


# Non-supervised Learning & Decision Making

Danilo J. Rezende



**Hammers & Nails 2019, Weizmann Institute of Science, Israel**


**"The revolution will not be supervised"** (Yann Lecun, 2017)



# Topics


- General discussion about Reinforcement Learning
- Why build models?
- Using generative models for representations learning:
  - GQN
  - SimCore
- Some thoughts on applications of group theory in ML

# Reinforcement Learning





# Let's consider a discrete Markov Decision Process (MDP)





# RL in a discrete world

Set of States

$$\mathcal{S} = \{1, \dots, N_s\}$$

Set of Rewards

$$\mathcal{R} = \{1, \dots, N_r\}$$

Set of Actions

$$\mathcal{A} = \{1, \dots, N_a\}$$

Reward function

$$r : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$$

Model/Environment

$$m : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$$



Let's consider a discrete world MDP

$$Q(s_0, a_0) = \max_{\text{all possible paths}} \sum_{t=1}^H \gamma^t r(s_t, a_t)$$

Optimal decision

$$a^*(s_0) = \operatorname{argmax}_a Q(s_0, a)$$



# Let's consider a discrete world MDP

$$Q(s_0, a_0) = \max_{\text{all possible paths}} \sum_{t=1}^H \gamma^t r(s_t, a_t)$$

Set of States

$$\mathcal{S} = \{1, \dots, N_s\}$$

Set of Rewards

$$\mathcal{R} = \{1, \dots, N_r\}$$

Set of Actions

$$\mathcal{A} = \{1, \dots, N_a\}$$

Reward function

$$r : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$$

Model/Environment

$$m : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$$

Q-function

$$Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}^H$$



# How much data do we need to learn $r$ , $m$ and $Q$ ?

Amount of Data to learn  $f \approx e^{\text{Description Length of } f}$

$$\text{DL}(r) = N_s N_a \log_2(N_r)$$

$$\text{DL}(m) = N_s N_a \log_2(N_s)$$

$$\text{DL}(Q) = N_s N_a H \log_2(N_r)$$




# When is it worth learning a model?

$$\text{DL}(r + m) = N_s N_a \log_2(N_r) + N_s N_a \log_2(N_s)$$

$$\text{DL}(Q) = N_s N_a H \log_2(N_r)$$

**Hypothesis: It will be favorable to learn a model when**





# When is it worth learning a model?

$$\text{DL}(r + m) = N_s N_a \log_2(N_r) + N_s N_a \log_2(N_s)$$

$$\text{DL}(Q) = N_s N_a H \log_2(N_r)$$

**Hypothesis: It will be favorable to learn a model when**

$$\text{DL}(Q) > \text{DL}(r + m)$$

$$(H - 1) \log_2(N_r) > \log_2(N_s)$$



# When is it worth learning a model?

$$(H - 1) \log_2(N_r) > \log_2(N_s)$$

In general, **we operate in a regime where  $N_s \gg N_r, H$** . So it seems that we would always prefer to learn  $Q$  from scratch rather than learning a model first.



When is it worth learning a model?  
Many tasks, a single world

$$L(H - 1) \log_2(N_r) > \log_2(N_s)$$

**For a sufficiently large number of different tasks  $L$ , it will require less data if we learn a model first.**




# Summary

- Learning a model first and using that to compute Q-functions via Monte-Carlo can be more data-efficient in some cases, specially in multi-task problems
- Models offer other advantages in addition to planning:
  - Notion of uncertainty or novelty
  - Unsupervised learning of useful features
  - Unsupervised learning to use complex memory architectures




# RL in a stochastic and partially observed world





# RL in a stochastic and partially observed world





# RL in a stochastic world

**Policy**

$$\pi(a_t | x_0, \dots, x_t, a_0, \dots, a_{t-1})$$

**Model/Simulator**

$$m(x_t | x_0, \dots, x_{t-1}, a_0, \dots, a_{t-1})$$

**Path/trajectory**

$$\tau = \{(x_1, a_1), \dots, (x_t, a_t), \dots, (x_H, a_H)\}$$

$$Q^\pi(s_0, a_0) = \mathbb{E}_{p(\tau)} \left[ \sum_t \gamma^t r(s_t, a_t) \mid s_0, \text{do}(a_0) \right]$$



# RL in a stochastic world

**Policy**

$$\pi(a_t | x_0, \dots, x_t, a_0, \dots, a_{t-1})$$

**Model/Simulator**

$$m(x_t | x_0, \dots, x_{t-1}, a_0, \dots, a_{t-1})$$

**Path/trajectory**

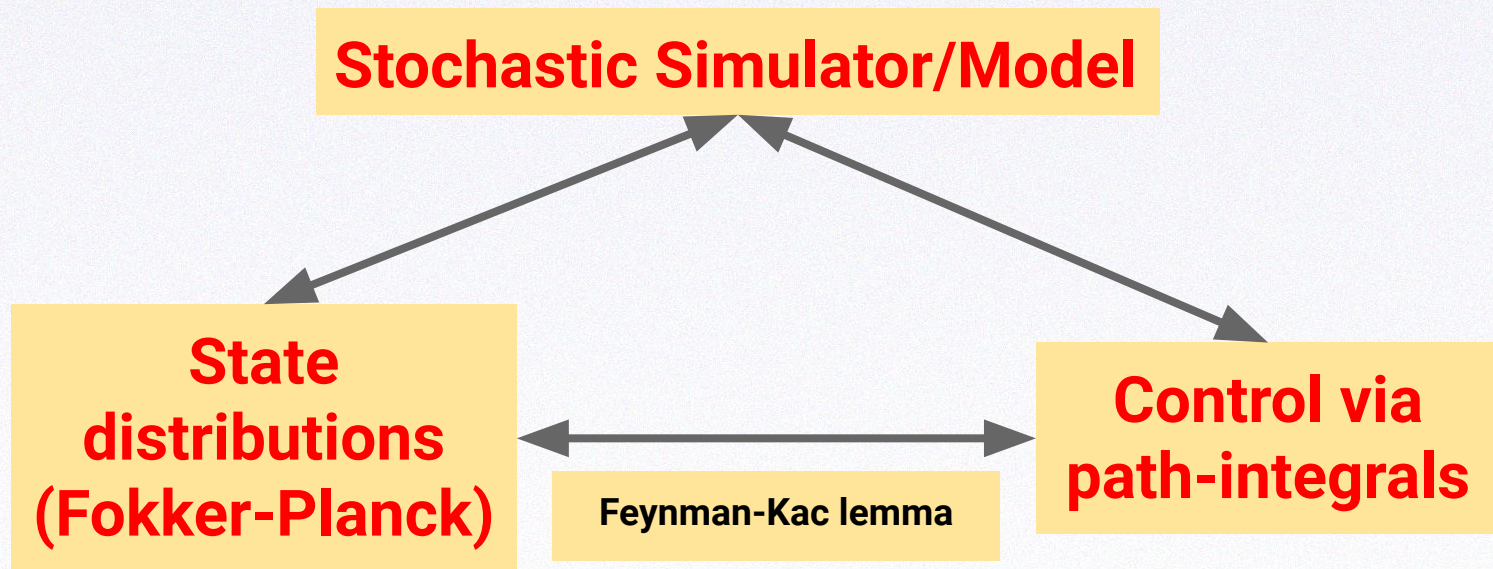
$$\tau = \{(x_1, a_1), \dots, (x_t, a_t), \dots, (x_H, a_H)\}$$

$$Q^\pi(s_0, a_0) = \mathbb{E}_{p(\tau)} \left[ \sum_t \gamma^t r(s_t, a_t) \mid s_0, \text{do}(a_0) \right]$$

**Path Integral**



# RL in a stochastic world



[Path Integral Formulation of Stochastic Optimal Control with Generalized Costs](#) Yang et al  
[A Generalized Path Integral Control Approach to Reinforcement Learning](#) Theodorou et al

## RL in a stochastic world

$$Q^{\pi}(s_0, a_0) = \mathbb{E}_{p(\tau)} \left[ \sum_t \gamma^t r(s_t, a_t) \middle| s_0, \text{do}(a_0) \right]$$

