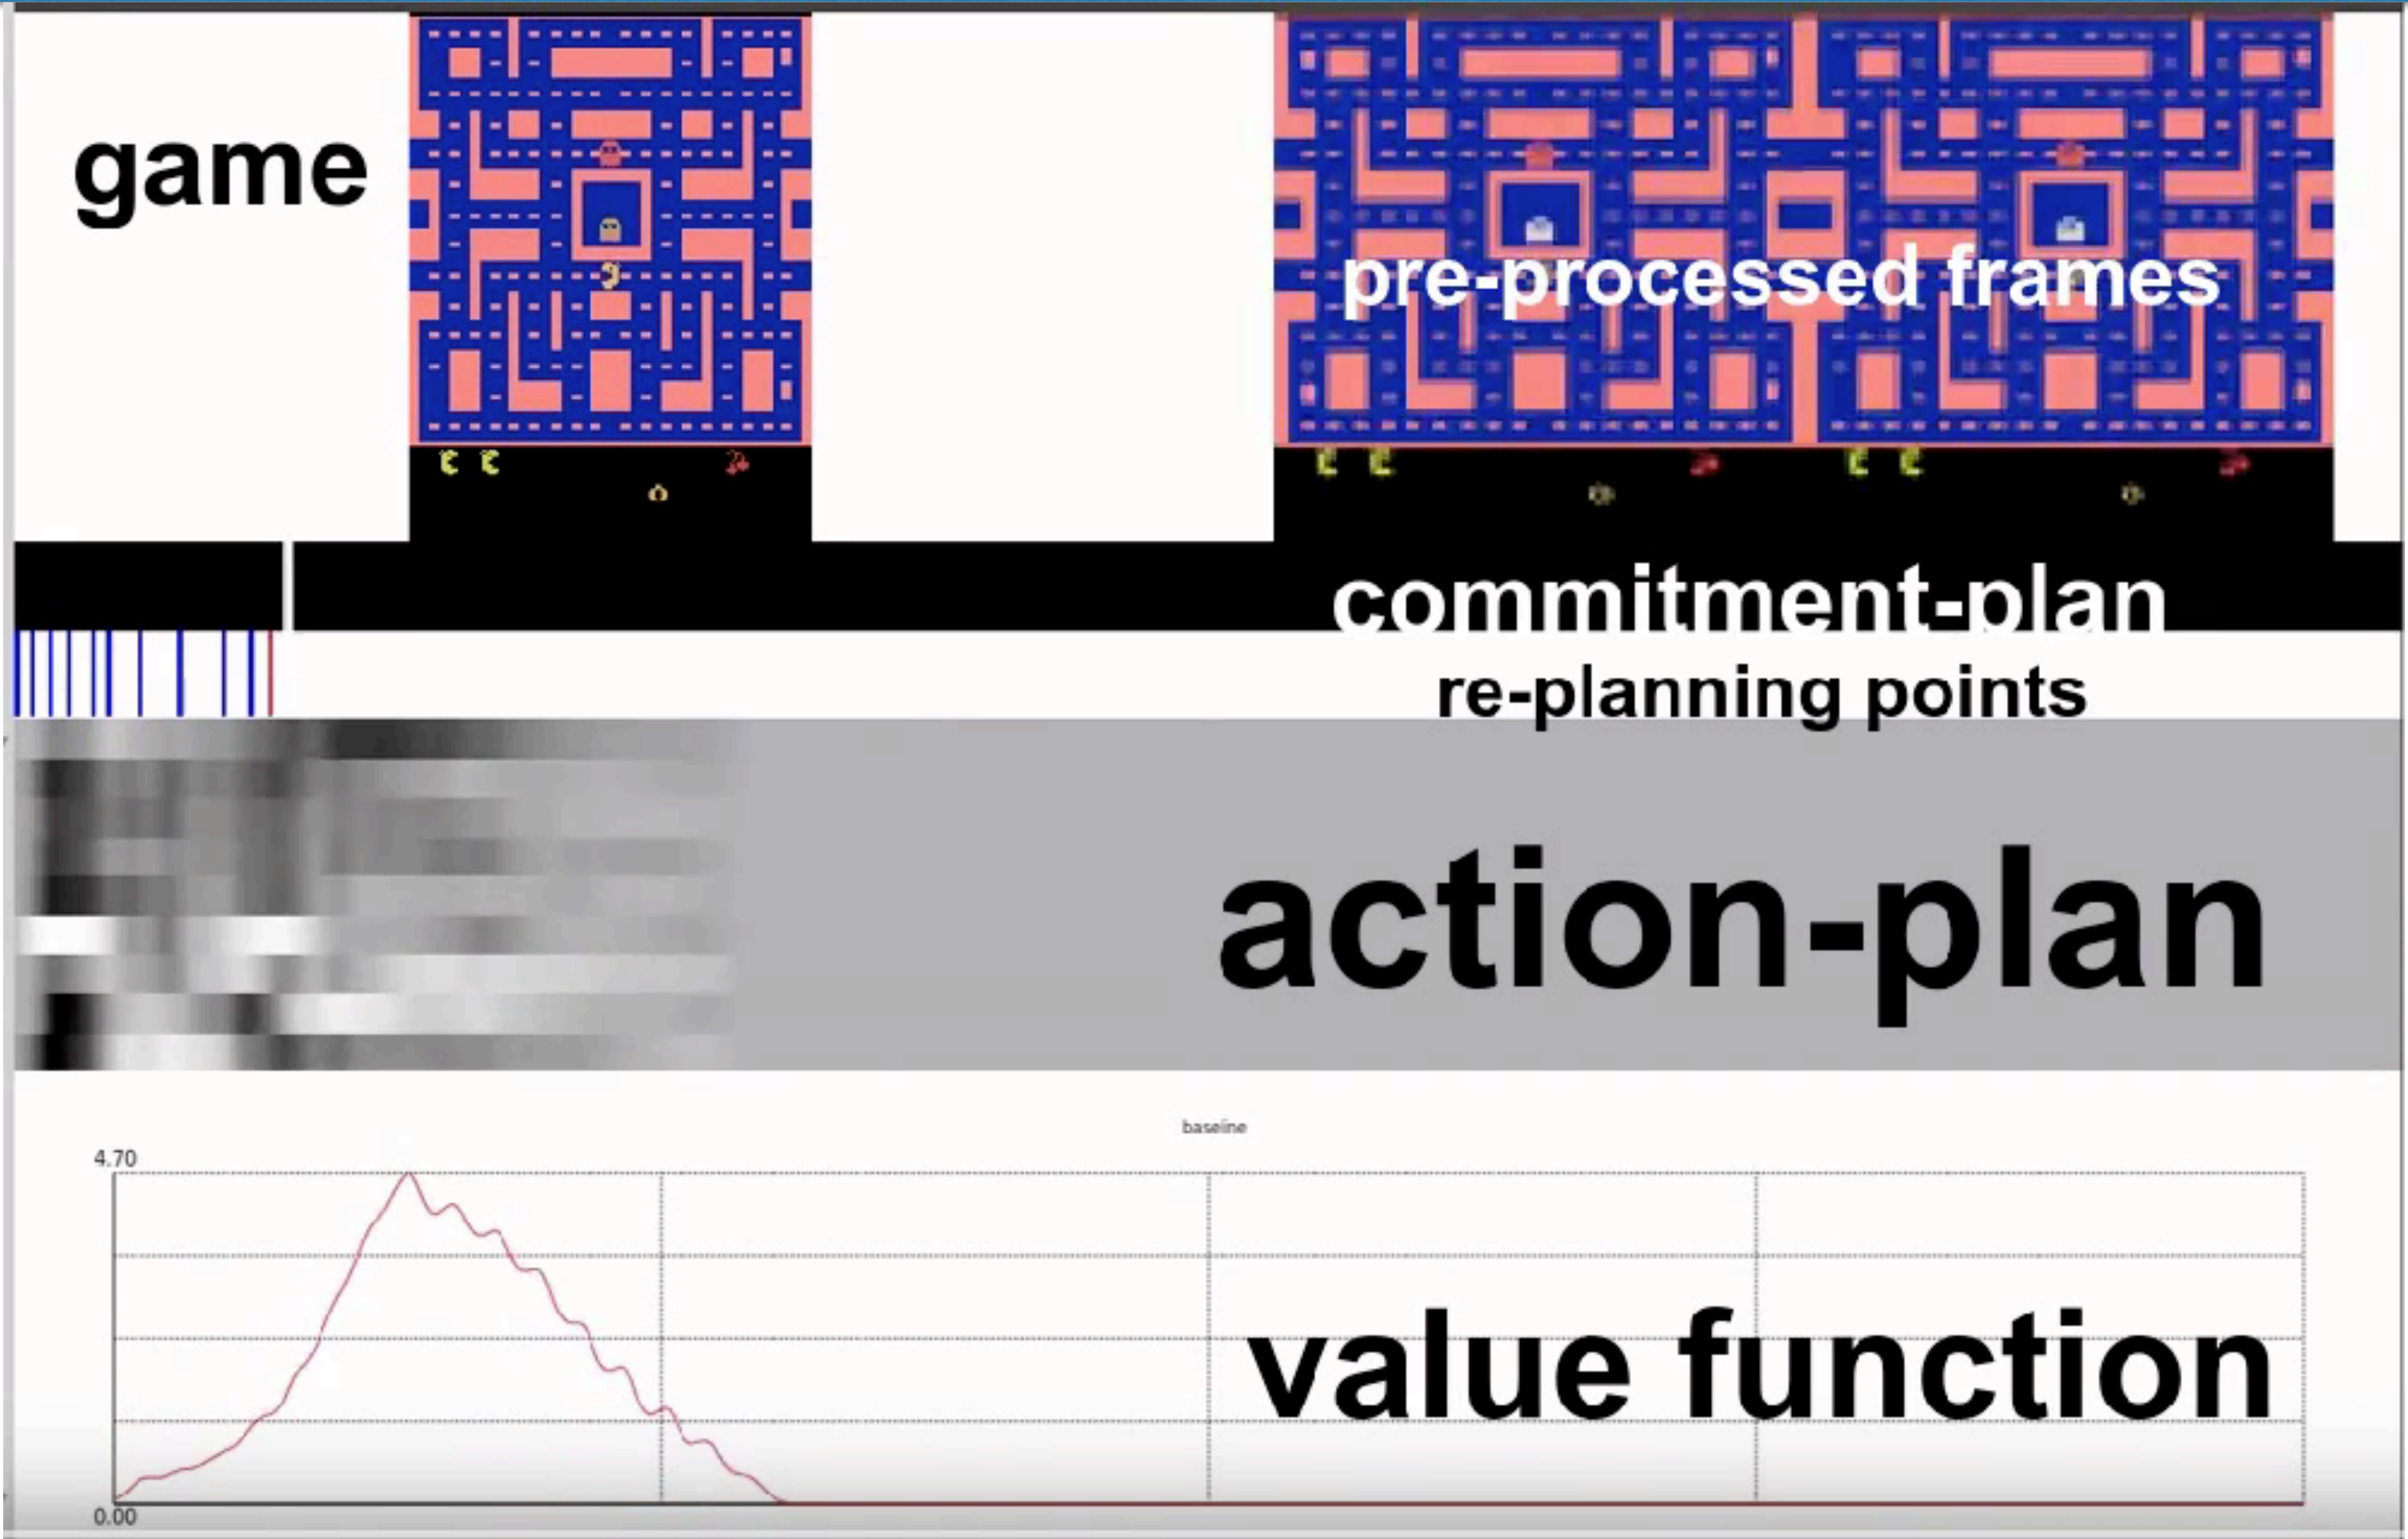

Evaluation | Integration | Memory | Discrete models | Continual learning

Bayesian Agents

Bayesian Reasoning and Deep Learning in Agent-based Systems

Thanks to many people:
Danilo Rezende, Theophane Weber, Andriy Mnih, Ali Eslami, Karol Gregor, Sasha Veznevehts, Irina Higgins, Balaji Lakshminarayanan, Lars Buesing, Daan Wierstra, and many others at DeepMind.

Shakir Mohamed

DeepMind

@shakir_za shakir@google.com