

Partial Observability in DRL

Part 2

POMDP Models so far

Short overview

Model free

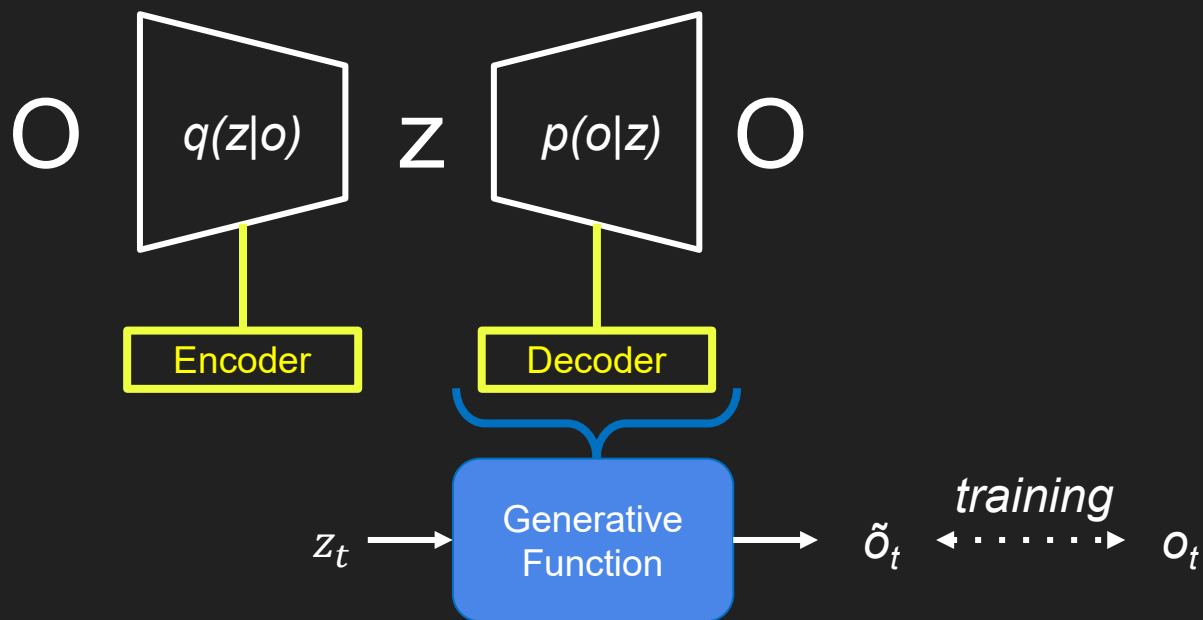
RNN
DRQN
ADRQN

**Explicit
Belief tracking**

DVRL

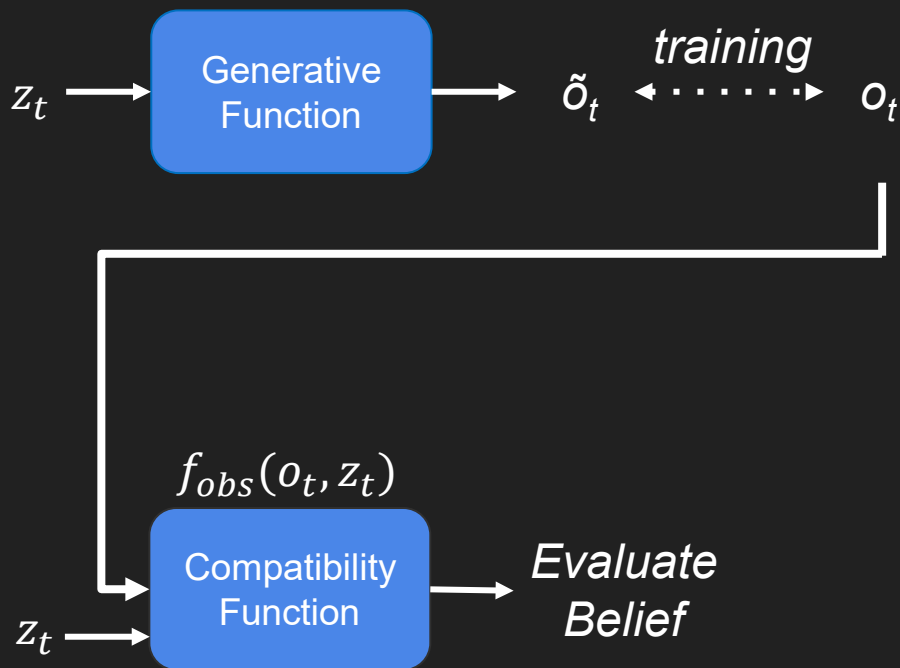
**Implicit
Belief tracking**

DVRL



DVRL


- Requires modeling all observations that define $p(o | h_t)$
 - Needs to learn features irrelevant for RL