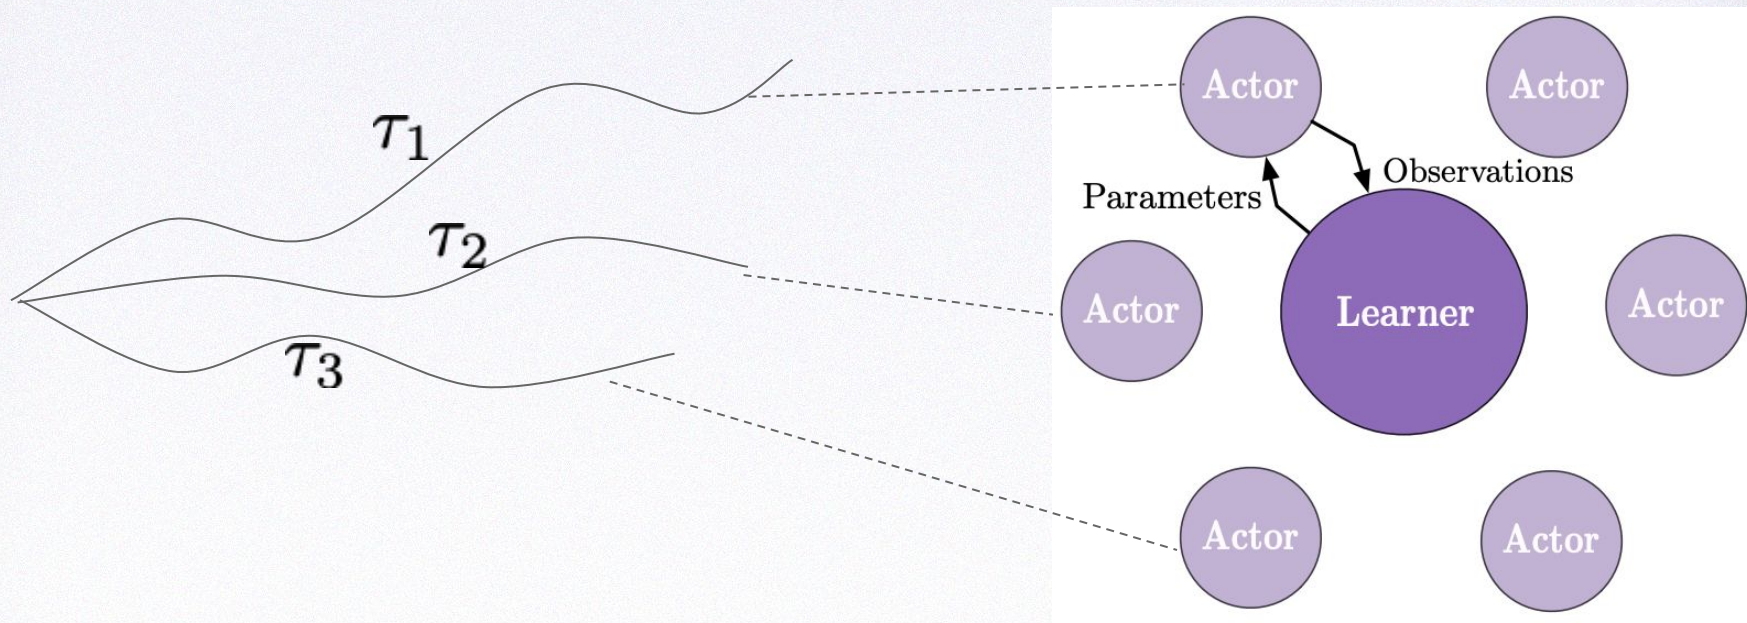


## RL in a stochastic world

$$Q^{\pi}(s_0, a_0) = \mathbb{E}_{p(\tau)} \left[ \sum_t \gamma^t r(s_t, a_t) \middle| s_0, \text{do}(a_0) \right]$$

$$Q^{\pi}(s_0, a_0) \approx \frac{1}{N} \sum_k \sum_t \gamma^t r(s_t^k, a_t^k)$$

# RL in a complex world



[IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures](#)

Lasse Espeholt, Hubert Soyer



## RL in a stochastic world

$$Q^\pi(s_0, a_0) = \mathbb{E}_{p(\tau)} \left[ \sum_t \gamma^t r(s_t, a_t) \middle| s_0, \text{do}(a_0) \right]$$

$$Q^\pi(s_0, a_0) \approx \frac{1}{N} \sum_k \sum_t \gamma^t r(s_t^k, a_t^k)$$

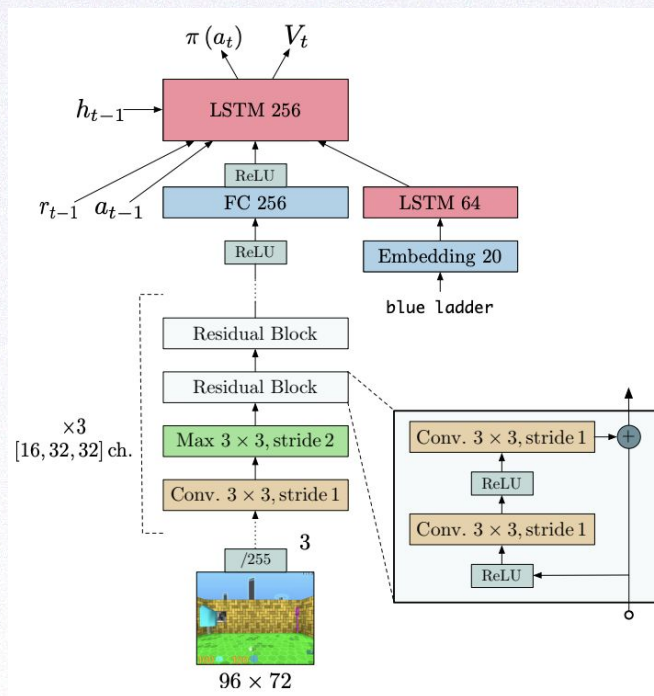
$$Q^\pi(s_0, a_0) \approx \frac{1}{N} \sum_k \sum_t \gamma^t r(s_t^k, a_t^k) \frac{\pi_{\text{current}}(a_t^k)}{\pi_{\text{old}}(a_t^k)}$$

[IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures](#)

Lasse Espeholt, Hubert Soyer



# "Canonical" Agent



[IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures](#)

Lasse Espeholt, Hubert Soyer

# Models and Reinforcement Learning



DeepMind



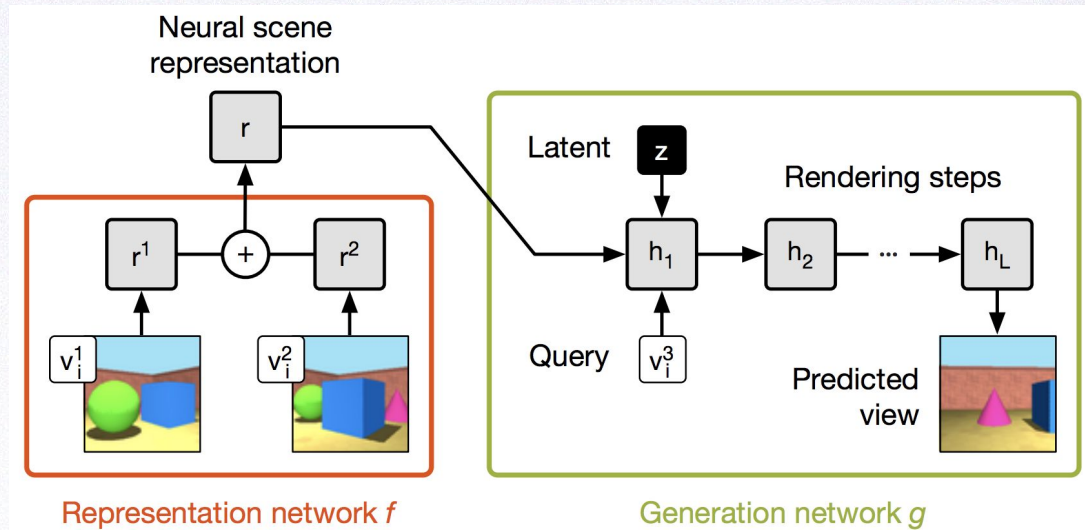

# Neural Scene Representation and Rendering

S. M. Ali Eslami, Danilo J. Rezende, et al. Science, 2018



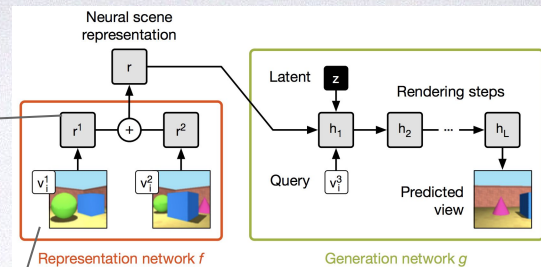
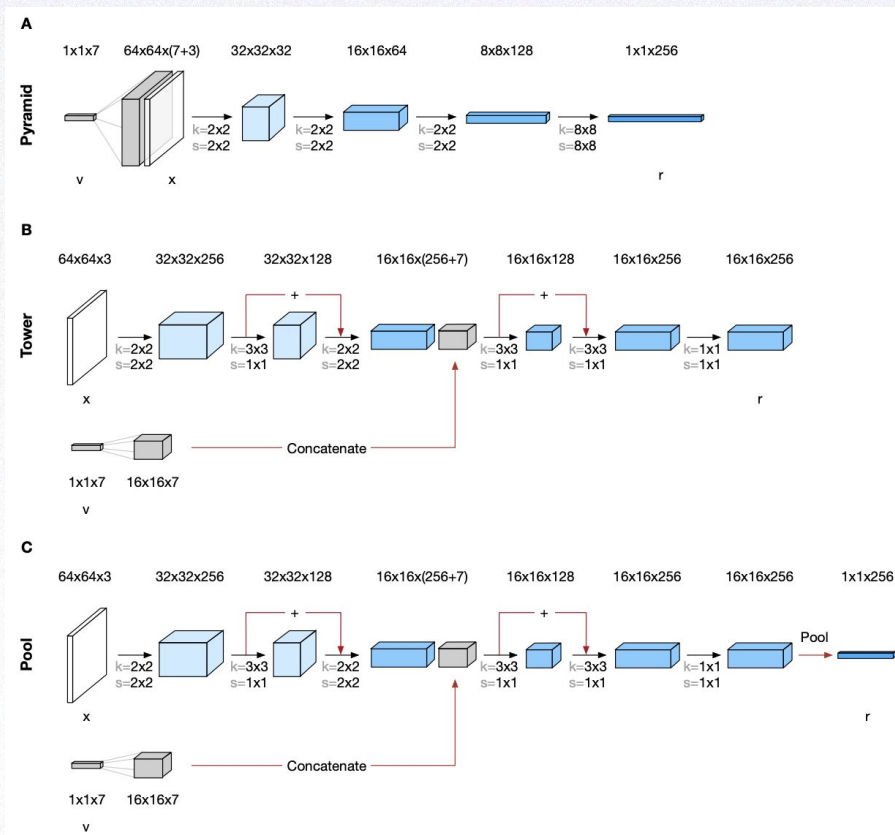
DeepMind

# The Model

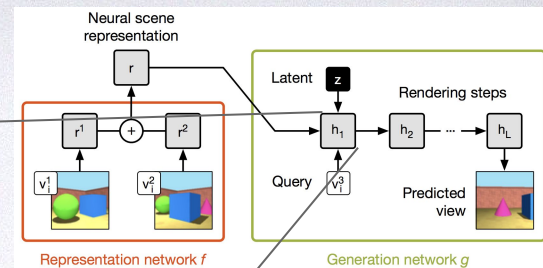
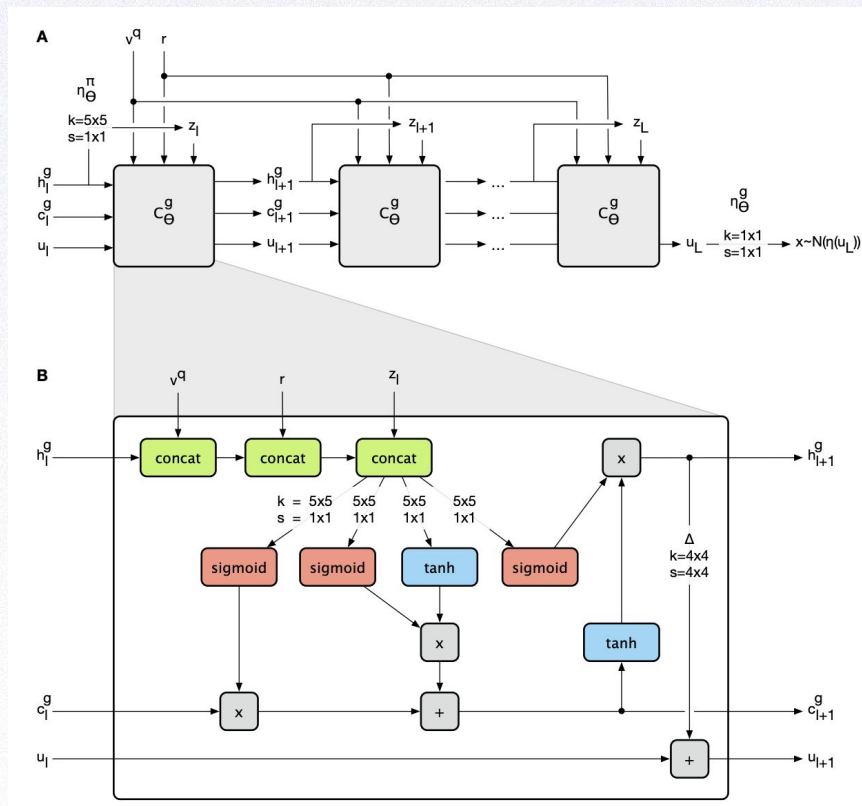




# The Model




# The Model







# The Model




# Near-Deterministic Predictions






# Data fusion and predicting with uncertainty



observation




neural rendering



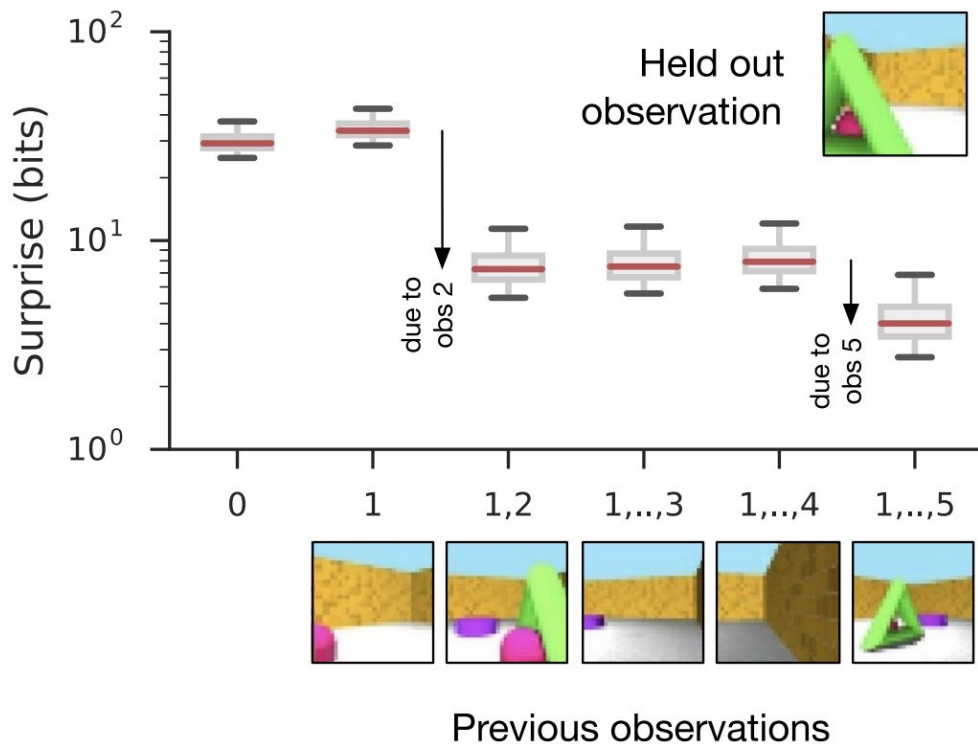
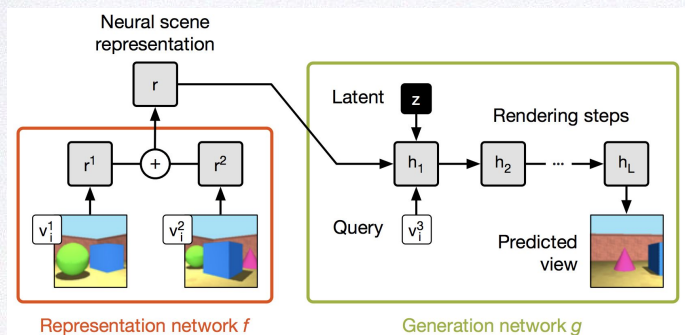
observation