- + Directly updates belief based on observation

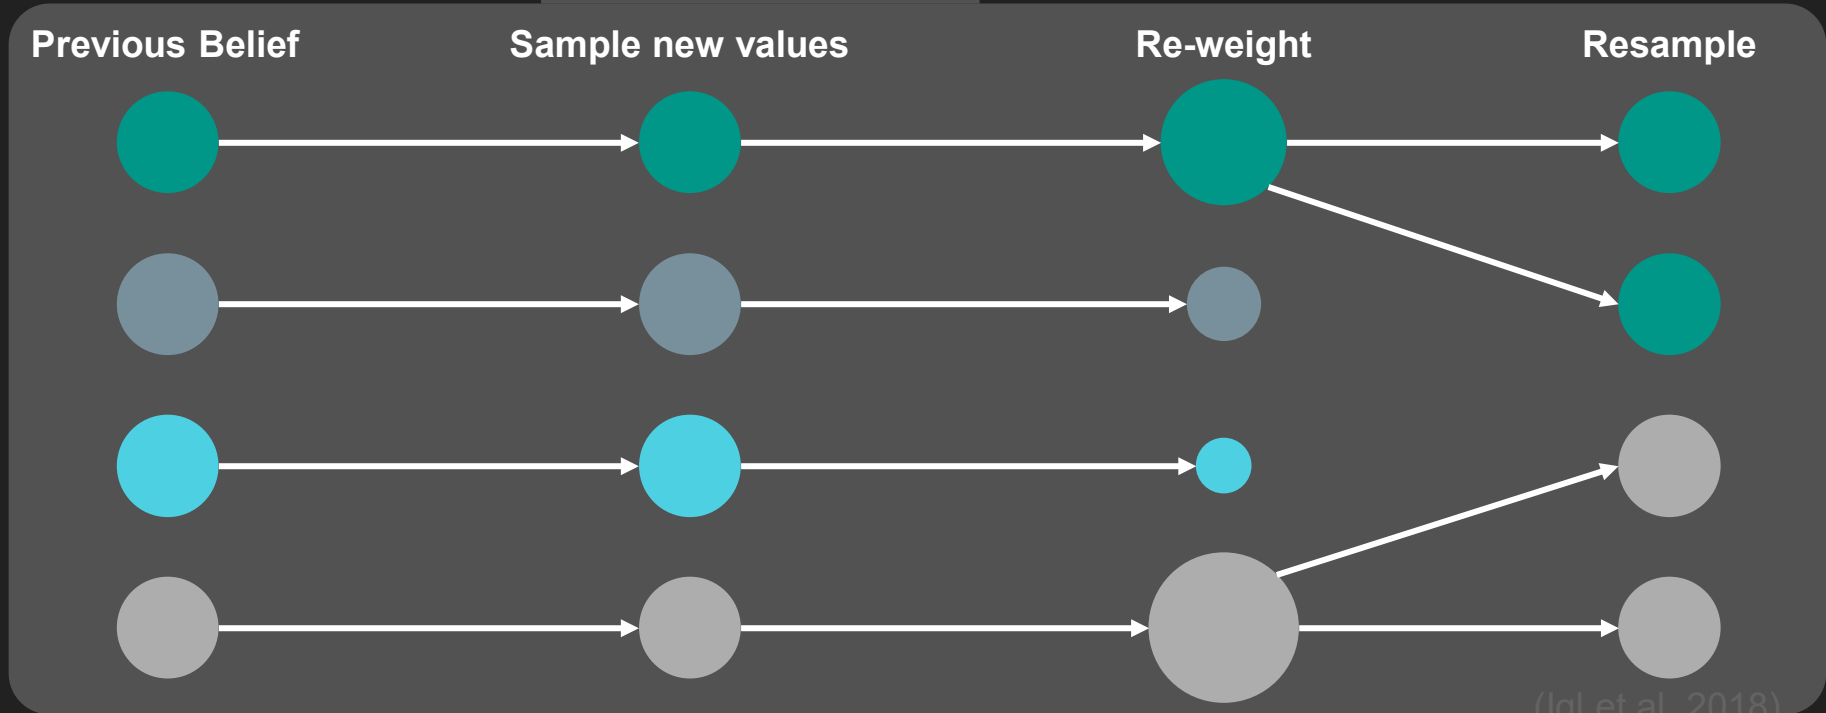
DPFRL

Discriminative Particle Filter Reinforcement Learning



DPFRL

Discriminative Particle Filter Reinforcement Learning



DPFRL


Discriminative Particle Filter Reinforcement Learning

Previous Belief

Sample new values

Re-weight

Resample


$$b_{t-1} \approx \{h_{t-1}, w_{t-1}\}_{k=1}^K$$

DPFRL


Discriminative Particle Filter Reinforcement Learning

Previous Belief

Sample new values

Re-weight

Resample



DPFRL

Discriminative Particle Filter Reinforcement Learning

Previous Belief

Sample new values

Re-weight

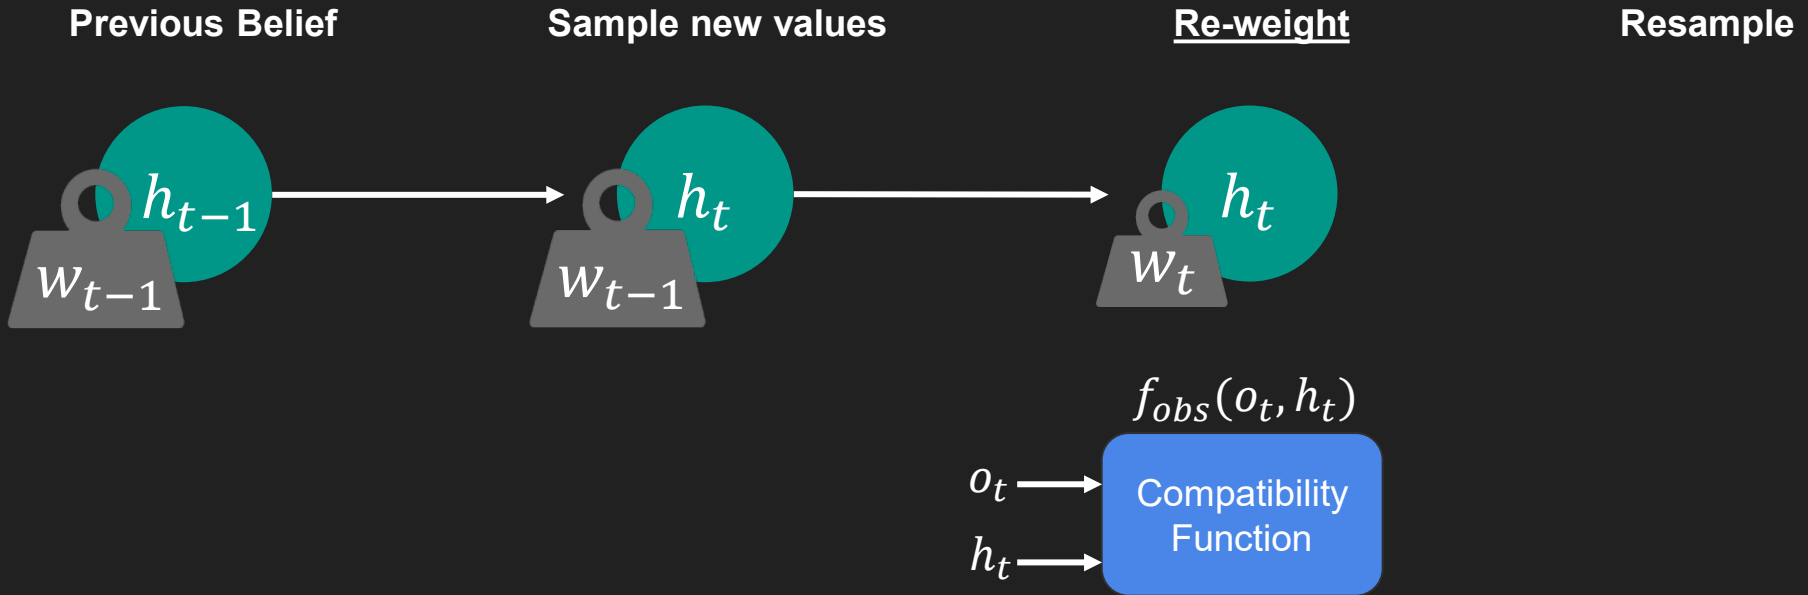
Resample



$$h_t \sim f_{trans}(h_{t-1}, a_t, o_t)$$

DPFRL

Discriminative Particle Filter Reinforcement Learning



DPFRL

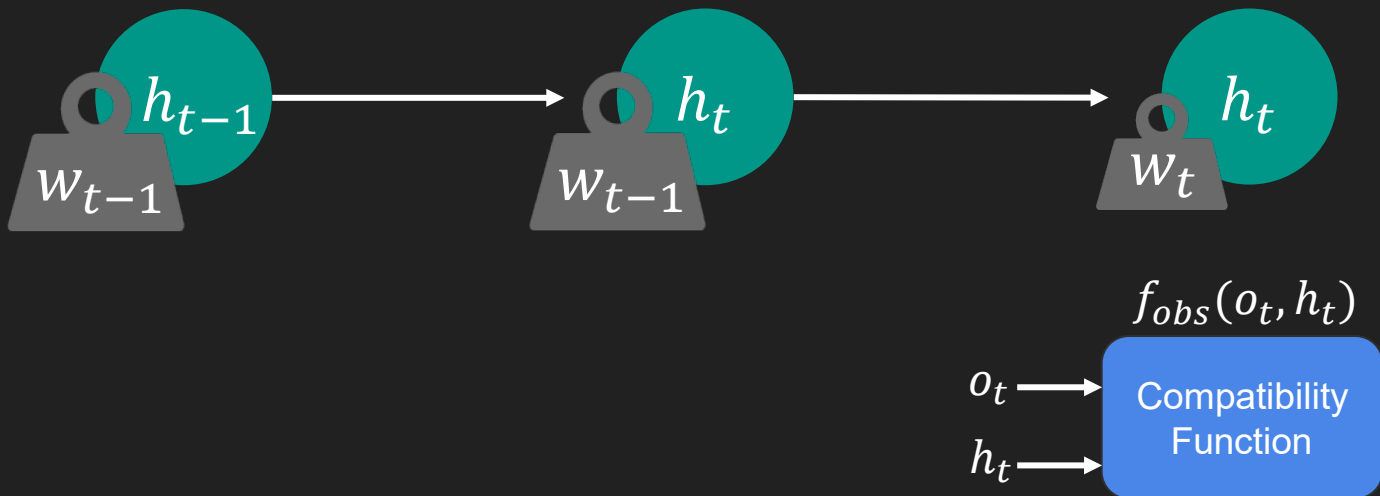
Discriminative Particle Filter Reinforcement Learning

Previous Belief

Sample new values

Re-weight

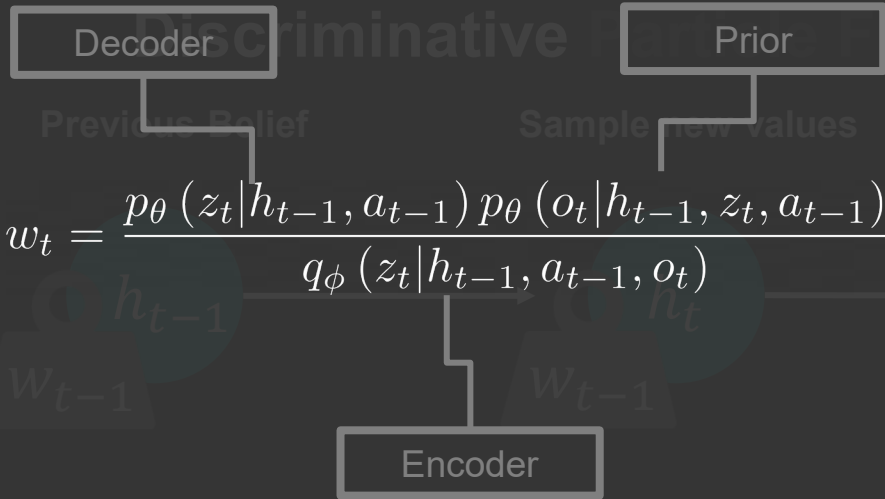
Resample



DVRL


FRL

Filter Reinforcement Learning



Re-weight

Resample



$f_{obs}(o_t, h_t)$

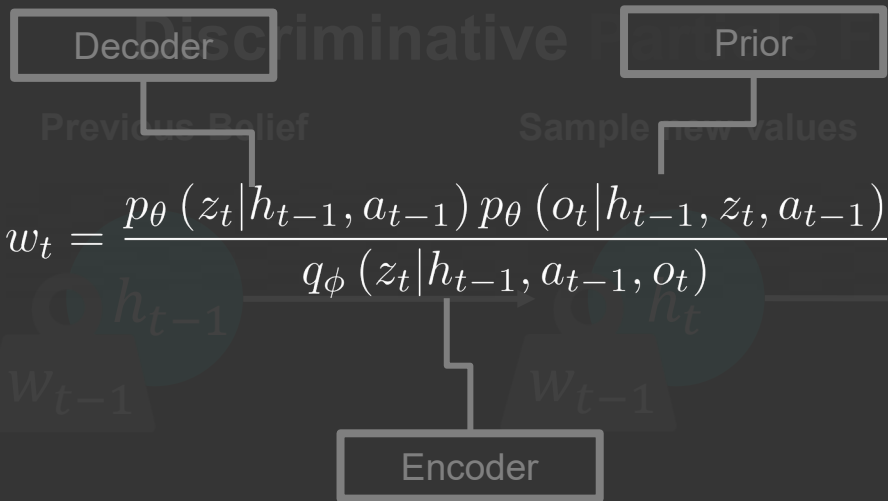
o_t

h_t

Compatibility
Function

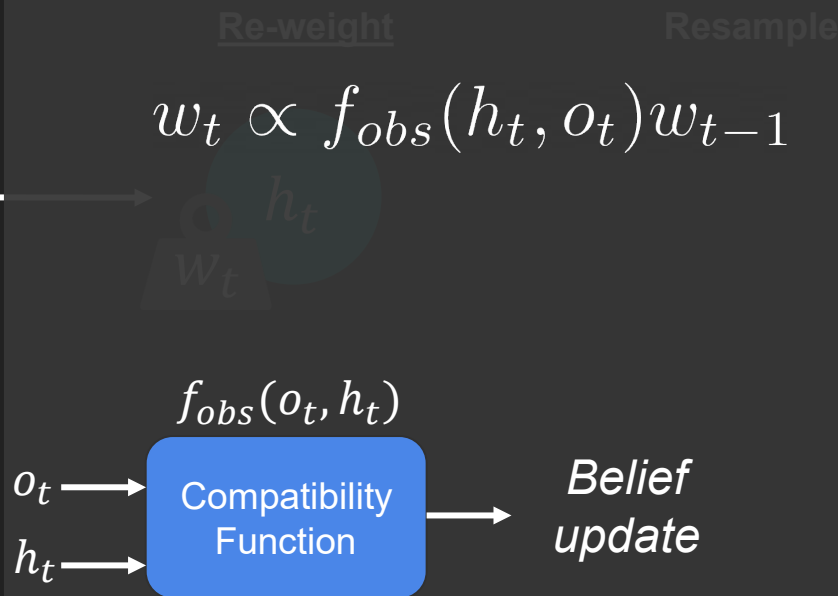
$$ELBO(\theta, \phi) \approx \sum_{t=1}^T \log \left(\frac{1}{K} \sum_{k=1}^K w_t^k \right)$$

DVRL



$$ELBO(\theta, \phi) \approx \sum_{t=1}^T \log \left(\frac{1}{K} \sum_{k=1}^K w_t^k \right)$$