neural rendering





# Quantifying uncertainty

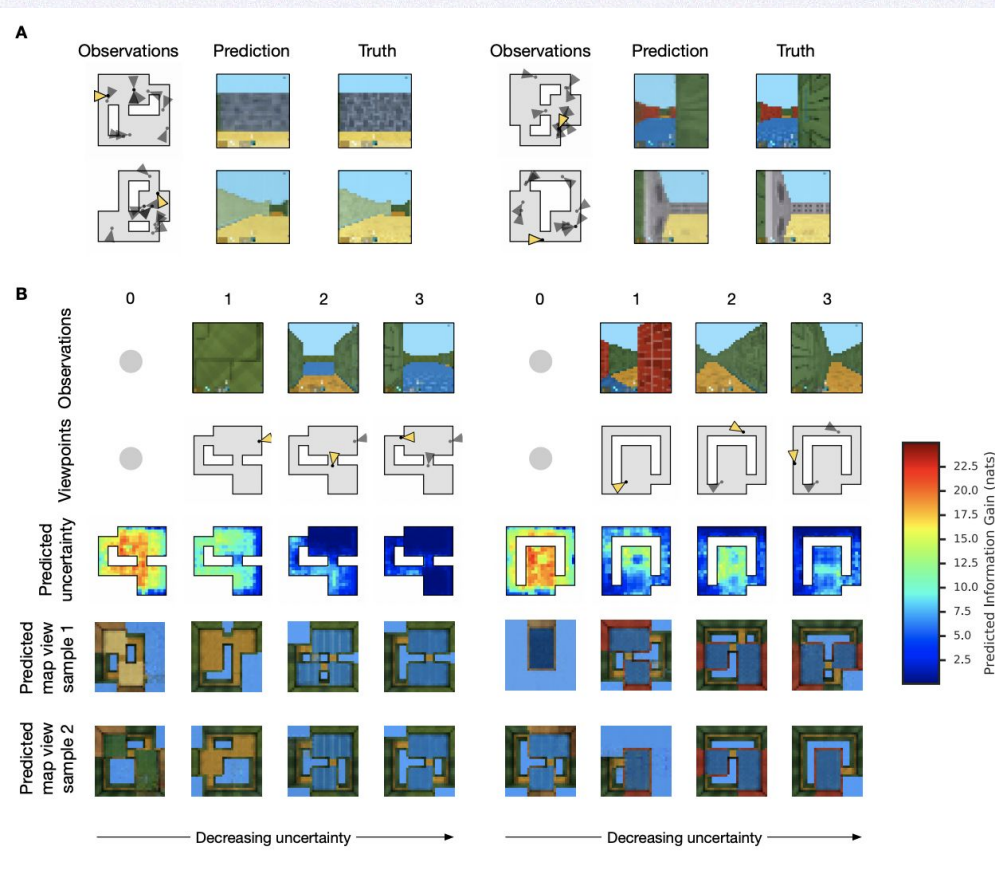




# Quantifying uncertainty




# Quantifying uncertainty






# Viewpoint invariance of the learned scene representations




# Factorization of the learned scene representations









# Factorization of the learned scene representations




# Learning useful representations for control





# Deep-Learning with Uncertainty



Slide credit: Marta Garnelo

“Conditional Neural Processes”. Marta Garnelo, Dan Rosenbaum, et al. ICML, 2018.

# TL;DL

- GQN can learn factored scene representations
- It is also a kind of meta-learning model (learning to do one-shot scene inference)
- It doesn't have a notion of time
- It requires knowledge about "camera locations"



# Shaping Belief States with Generative Environment Models for RL

Karol Gregor, Danilo Rezende, Frederic Besse,  
Yan Wu, Hamza Merzic, Aaron van den Oord,  
<https://arxiv.org/abs/1906.09237>

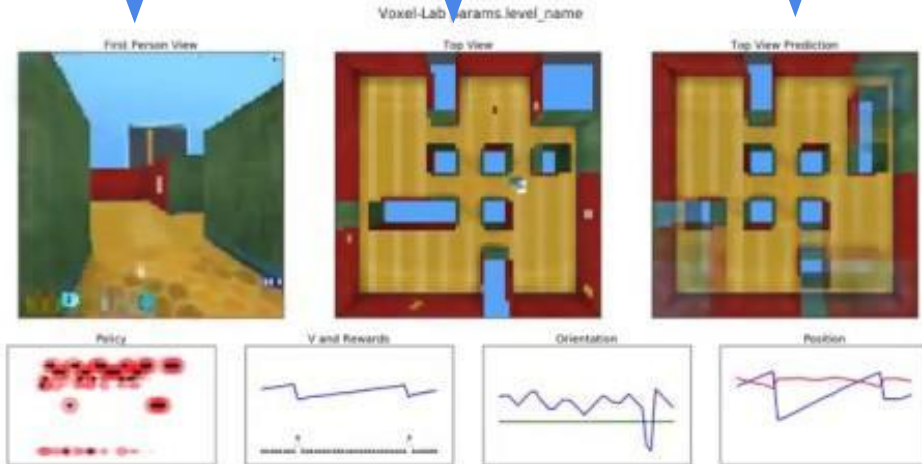



# Example (actual result)

What agent sees


Actual top down view

Top down view reconstructed from agent rep.