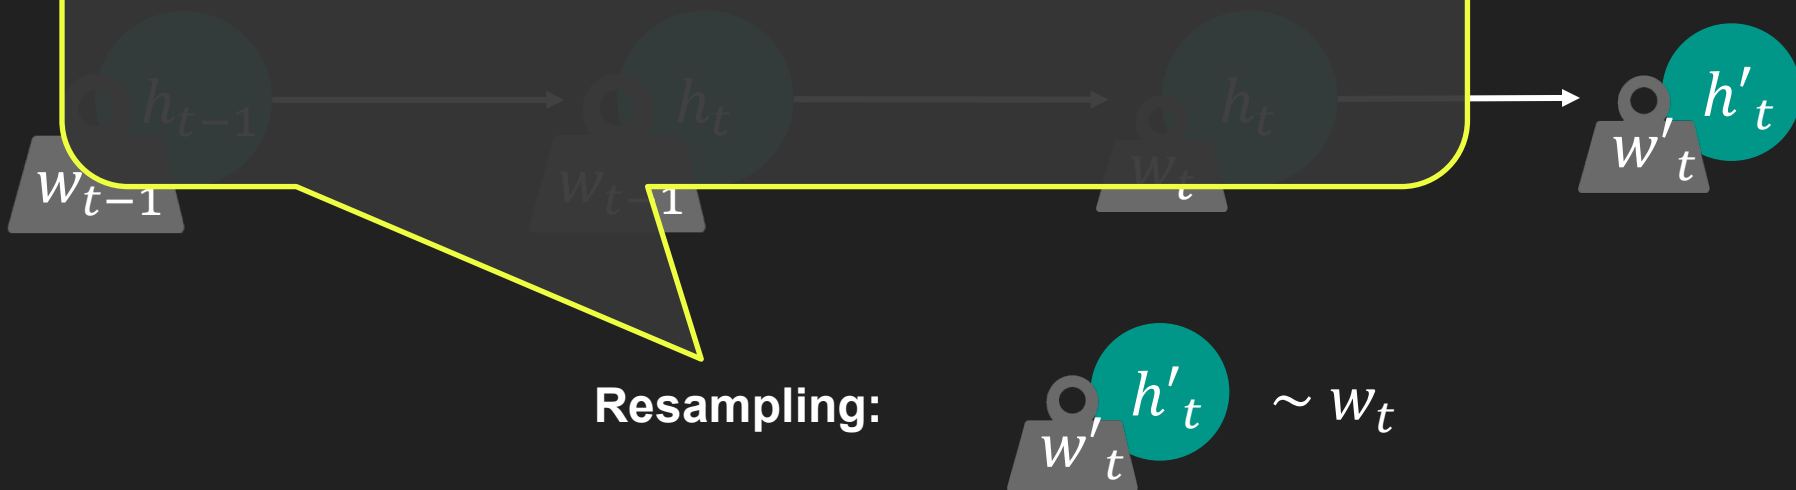
DPFRL



DPFRL

Problem: Particle degeneracy

Most of the particles have a near-zero weight



DPFRL

Problem: Particle degeneracy

Most of the particles have a near-zero weight

Solution: Soft-resampling

Provides approximate gradients for the non-differentiable resampling step

Learning

Resample

w_{t-1}

1




Soft-Resampling:

A diagram showing a particle with a weight w'_t and a value h'_t . The weight is represented by a grey trapezoid with a small circle on top, and the value is represented by a teal circle.

$$w'_t \sim \alpha w_t + (1 - \alpha) \frac{1}{K}$$

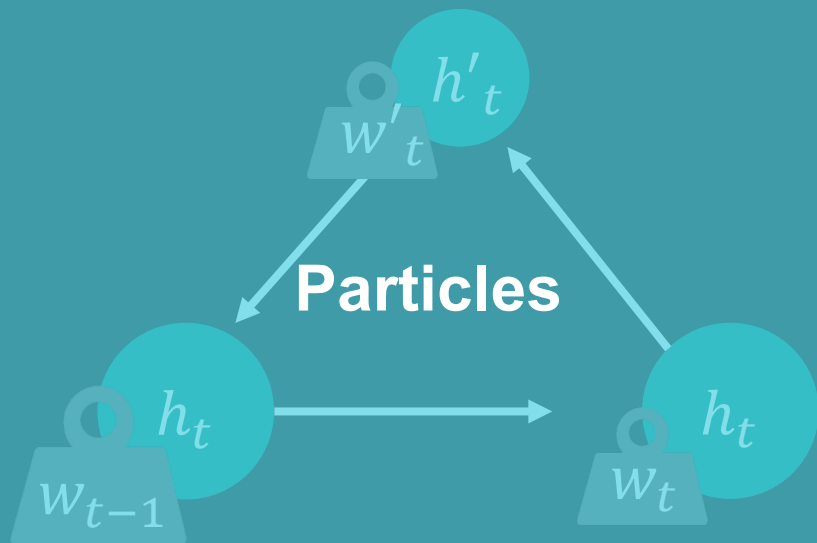
DPFRL

Discriminative Particle Filter Reinforcement Learning



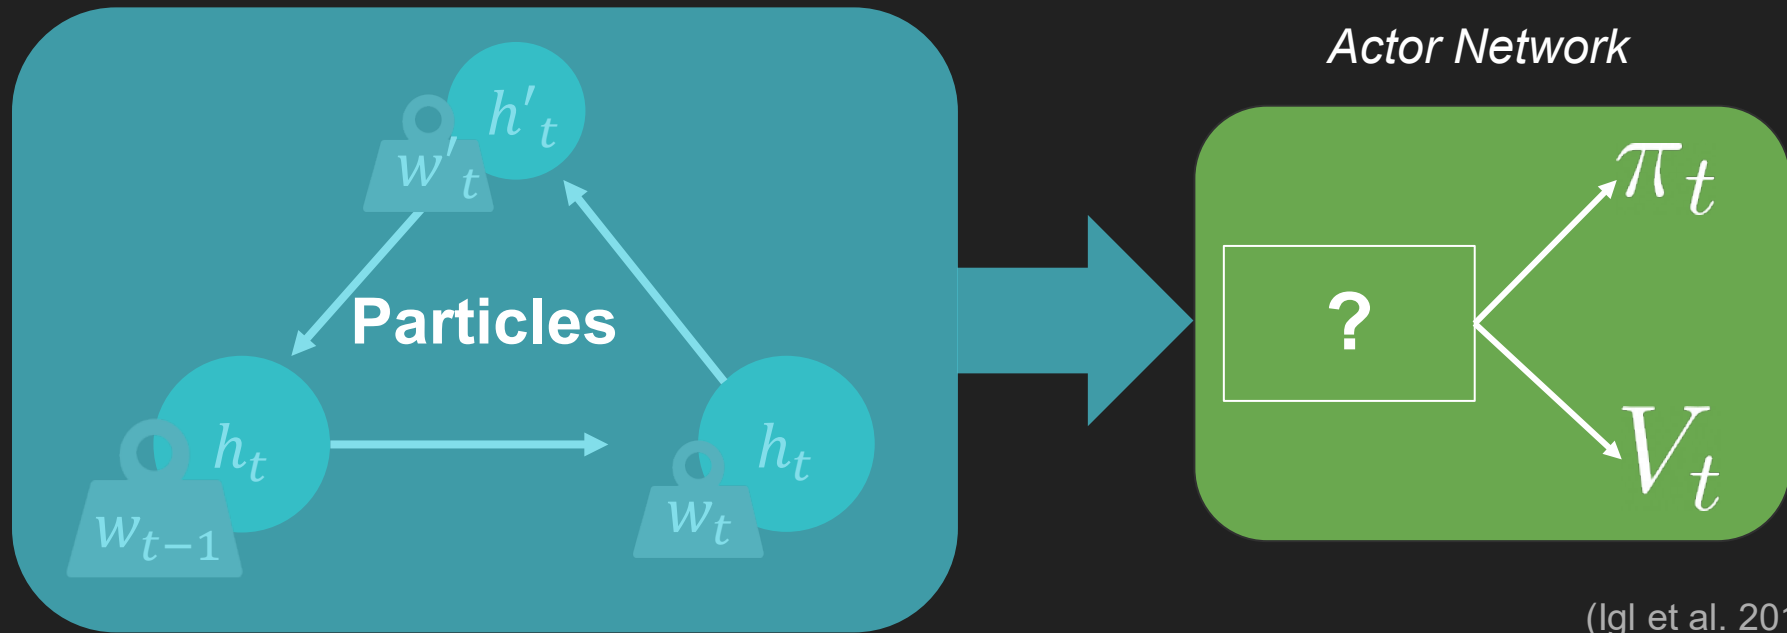
DPFRL

Discriminative Particle Filter Reinforcement Learning



DPFRL

Discriminative Particle Filter Reinforcement Learning



Recall Reinforcement Learning in DVRL

Particles

Properties

- Ordering is irrelevant
- Particle set may be large

Actor Network

RNN

π_t

V_t

Simple but powerful solution (used in PF-RNNs)

Particles

Properties

- Ordering is irrelevant
- Particle set may be large

Actor Network

Mean

π_t

V_t

DPFRL

Reinforcement Learning in DPFRL

Particles

Properties

- Ordering is irrelevant
- Particle set may be large

Actor Network

MGF

π_t

V_t

Short excursion

Moment-Generating Functions (MGFs)

0th Total probability

1st Expectation

2nd Variance

3rd Skewness

...

...

Short excursion

Moment-Generating Functions (MGFs)

Total probability

$$\mathbb{P} [X] = 1$$

Centralized Moment

Expectation

$$\mathbb{E} [X]$$

Variance

$$\mathbb{E} \left[(X - \mathbb{E}(X))^2 \right]$$

Skewness

$$\mathbb{E} \left[(X - \mathbb{E}(X))^3 \right]$$

...

$$\mathbb{E} \left[(X - \mathbb{E}(X))^n \right]$$

Short excursion

Moment-Generating Functions (MGFs)

Total probability $\mathbb{E}(X^0)$

Expectation $\mathbb{E}(X^1)$

Variance $\mathbb{E}(X^2)$

Skewness $\mathbb{E}(X^3)$

... $\mathbb{E}(X^n)$



Raw Moment

Short excursion

Moment-Generating Functions (MGFs)

$$\mathbb{E}(X^0)$$

$$\mathbb{E}(X^1) \quad E[e^{tX}] = 1 + t E[X] + \frac{t^2 E[X]^2}{2!} + \frac{t^3 E[X]^3}{3!} + \dots$$

$$\mathbb{E}(X^2)$$

$$\mathbb{E}(X^3)$$

$$\mathbb{E}(X^n)$$



Raw Moment

Short excursion

Moment-Generating Functions (MGFs)

$$\begin{array}{l} \mathbb{E}(X^0) \\ \mathbb{E}(X^1) \\ \mathbb{E}(X^2) \\ \mathbb{E}(X^3) \\ \mathbb{E}(X^n) \end{array} \quad E[e^{tX}] = \boxed{1} + t \boxed{E[X]} + \frac{t^2 \boxed{E[X]^2}}{2!} + \frac{t^3 \boxed{E[X]^3}}{3!} + \dots$$

Raw Moment

Short excursion

Moment-Generating Functions (MGFs)

$$\mathbb{E}(X^0)$$

$$\mathbb{E}(X^1) \quad E[e^{tX}] = 1 + t E[X] + \frac{t^2 E[X]^2}{2!} + \frac{t^3 E[X]^3}{3!} + \dots$$

$$\mathbb{E}(X^2)$$

$$\mathbb{E}(X^3)$$

$$\mathbb{E}(X^n)$$

$$E(X^n) = \left. \frac{d^n E[e^{tX}]}{dt^n} \right|_{t=0}$$

Short excursion

Moment-Generating Function

Take home message:

A Moment-generating function $M_X(t)$ is an alternative specification of the probability distribution

$$M_X(t) := E[e^{tX}]$$

$E(X^n) = \frac{d^n E[e^{tX}]}{dt^n} \Big|_{t=0}$

Properties

- + Permutation invariant
- + Computationally efficient
- + Easy to optimize
- + Also works well with a large number of particles

Trainable vector \mathbf{v}



Allows model to extract useful moment features for decision making

Implementation in DPFRL

$$M_b = E[e^{\mathbf{v}^T h}] = \sum_i^K w_i e^{\mathbf{v}^T h_i}$$

DPFRL: Ablation study

Discriminative Particle Filter Reinforcement Learning



Flickering Atari Games

Natural Flickering Atari Games

DPFRL: Ablation study

Discriminative Particle Filter Reinforcement Learning

RNN

Mean

MGF

Environments

Envs
Pong
ChopperCommand
MsPacman
Centipede
BeamRider
Frostbite
Bowling
IceHockey
DDunk
Asteroids

DPFRL-GRUmerge
13.14±4.01
1,530±29.31
1,930±48.54
4,093±76.4
603.8±40.25
252.1±0.48
29.50±0.33
-5.85±0.30
-14.39±0.24
1,397±11.44

DPFRL-mean
-5.53±14.35
1,091±109.9
1,878±63.86
3,599±439.8
645.5±227.4
178.4±81.70
26.0±0.81
-6.25±1.96
-14.42±0.18
1,433±40.73

DPFRL
15.65±1.99
1,566±67.03
2,106±123.9
4,164±23.0
682.9±37.42
260.2±4.60
29.45±0.13
-6.08±0.18
-15.36±0.96
1,406±132.3

Results on the Natural Flickering Atari Games dataset

DPFRL: Ablation study

Discriminative Particle Filter Reinforcement Learning

Generative

Discriminative

Environments

Envs
Pong
ChopperCommand
MsPacman
Centipede
BeamRider
Frostbite
Bowling
IceHockey
DDunk
Asteroids

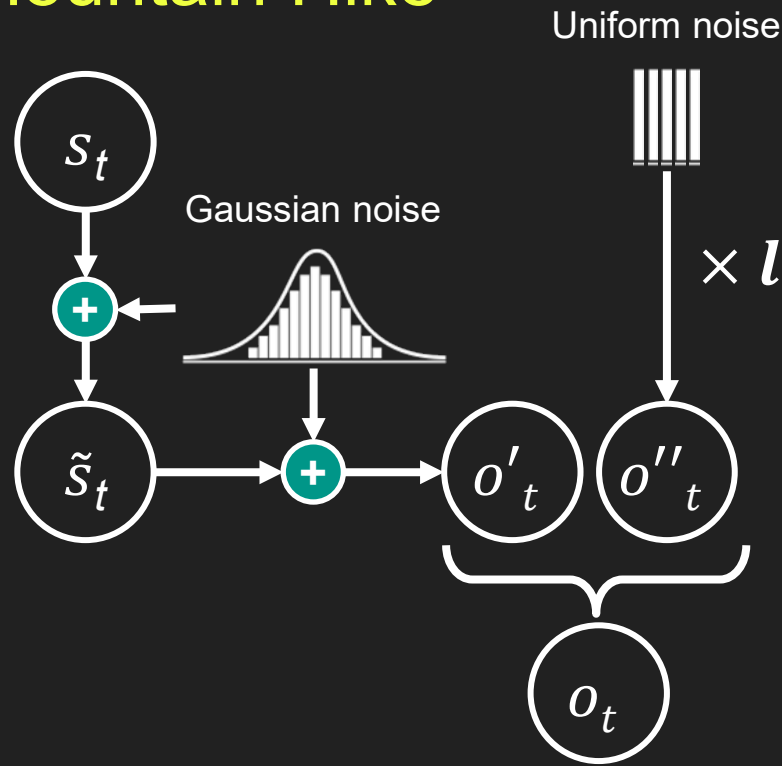

DPFRL-generative
-20.21±0.02
1,027± 12.94
2,130±182.3
3,194±339.4
498.1±8.38
255.9±2.78
24.68±0.13
-7.88±0.30
-15.59±0.06
1,415±5.33

DPFRL
15.65±1.99
1,566±67.03
2,106±123.9
4,164±23.0
682.9±37.42
260.2±4.60
29.45±0.13
-6.08±0.18
-15.36±0.96
1,406±132.3

Results on the Natural Flickering Atari Games dataset

DPFRL: Noise robustness - Mountain Hike

Discriminative Particle Filter Reinforcement Learning



DPFRL: Noise robustness - Mountain Hike

Discriminative Particle Filter Reinforcement Learning

RNN

DVRL





DPFRL

Effect of varying l

$l = 0$


$l = 50$

$l = 100$



DPFRL: Real world environment

Discriminative Particle Filter Reinforcement Learning



Natural habitat dataset

Receiving a first-person RGB-D image in each step.

A robot needs to navigate in previously unseen environments.

DPFRL: Real world environment

Discriminative Particle Filter Reinforcement Learning

Results	SPL	Success Rate	Reward
DPFRL	0.79	0.88	12.82±5.82
DVRL	0.09	0.11	5.22±2.24
RNN	0.63	0.74	10.14±2.82

Natural habitat dataset

Receiving a first-person RGB-D image in each step.

A robot needs to navigate in previously unseen environments.

DPFRL: Summary

Advantages

- + No need to infer an observation model
- + Better robustness to noise in comparison to generative models
- + Simpler than DVRL

Weaknesses

- No reconstruction loss on observation restricts the learning signal
 - Limits sample efficiency and accuracy

POMDP

Approaches

Model free

RNN
DRQN
ADRQN

Explicit Belief tracking

DVRL
DPFRL

Implicit Belief tracking

VRN

Variational Recurrent Models

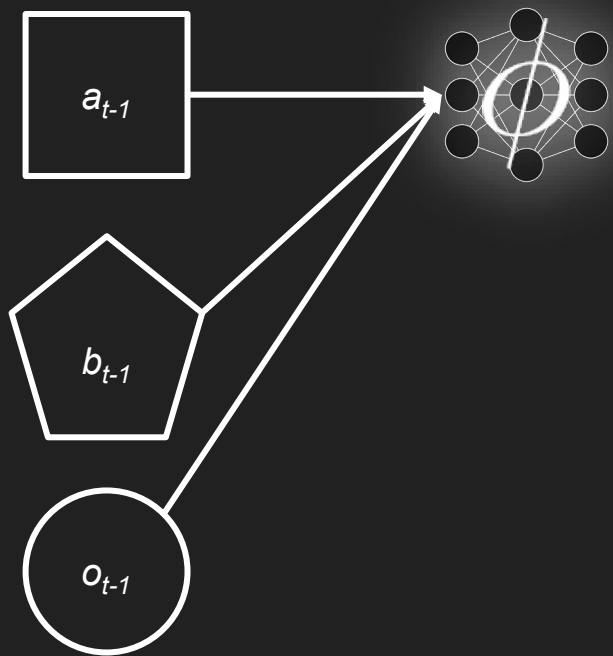
Objective

Piggyback on existing Algorithms for fully observable tasks


Idea

1. Conjecture a state (*Believe state*)
2. Solve the problem as if it was a Fully Observable MDP


Variational Recurrent Models



Variational Recurrent Models



Variational Recurrent Models



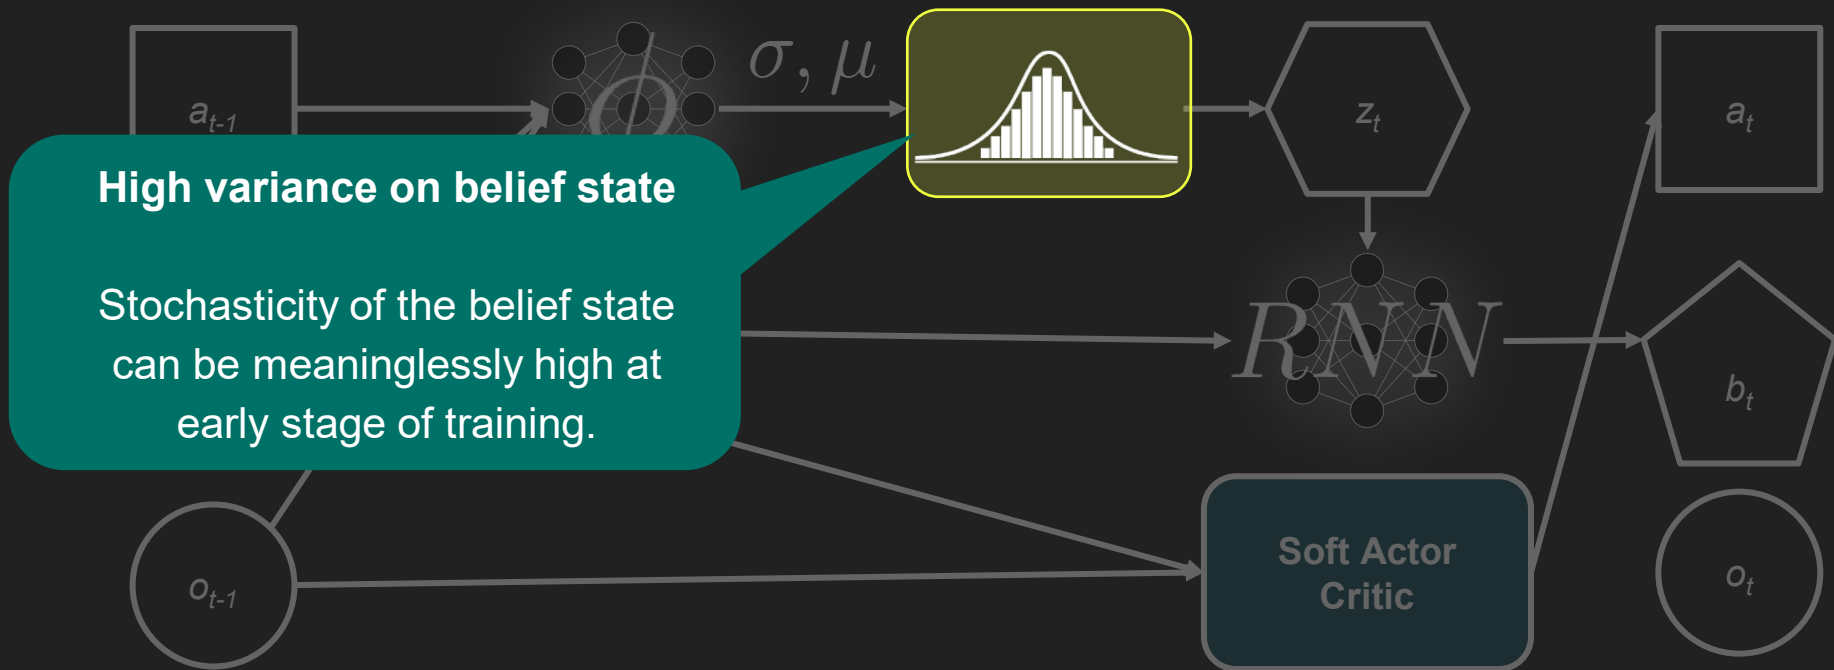
Variational Recurrent Models



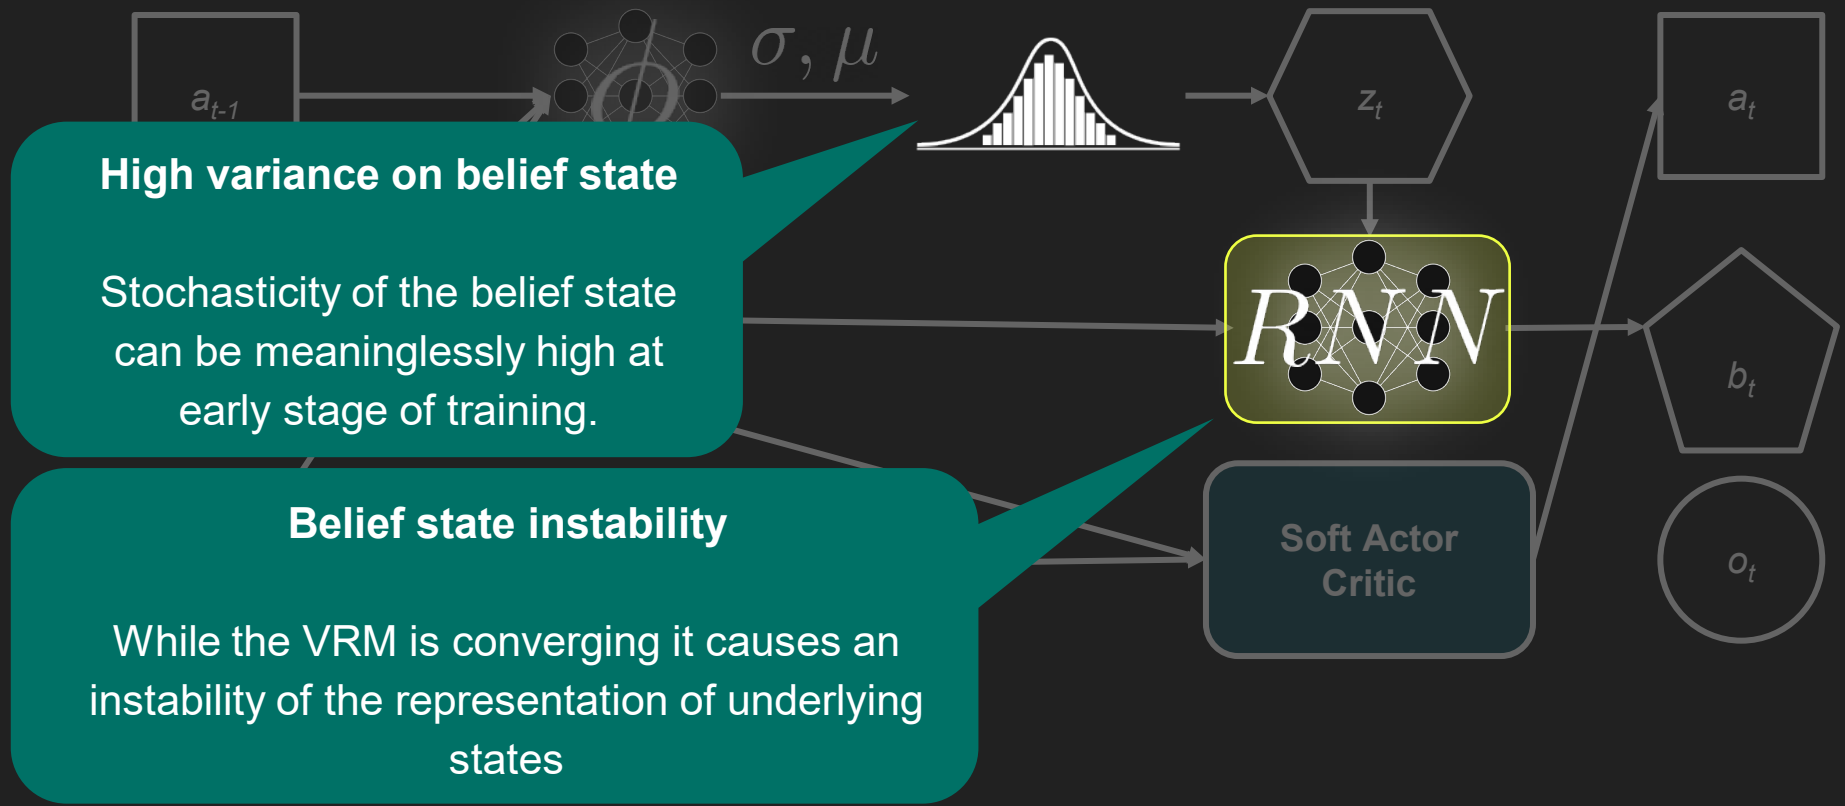
Variational Recurrent Models



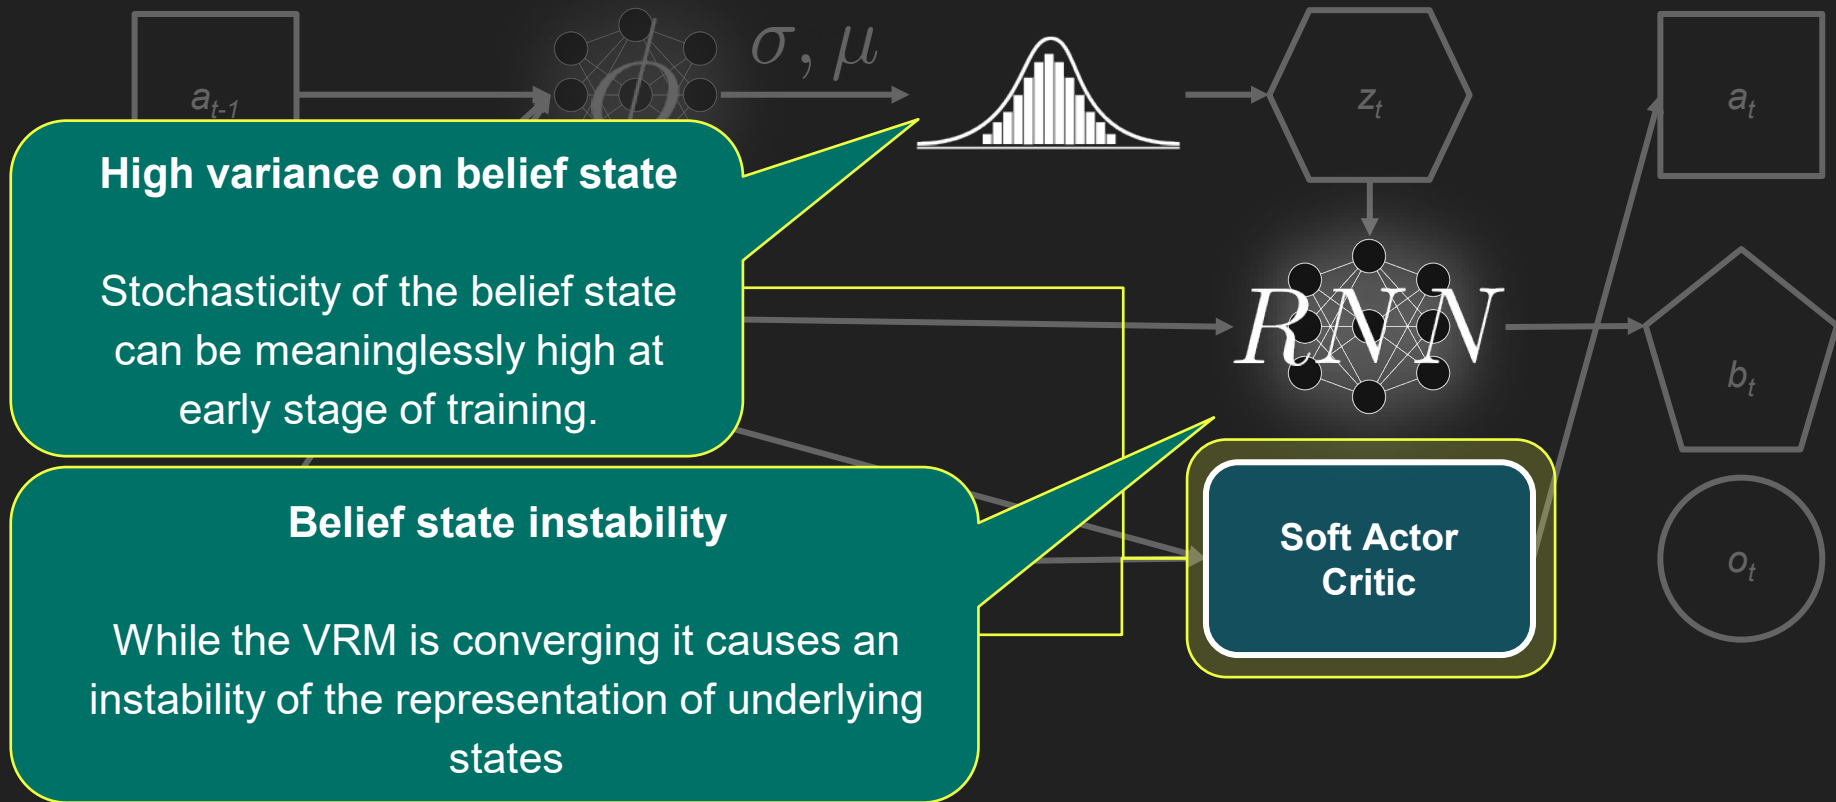
Variational Recurrent Models: Challenges



Variational Recurrent Models: Challenges



Variational Recurrent Models: Challenges



Variational Recurrent Models: Solution

High variance on belief state

VRM

First-impression model

- Keeps the representation stable for the RL during training
- Pre-trained and not updated during exploration

Keep-learning model

- Learn from new observations obtained after the policy update by the RL
- Trained during exploration

Variational Recurrent Models

Algorithm

Algorithm 1 Variational Recurrent Models with Soft Actor Critic

Initialize the first-impression VRM \mathcal{M}_f and the keep-learning VRM \mathcal{M}_k , the RL controller \mathcal{C} , and the replay buffer \mathcal{D} , global step $t \leftarrow 0$.

repeat

Initialize an episode, assign \mathcal{M} with zero initial states.

while episode not terminated **do**

Sample an action \mathbf{a}_t from $\pi(\mathbf{a}_t | \mathbf{d}_t, \mathbf{x}_t)$ and execute $a_t, t \leftarrow t + 1$.