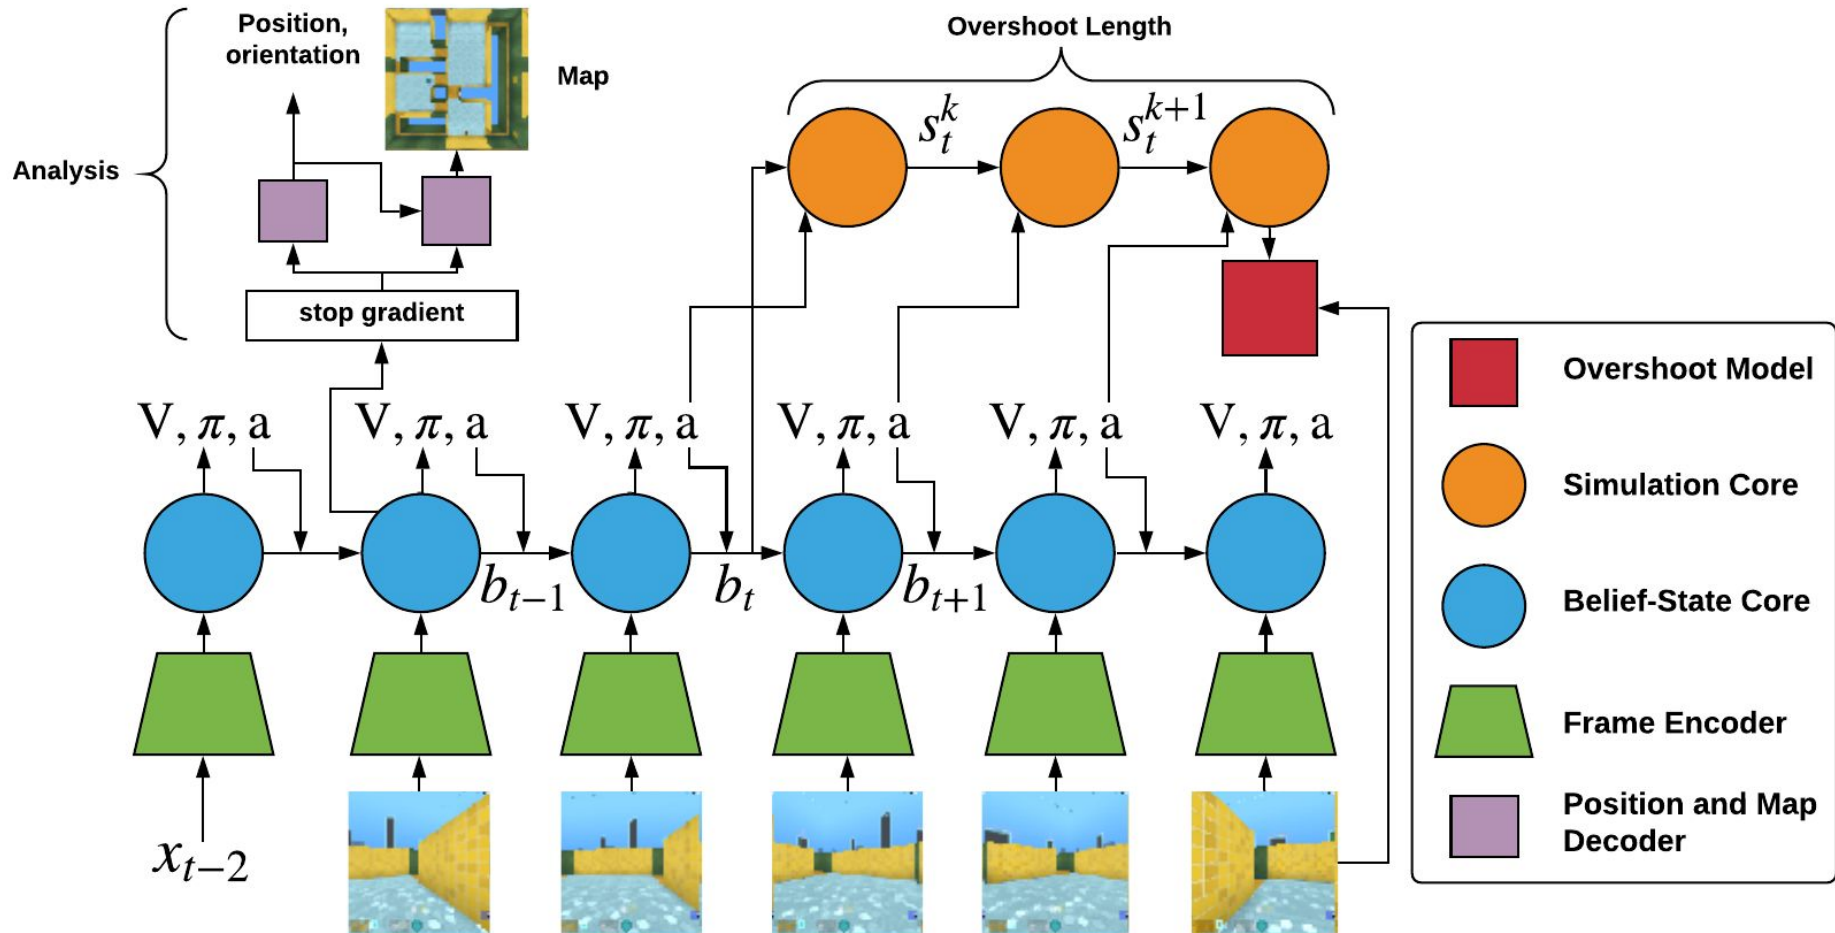


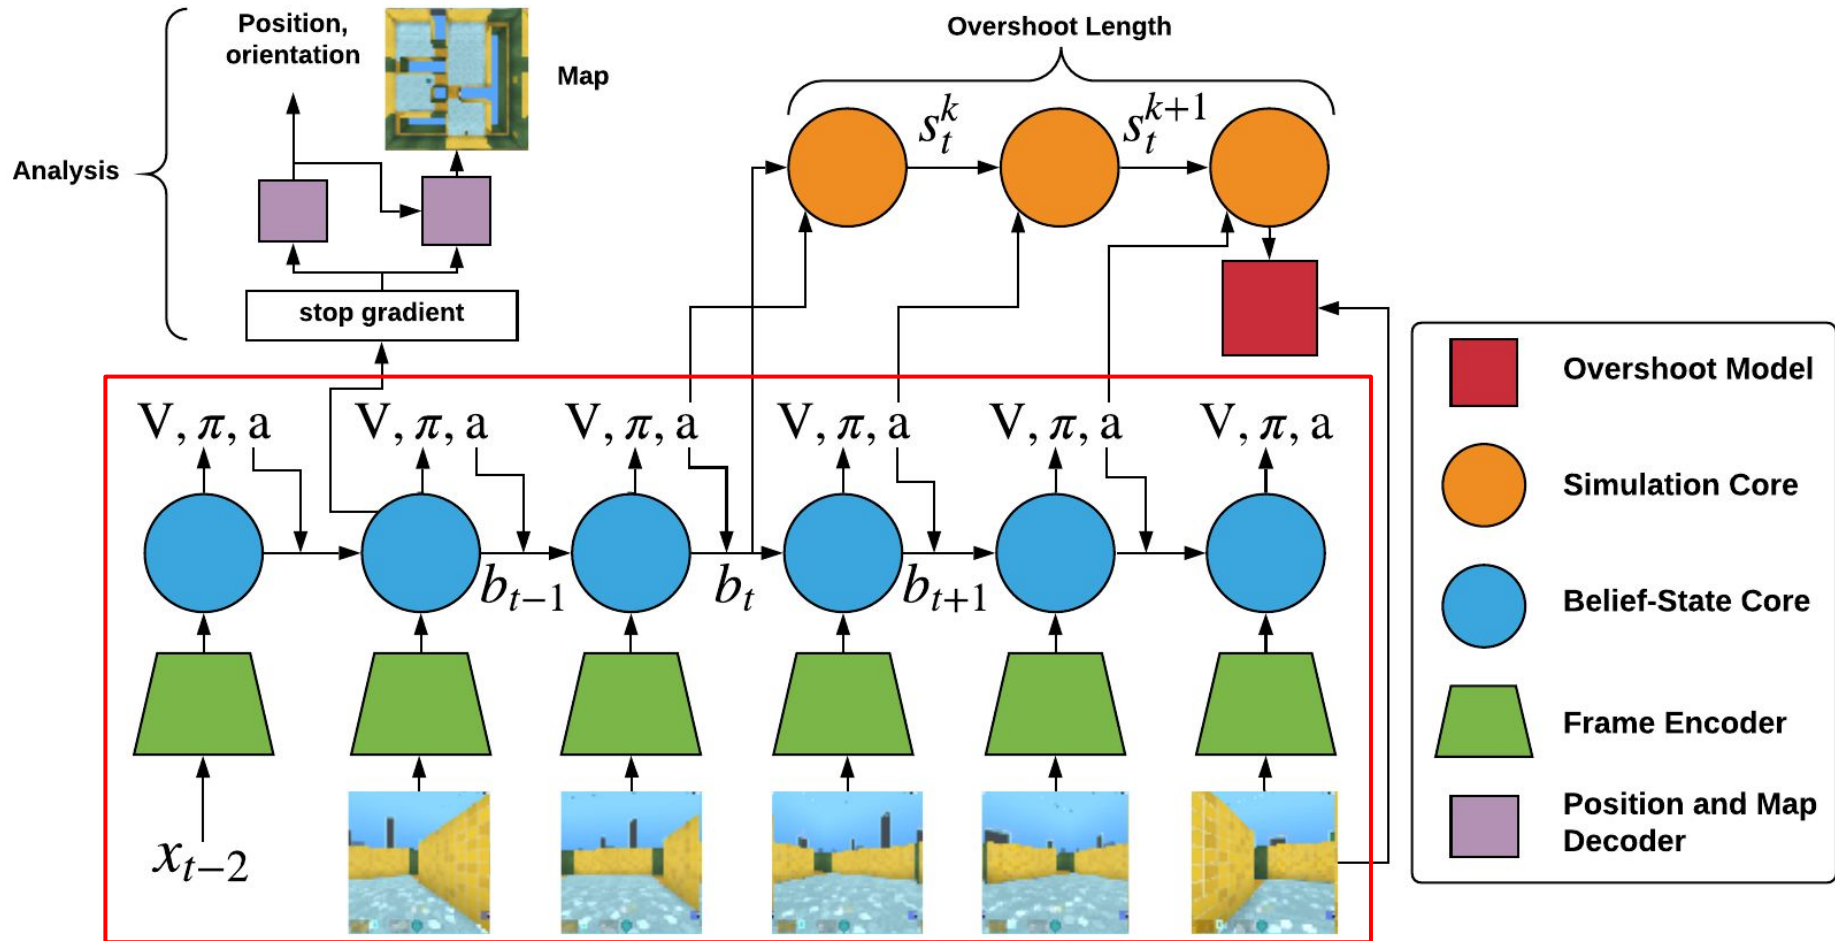
Without belief state training

With belief state training

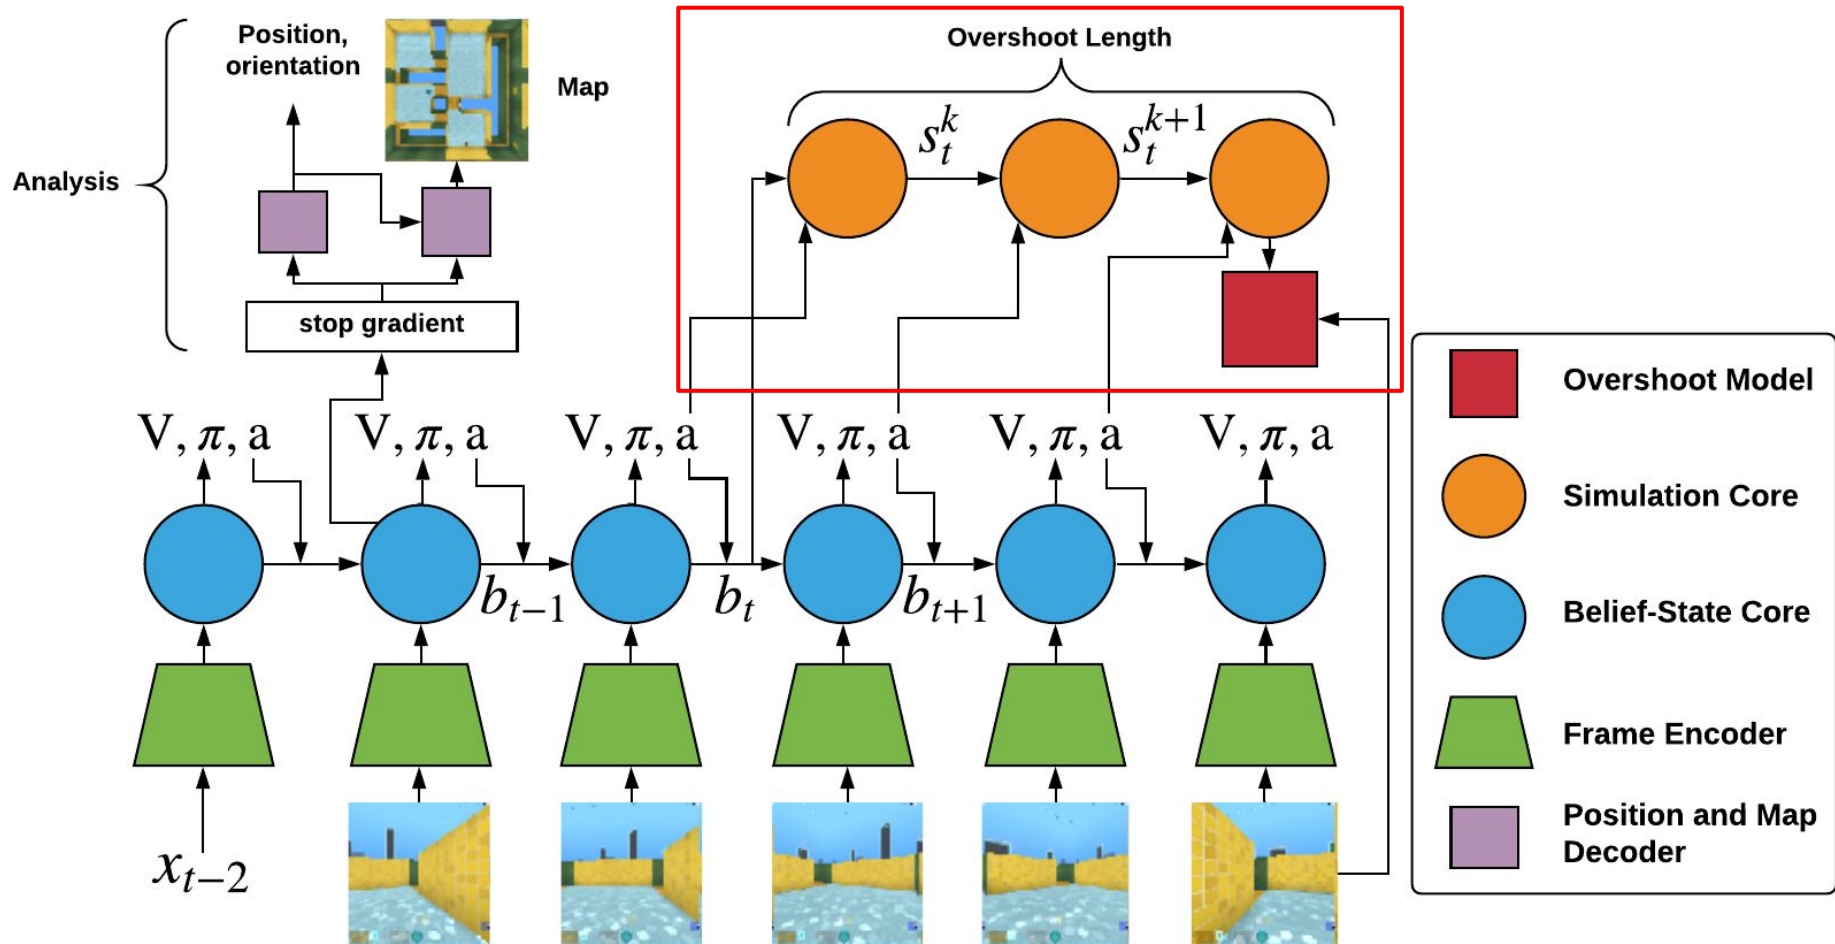


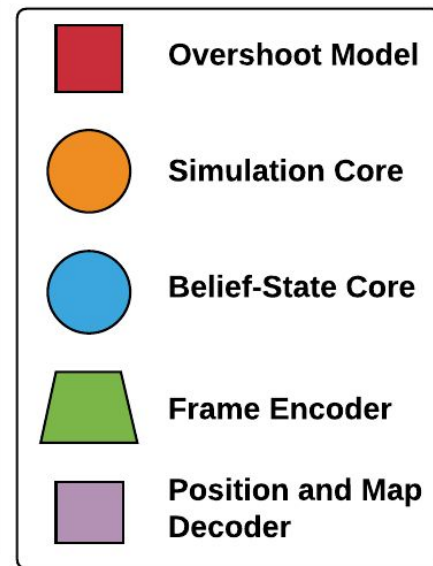