Record $(\mathbf{x}_t, \mathbf{a}_t, done_t)$ into \mathcal{B} .

Compute 1-step forward of both VRMs using inference models.

if $t == step_start_RL$ **then**

For N epochs, sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_f (Eq. [II](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_KLVRM) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_k (Eq. [5](#), [6](#), [7](#), [8](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_RL) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{R} (Eq. [II](#)).

end if

end while

until training stopped

Initialization

Variational Recurrent Models

Algorithm

Algorithm 1 Variational Recurrent Models with Soft Actor Critic

Initialize the first-impression VRM \mathcal{M}_f and the keep-learning VRM \mathcal{M}_k , the RL controller \mathcal{C} , and the replay buffer \mathcal{D} , global step $t \leftarrow 0$.

repeat

Initialize an episode, assign \mathcal{M} with zero initial states.

while episode not terminated **do**

Sample an action \mathbf{a}_t from $\pi(\mathbf{a}_t|\mathbf{d}_t, \mathbf{x}_t)$ and execute $\mathbf{a}_t, t \leftarrow t + 1$.

Record $(\mathbf{x}_t, \mathbf{a}_t, done_t)$ into \mathcal{B} .

Compute 1-step forward of both VRMs using inference models.

if $t == step_start_RL$ **then**

For N epochs, sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_f (Eq. [II](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_KLVRM) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_k (Eq. [5](#), [6](#), [7](#), [8](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_RL) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{R} (Eq. [II](#)).

end if

end while

until training stopped

Fill the replay buffer

Variational Recurrent Models

Algorithm

Algorithm 1 Variational Recurrent Models with Soft Actor Critic

Initialize the first-impression VRM \mathcal{M}_f and the keep-learning VRM \mathcal{M}_k , the RL controller \mathcal{C} , and the replay buffer \mathcal{D} , global step $t \leftarrow 0$.

repeat

Initialize an episode, assign \mathcal{M} with zero initial states.

while episode not terminated **do**

Sample an action \mathbf{a}_t from $\pi(\mathbf{a}_t | \mathbf{d}_t, \mathbf{x}_t)$ and execute $a_t, t \leftarrow t + 1$.

Record $(\mathbf{x}_t, \mathbf{a}_t, done_t)$ into \mathcal{B} .

Compute 1-step forward of both VRMs using inference models.

if $t == step_start_RL$ **then**

For N epochs, sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_f (Eq. II).

end if

if $t > step_start_RL$ and $mod(t, train_interval_KLVRM) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_k (Eq. 5, 6, 7, 8).

end if

if $t > step_start_RL$ and $mod(t, train_interval_RL) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{R} (Eq. II).

end if

end while

until training stopped

Train first-impression
Model

Variational Recurrent Models

Algorithm

Algorithm 1 Variational Recurrent Models with Soft Actor Critic

Initialize the first-impression VRM \mathcal{M}_f and the keep-learning VRM \mathcal{M}_k , the RL controller \mathcal{C} , and the replay buffer \mathcal{D} , global step $t \leftarrow 0$.

repeat

Initialize an episode, assign \mathcal{M} with zero initial states.

while episode not terminated **do**

Sample an action \mathbf{a}_t from $\pi(\mathbf{a}_t|\mathbf{d}_t, \mathbf{x}_t)$ and execute $a_t, t \leftarrow t + 1$.

Record $(\mathbf{x}_t, \mathbf{a}_t, done_t)$ into \mathcal{B} .

Compute 1-step forward of both VRMs using inference models.

if $t == step_start_RL$ **then**

For N epochs, sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_f (Eq. [II](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_KLVRM) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_k (Eq. [5](#), [6](#), [7](#), [8](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_RL) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{R} (Eq. [II](#)).

end if

end while

until training stopped

Train keep-learning
Model

Variational Recurrent Models

Algorithm

Algorithm 1 Variational Recurrent Models with Soft Actor Critic

Initialize the first-impression VRM \mathcal{M}_f and the keep-learning VRM \mathcal{M}_k , the RL controller \mathcal{C} , and the replay buffer \mathcal{D} , global step $t \leftarrow 0$.

repeat

Initialize an episode, assign \mathcal{M} with zero initial states.

while episode not terminated **do**

Sample an action \mathbf{a}_t from $\pi(\mathbf{a}_t|\mathbf{d}_t, \mathbf{x}_t)$ and execute \mathbf{a}_t , $t \leftarrow t + 1$.

Record $(\mathbf{x}_t, \mathbf{a}_t, done_t)$ into \mathcal{B} .

Compute 1-step forward of both VRMs using inference models.

if $t == step_start_RL$ **then**

For N epochs, sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_f (Eq. [II](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_KLVRM) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{M}_k (Eq. [5](#), [6](#), [7](#), [8](#)).

end if

if $t > step_start_RL$ and $mod(t, train_interval_RL) == 0$ **then**

Sample a minibatch of samples from \mathcal{B} to update \mathcal{R} (Eq. [II](#)).

end if

end while

until training stopped

Algorithm 1 Soft Actor-Critic

Initialize parameter vectors $\psi, \bar{\psi}, \theta, \phi$.

for each iteration **do**

for each environment step **do**

$\mathbf{a}_t \sim \pi_\phi(\mathbf{a}_t|\mathbf{s}_t)$

$\mathbf{s}_{t+1} \sim p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$

$\mathcal{D} \leftarrow \mathcal{D} \cup \{(\mathbf{s}_t, \mathbf{a}_t, r(\mathbf{s}_t, \mathbf{a}_t), \mathbf{s}_{t+1})\}$

end for

for each gradient step **do**

$\psi \leftarrow \psi - \lambda_V \hat{\nabla}_\psi J_V(\psi)$

$\theta_i \leftarrow \theta_i - \lambda_Q \hat{\nabla}_{\theta_i} J_Q(\theta_i)$ for $i \in \{1, 2\}$