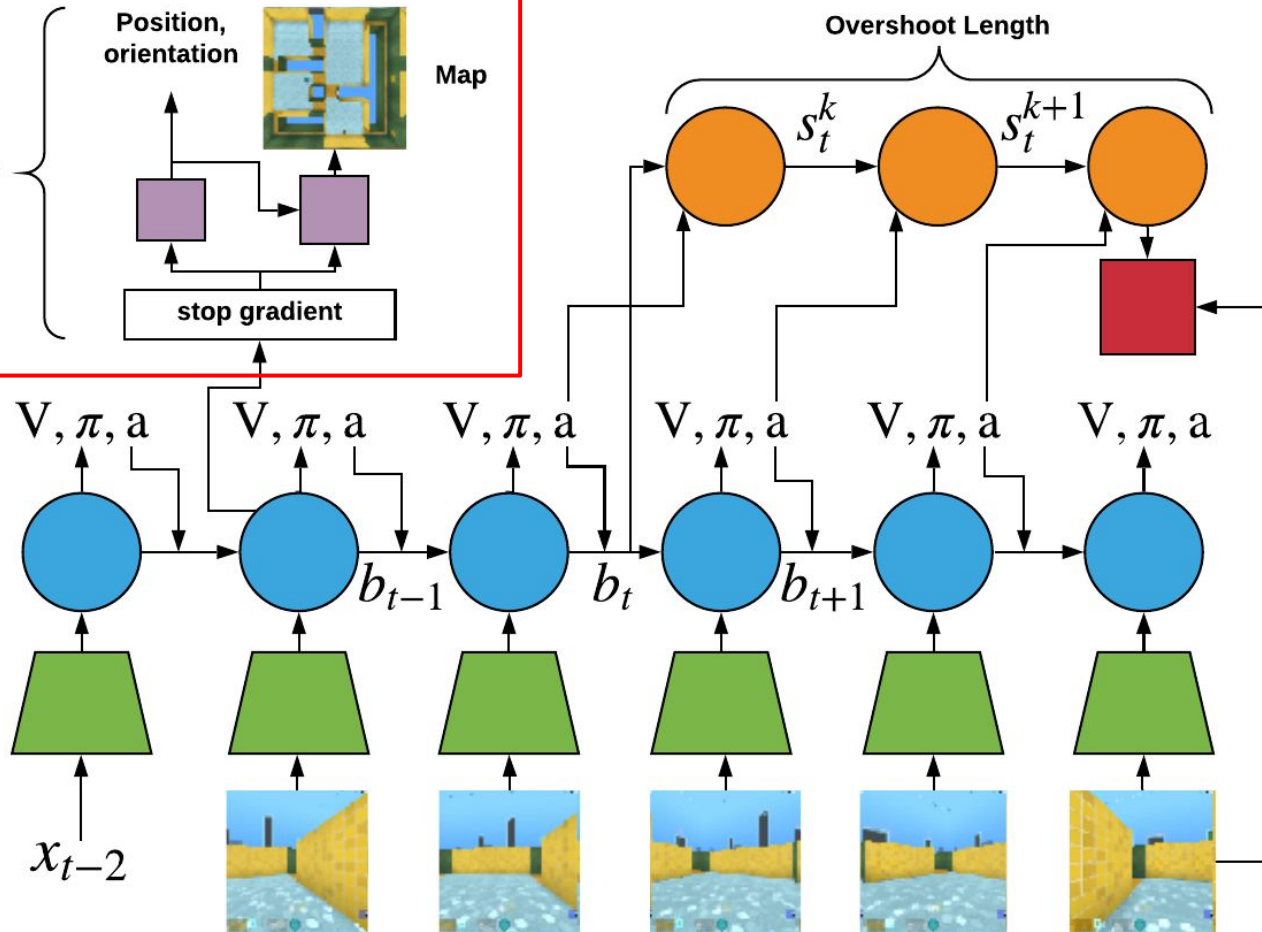


















# Model Rollouts





# Formation of belief states






# Top down view reconstruction error (mapping)

- 1) Longer simulations improve map decoding performance (if proper generative model is used)
- 2) Generative model works better than deterministic decoder and CPC
- 3) Kanerva Machine episodic memory works the best



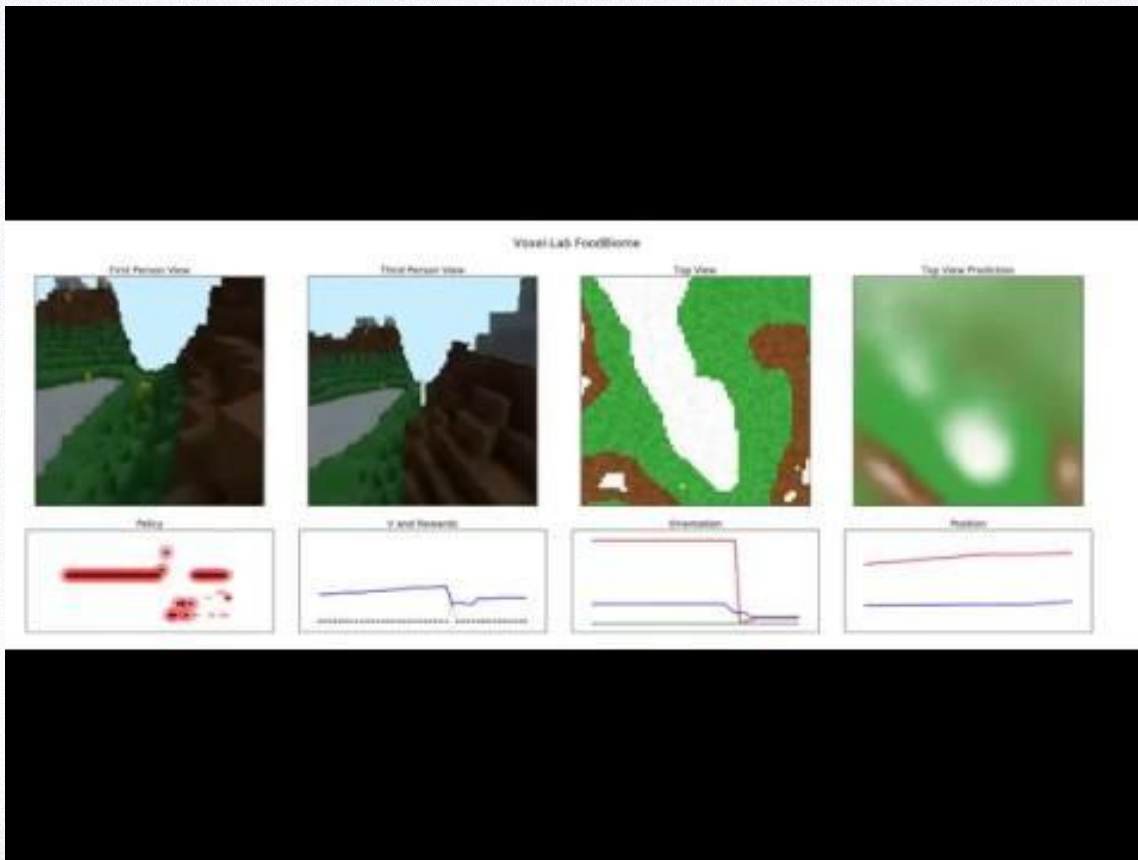
# Position and orientation prediction (localization)


- 1) Longer simulations improve localization
- 2) CPC works the best





# More complex naturalistic environment









# RL Performance

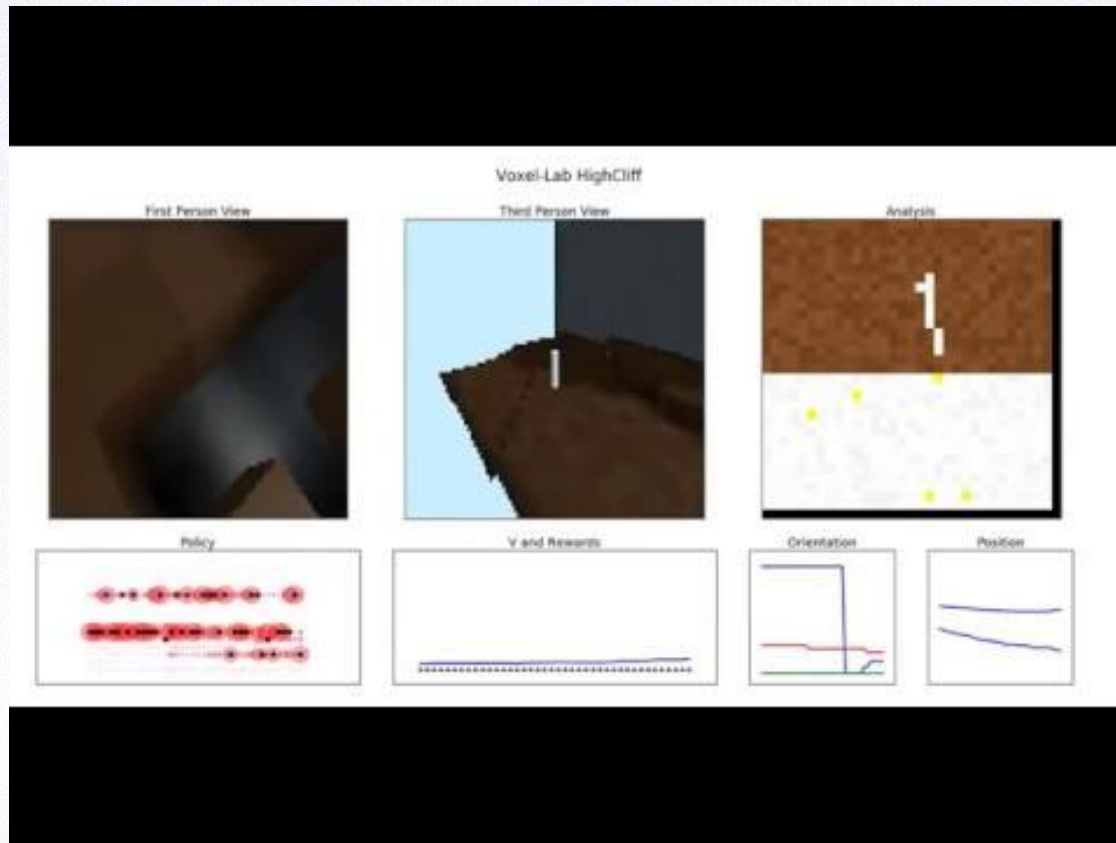
DM Lab



Voxel Lab



# Learns to build stairs and towers






Learns to imagine  
building them


Building stairs

Building towers

input  
rollout



# Naturalistic landscape - imagines eating the yellow block





# Towards a Definition of Disentangled Representations

Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, Alexander Lerchner



# Motivations

- What is the role Group theory in "feature/representation learning"?
- invariance vs covariance vs equivariance
- When is "Disentanglement" useful?
- Our intuition should not rely on a specific coordinate system
- Lack of universality => lack of generalisation to domains where we have no intuition



# Group theory 1:1

## What is a Group?

- A group  $G$  is a set endowed with a binary operator  $\circ$ . The operator must satisfy the following properties:
  - Closure:  $x_1, x_2$  in  $G$ , then  $x_1 \circ x_2$  is in  $G$
  - Identity: there is an  $e$  in  $G$  such that  $e \circ x = x$  (for all  $x$  in  $G$ )
  - Inverse: for any  $x$  in  $G$  there is a  $y$  such that  $y \circ x = e$
  - Associativity:  $x_1 \circ (x_2 \circ x_3) = (x_1 \circ x_2) \circ x_3$

**Example:**

**$G = (\text{Reals}, \circ=+, e = 0)$**

**Example:**

**$G = (\text{Rotation matrices}, \circ=\text{matrix product}, e = I)$**





# Group theory 1:1

What is a representation of a group?

**A representation of a group is a mapping from group element  $g$  to an operator  $\rho(g)$  acting on a different space  $H$  with operator  $x$ . Such that  $\rho(g_1 \circ g_2) = \rho(g_1) \times \rho(g_2)$**

	1 <sup>st</sup>	2 <sup>nd</sup>	3 <sup>rd</sup>	
QUARKS	$\begin{matrix} 2/3 \\ u \\ \text{up} \end{matrix}$	$\begin{matrix} 2/3 \\ c \\ \text{charm} \end{matrix}$	$\begin{matrix} 2/3 \\ t \\ \text{top} \end{matrix}$	$\begin{matrix} 0 \\ g \\ \text{gluons} \end{matrix}$ strong nuclear force
	$\begin{matrix} -1/3 \\ d \\ \text{down} \end{matrix}$	$\begin{matrix} -1/3 \\ s \\ \text{strange} \end{matrix}$	$\begin{matrix} -1/3 \\ b \\ \text{bottom} \end{matrix}$	
	$\begin{matrix} -1/3 \\ e \\ \text{electron} \end{matrix}$	$\begin{matrix} -1/2 \\ \mu \\ \text{muon} \end{matrix}$	$\begin{matrix} -1 \\ \tau \\ \text{tau} \end{matrix}$	
LEPTONS	$\begin{matrix} 0 \\ \nu_e \\ \text{e neutrino} \end{matrix}$	$\begin{matrix} 0 \\ \nu_\mu \\ \text{mu neutrino} \end{matrix}$	$\begin{matrix} 0 \\ \nu_\tau \\ \text{tau neutrino} \end{matrix}$	$\begin{matrix} 0 \\ \gamma \\ \text{photon} \end{matrix}$ electromagnetic force
FERMIONS				
GAUGE BOSONS				$\begin{matrix} 0 \\ W^\pm \\ \text{W bosons} \end{matrix}$ weak nuclear force
				$\begin{matrix} 0 \\ Z \\ \text{Z boson} \end{matrix}$





## Group theory 1:1

What is an irreducible representation of a group?

**Representations that leave invariant some subset of  $H$  that cannot be broken down into smaller invariant subsets.**



# Group theory 1:1

## What is a Lie Group?

**A Lie Group is a group that is also a manifold (e.g. Translations and Rotations)**



# Group theory 1:1

## Invariance and Equivariance

**We say that our features  $f$  are *invariant* under a symmetry group  $G$  iff  $f(g \circ \text{data}) = f(\text{data})$  for any  $g$  in  $G$ . Example: output of pooling layers**


**We say that we have *equivariant features*  $f$  under a symmetry group  $G$  if  $f(g \circ \text{data}) = J \times f(\text{data})$   
Example: output of convolution layers, vector fields**



Too abstract?

$$p(\hat{\mathbf{y}}|\mathbf{x}, \varphi) = \mathcal{N}(\mathbf{y}|\mathbf{W}\mathbf{R}(\varphi)\mathbf{W}^T\mathbf{x}, \sigma^2)$$

$$p(\varphi_j) = \frac{1}{2\pi I_0(\|\eta_j\|)} \exp(\eta_j^T T(\varphi_j)).$$



[Learning the Irreducible Representations of Commutative Lie Groups](#)

[2D/3D Rotation-Invariant Detection using Equivariant Filters and Kernel Weighted Mapping](#)



# Invariance, Equivariance, Classifiers and Generative Models

When building classifiers it is desirable that its output to be invariant under some groups (e.g. translation, rotation).

That is, we want to build ***both features and losses*** such that  $\text{Loss}(g \circ \text{data}) = \text{Loss}(\text{data})$ , where  $g = 2\text{D Translations} \times 2\text{D Rotations}$

However, when building generative models, we should seek ***invariant losses***, but the ***representations should be equivariant***. This is because invariant representations lose information, destroying the ability to reconstruct the data.



# General Recipe to build invariant networks

First, build equivariant features

$$\phi(g \circ x) = T_g \circ \phi(x)$$

Second, apply a 'pooling' operator

$$\psi(x) = \int d\mu(g) T_g \circ \phi(x)$$

**Example: convnets**



## Weyl's Principle

**The elementary components of a system are the irreducible representations of the symmetry group of the system. Example: entire physics.**



Why does it matter to AI?

**3D OBJECTS are irreducible representations  
of 3D Translation x Rotation x Scale groups  
(Galileo Group)**



## Group theory 1:1

**Weyl principle  $\Rightarrow$  Irreducible representations  
 $\Leftrightarrow$  Disentanglement (since we have broken  
down our representations into the smallest  
possible invariant sets)**

# Towards a Definition of Disentangled Representations


A space  $Z$  is **disentangled** with respect to a group decomposition  $G = G_1 \times G_2 \times \dots \times G_N$  if:

- 1) There is an action  $\cdot : G \times Z \rightarrow Z$
- 2) The map  $f : W \rightarrow Z$  is equivariant between the actions on  $W$  and  $Z$
- 3) There is a decomposition  $Z = Z_1 \times Z_2 \times \dots \times Z_N$  such that each  $Z_i$  is fixed by the action of all  $G_j$   $j \neq i$  and affected only by  $G_i$


The actions are assumed to preserve any structure of  $Z$  (e.g. be linear or continuous).




# Towards a Definition of Disentangled Representations




generative factors



X vs y



X vs C



# Summary

- The language of symmetry groups and differential geometry allows for generalisation of useful tools such as conv operators beyond classification of images.
- Such ideas have only been superficially explored in ML
- Do you want objects to emerge from neural nets + data? Let's build "Weyl machines" (or automated physicists).



# Thanks to great collaborators

Karol Gregor

Marta Garnelo

Ali Eslami

Frederic Besse

Fabio Viola

Hamza Merzic

Yan Wu

Theophane Weber

Daniel Zoran

Alex Mott

Irina Higgins

Sebastian Racaniere

David Pfau

Mihaela Rosca

Marco Fraccaro

Aaron Van der Oord

Daan Wierstra