$\phi \leftarrow \phi - \lambda_\pi \hat{\nabla}_\phi J_\pi(\phi)$

$\bar{\psi} \leftarrow \tau\psi + (1 - \tau)\bar{\psi}$

end for

end for

Variational Recurrent Models

Results

Pendulum



Goal: Swing the pendulum up and maintain highest position

Angle

Angular velocity

Variational Recurrent Models

Results

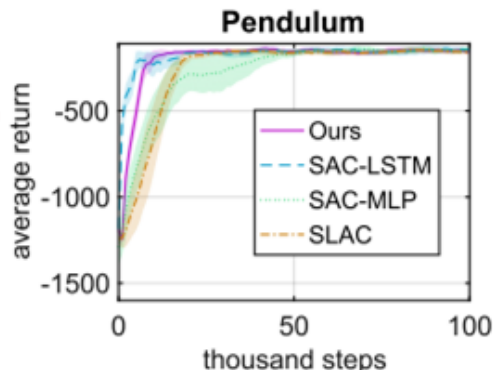
VRM

SAC-MLP

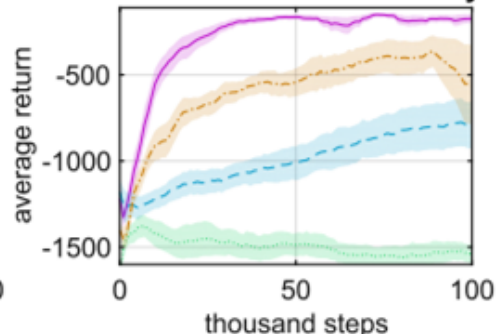
SAC-RNN

SLAC

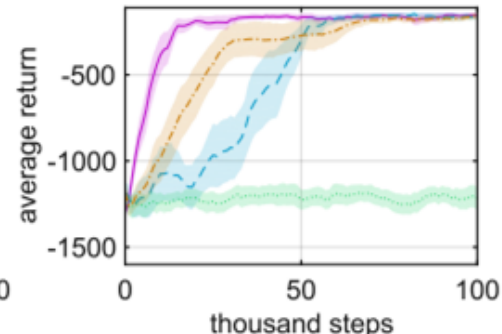
Pendulum



Pendulum - velocities only



Pendulum - no velocities



Variational Recurrent Models

Results


VRM

SAC-MLP

SAC-RNN

SLAC

CartPole



Goal: Balance pole upright without running away

Position

Velocity

Angle

Angular velocity

Variational Recurrent Models


Results

VRM

SAC-MLP

SAC-RNN

SLAC



Variational Recurrent Models

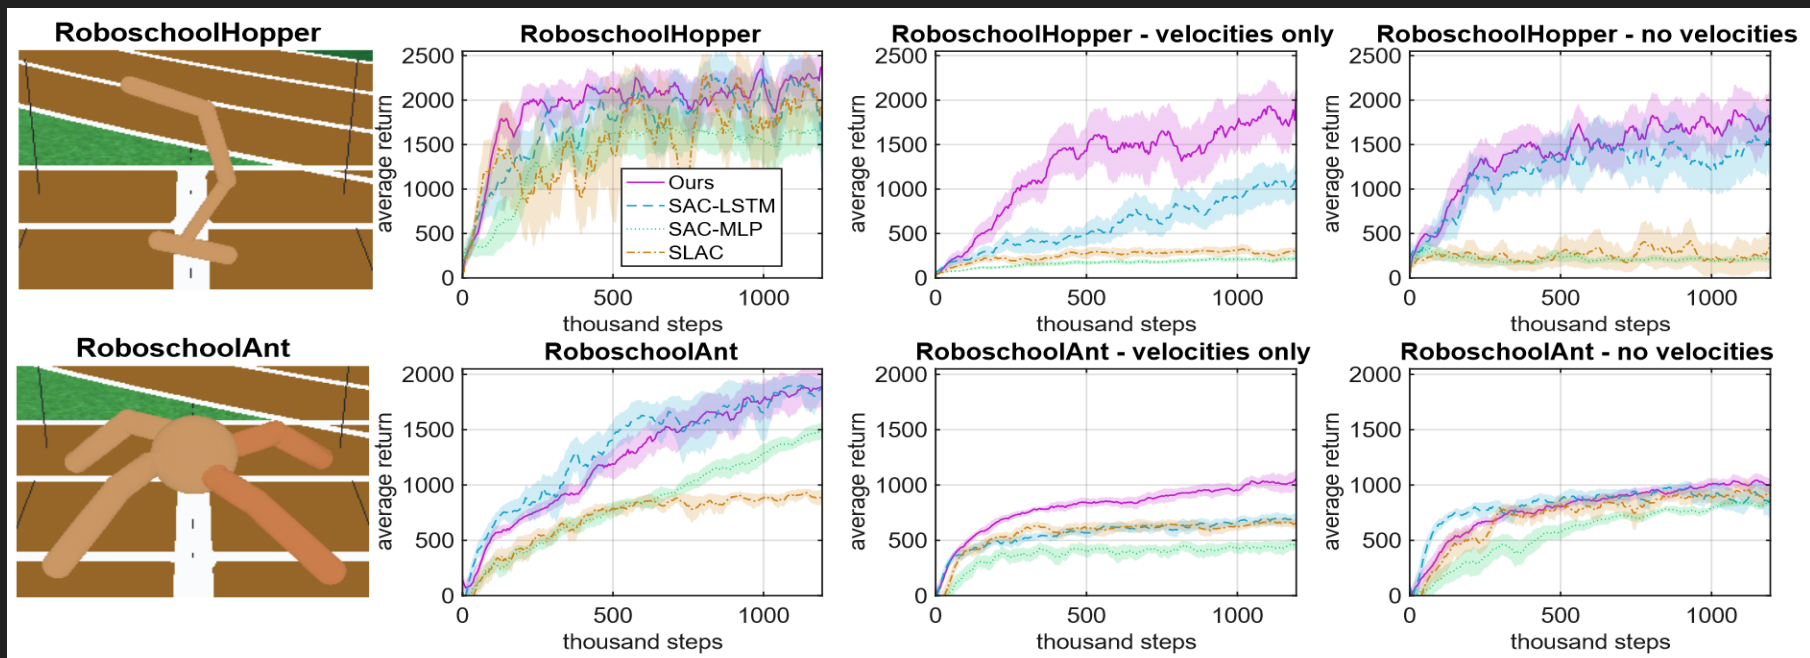
Results

VRM

SAC-MLP

SAC-RNN

SLAC



Variational Recurrent Models

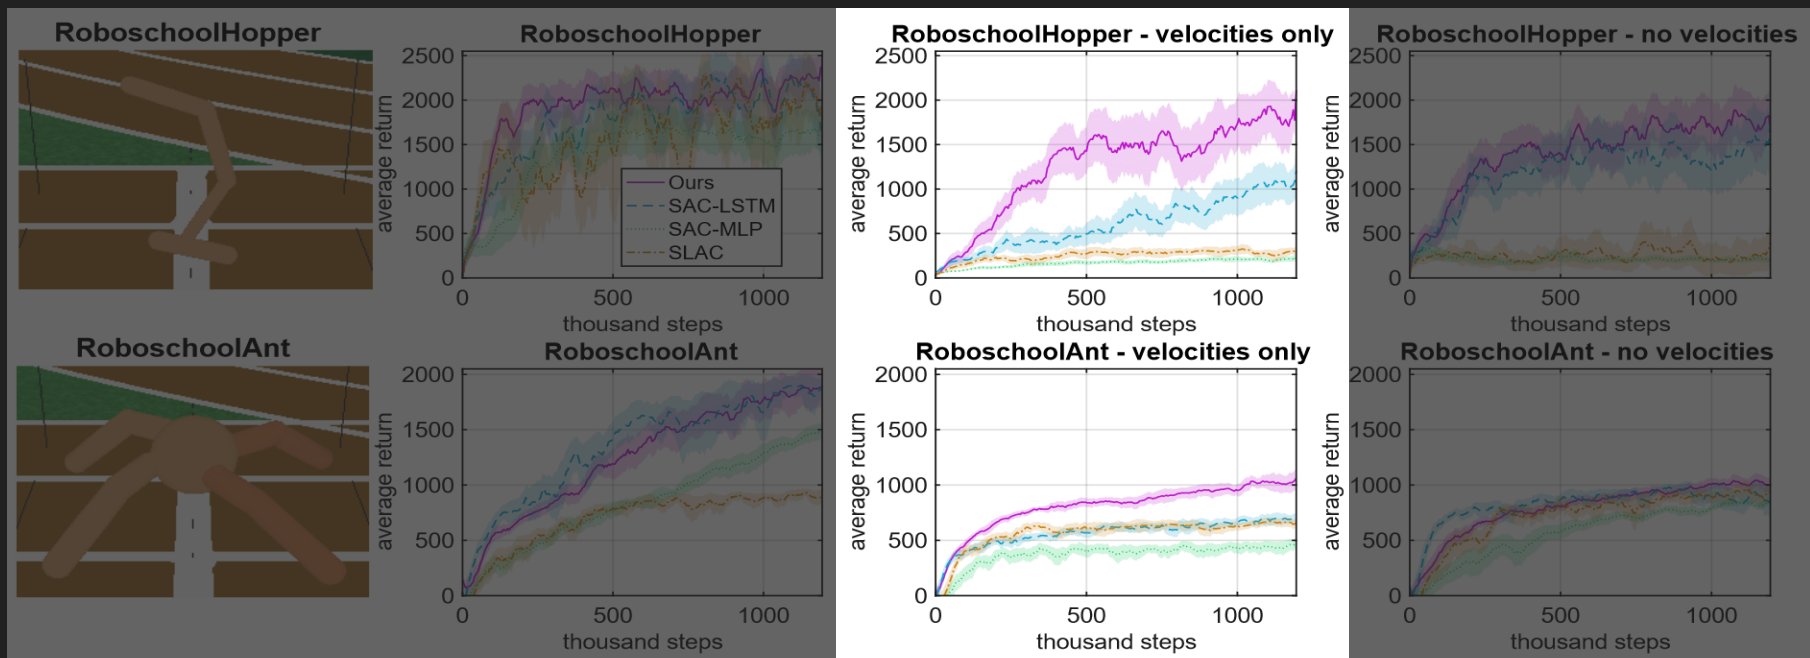
Results

VRM

SAC-MLP

SAC-RNN

SLAC



Variational Recurrent Models


Results

VRM

SAC-MLP

SAC-RNN

SLAC



Variational Recurrent Models

Results


VRM

SAC-MLP

SAC-RNN

SLAC

Sequential target reaching task



Goal

An agent needs to reach 3 different targets in a certain sequence

Requires long-term memorization of past events

Variational Recurrent Models

Results


VRM

SAC-MLP

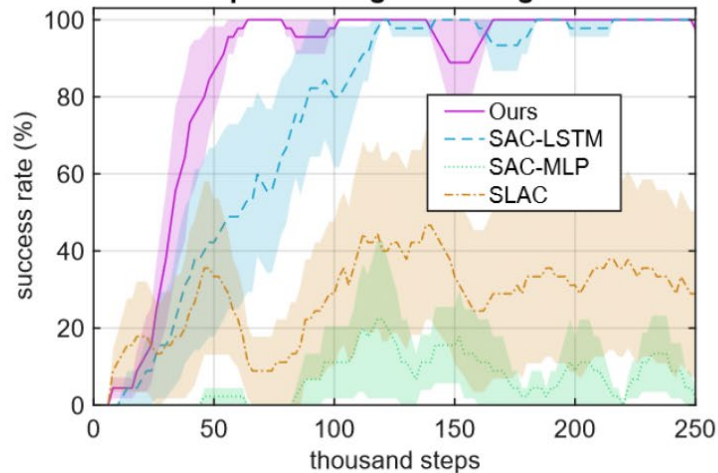
SAC-RNN

SLAC

Sequential target reaching task



Sequential target reaching task



Variational Recurrent Models (ICLR 2020)

Summary

Advantages

- + Simple to implement
- + Gaussian instead of Particle Filter
 - More sample efficient
 - Includes prior knowledge on belief distribution
- + Can make use of advancements in Fully Observable MDP tasks

Weaknesses

- The training of a RNN is still data intensive
- The use of a Gaussian distribution makes the model less suitable for general POMDPs
- The model was only tested against Model-free algorithms

Variational Recurrent Models (ICLR 2020)

Summary

Advantages

Weaknesses

+

Personal note

+

VRM seems to aim at solving a POMDP where the underlying state is not stochastic in nature. The variational part in the VRNN may mostly help to get a continuous latent variable like it is the case for a VAE.

+

Can make use of advancements in Fully Observable MDP tasks

-

The model was only tested against Model-free algorithms

POMDP

Approaches

Model free

RNN
DRQN
ADRQN

Explicit Belief tracking

DVRL
DPFRL

Implicit Belief tracking

VRN

Belief

VRN

DVRL

DPFRL

Belief tracked by RNN

Belief tracked by
Particle and
summarized by RNN

Belief tracked by
Particle and
summarized by MGF

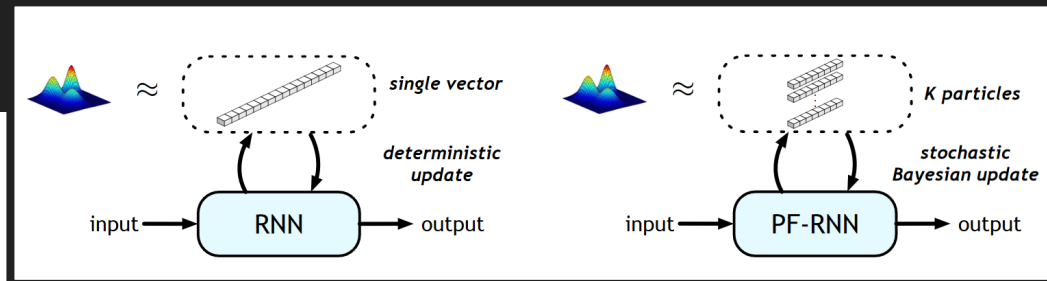
Belief updated using a
Gaussian

Belief updated using a
latent representation of
the observation

Belief updated using
discriminative function

POMDP

Further discussion



Particle Filter Recurrent Neural Networks

Xiao Ma*, Peter Karkus*, David Hsu, Wee Sun Lee
National University of Singapore
{xiao-ma, karkus, dyhsu, leews}@comp.nus.edu.sg

Abstract

Recurrent neural networks (RNNs) have been extraordinarily successful for prediction with sequential data. To tackle highly variable and multi-modal real-world data, we introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new RNN family that explicitly models uncertainty in its internal structure: while an RNN relies on a long, deterministic latent state vector, a PF-RNN maintains a latent *state distribution*, approximated as a set of particles. For effective learning, we provide a fully differentiable particle filter algorithm that updates the PF-RNN latent state distribution according to the Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.

Introduction


Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

thus increasing the number of network parameters and amount of data required for training.

We introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new family of RNNs that seek belief approximation without lengthening the h_t , thus reducing the data required for learning (Del Moral 1996) is a model-based belief approximation. It approximates the belief as a set of particles that typically have well-understood meaning. Particle filtering is the idea of approximating belief as a set of weighted particles, and combining the powerful approximation capacity of RNNs with particle filtering. Like standard RNNs, PF-RNNs approximate the variable and multi-modal belief as a set of weighted latent vectors $\{h^1, h^2, \dots\}$ sampled from the same distribution. Like standard RNNs, PF-RNNs use a model-free approach: PF-RNNs’ latent vector representations are distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian

e.g., Kalman filters, particle filtering is a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2013). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,



POMDP

Further discussion

Particle Filter Recurrent Neural Networks

Xiao Ma*, Peter Karkus*, David Hsu, Wee Sun Lee
National University of Singapore
{xiao-ma, karkus, dyhsu, leews}@comp.nus.edu.sg

Abstract

Recurrent neural networks (RNNs) have been extraordinarily successful for prediction with sequential data. To tackle highly variable and multi-modal real-world data, we introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new RNN family that explicitly models uncertainty in its internal structure: while an RNN relies on a long, deterministic latent state vector, a PF-RNN maintains a latent *state distribution*, approximated as a set of particles. For effective learning, we provide a fully differentiable particle filter algorithm that updates the PF-RNN latent state distribution according to the Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.



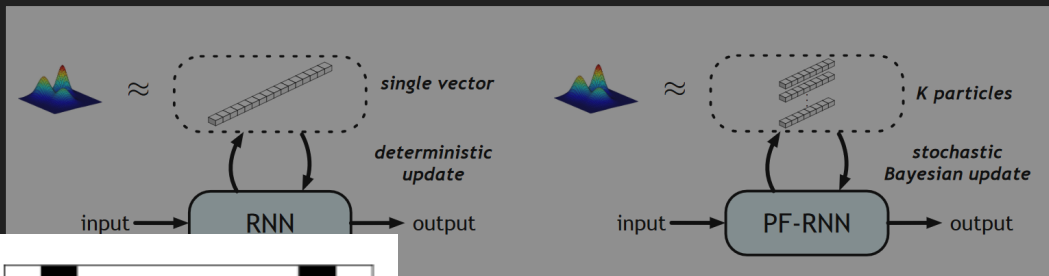
Introduction

Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

thus increasing the number of network parameters and amount of data required for training.

We introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new family of RNNs that approximate belief without lengthening h_t , thus reducing the data required for learning (Del Moral 1996) is a model-based algorithm. It approximates the belief as a set of particles that typically have well-understood meaning from particle filtering: the idea of approximating belief as a set of weighted particles, and the powerful approximation capacity of particle filtering. It approximates the variable and multi-modal latent vectors $\{h^1, h^2, \dots\}$ with the same distribution. Like standard RNNs, PF-RNNs follow a model-free approach: PF-RNNs’ latent vectors are learned distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian based filters, e.g., Kalman filters, particle filtering is a non-parametric approximator that offers a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2015). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,



POMDP

Further discussion

Particle Filter Recurrent Neural Networks

Xiao Ma*, Peter Karkus*, David Hsu, Wee Sun Lee
National University of Singapore
{xiao-ma, karkus, dyhsu, leews}@comp.nus.edu.sg

Abstract

Recurrent neural networks (RNNs) have been extraordinarily successful for prediction with sequential data. To tackle highly variable and multi-modal real-world data, we introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new RNN family that explicitly models uncertainty in its internal structure: while an RNN relies on a long, deterministic latent state vector, a PF-RNN maintains a latent *state distribution*, approximated as a set of particles. For effective learning, we provide a fully differentiable particle filter algorithm that updates the PF-RNN latent state distribution according to the Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.


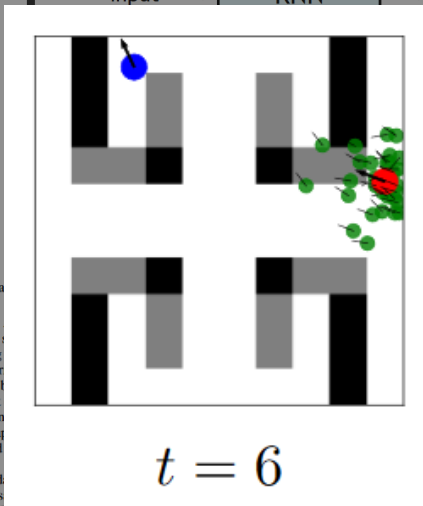
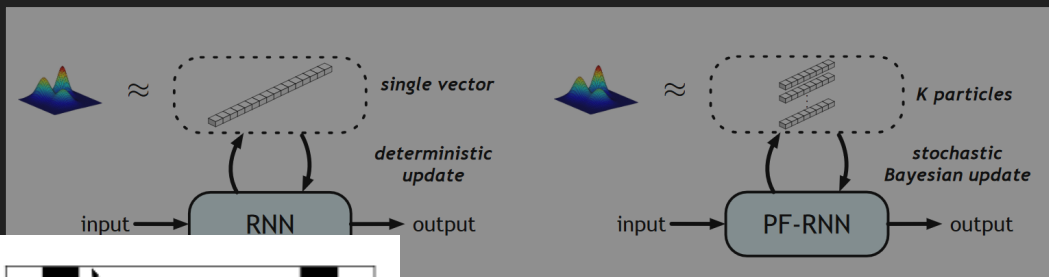
Introduction

Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

thus increasing the number of network parameters and amount of data required for training.

We introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new family of RNNs that approximate belief without lengthening h_t , thus reducing the data required for learning (Del Moral 1996) is a model-based algorithm. It approximates the belief as a set of particles that typically have well-understood meaning from particle filtering: the idea of approximating belief as a set of weighted particles, and the powerful approximation capacity of particles. PF-RNNs approximate the variable and multi-modal distribution of weighted latent vectors $\{h^1, h^2, \dots\}$ with the same distribution. Like standard RNNs, PF-RNNs follow a model-free approach: PF-RNNs’ latent vectors are learned distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian based filters, e.g., Kalman filters, particle filtering is a non-parametric approximator that offers a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2015). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,



POMDP

Further discussion

Particle Filter Recurrent Neural Networks

Xiao Ma*, Peter Karkus*, David Hsu, Wee Sun Lee
National University of Singapore
{xiao-ma, karkus, dyhsu, leews}@comp.nus.edu.sg

Abstract

Recurrent neural networks (RNNs) have been extraordinarily successful for prediction with sequential data. To tackle highly variable and multi-modal real-world data, we introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new RNN family that explicitly models uncertainty in its internal structure: while an RNN relies on a long, deterministic latent state vector, a PF-RNN maintains a latent *state distribution*, approximated as a set of particles. For effective learning, we provide a fully differentiable particle filter algorithm that updates the PF-RNN latent state distribution according to the Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.


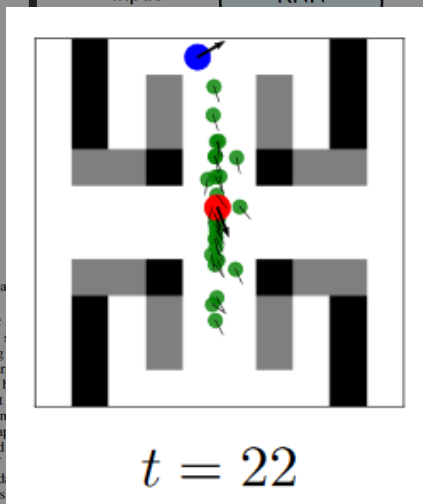
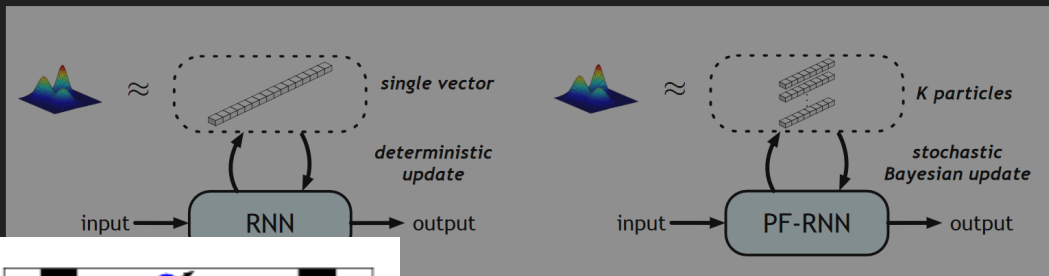
Introduction

Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

thus increasing the number of network parameters and amount of data required for training.

We introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new family of RNNs that approximate belief without lengthening h_t , thus reducing the data required for learning (Del Moral 1996) is a model-based filtering algorithm. It approximates the belief as a set of particles that typically have well-understood meaning from particle filtering: the idea of approximating belief as a set of weighted particles, and the powerful approximation capacity of particles approximates the variable and multi-modal distribution of weighted latent vectors $\{h^1, h^2, \dots\}$ with the same distribution. Like standard RNNs, PF-RNNs follow a model-free approach: PF-RNNs’ latent vectors are learned distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian based filters, e.g., Kalman filters, particle filtering is a non-parametric approximator that offers a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2015). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,



POMDP

Further discussion

Particle Filter Recurrent Neural Networks

Xiao Ma*, Peter Karkus*, David Hsu, Wee Sun Lee
National University of Singapore
{xiao-ma, karkus, dyhsu, leews}@comp.nus.edu.sg

Abstract

Recurrent neural networks (RNNs) have been extraordinarily successful for prediction with sequential data. To tackle highly variable and multi-modal real-world data, we introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new RNN family that explicitly models uncertainty in its internal structure: while an RNN relies on a long, deterministic latent state vector, a PF-RNN maintains a latent *state distribution*, approximated as a set of particles. For effective learning, we provide a fully differentiable particle filter algorithm that updates the PF-RNN latent state distribution according to the Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.



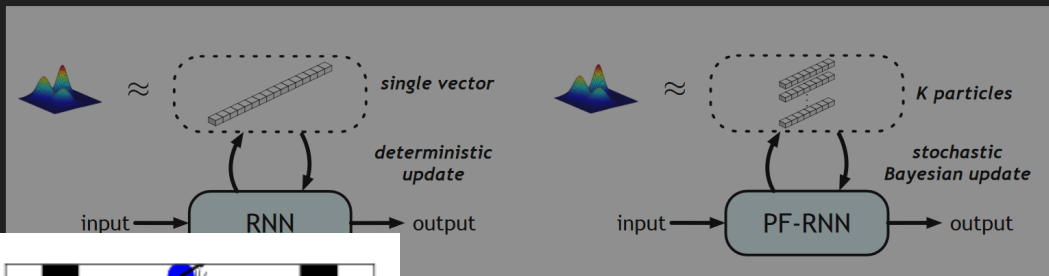
Introduction

Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

thus increasing the number of network parameters and amount of data required for training.

We introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new family of RNNs that approximate belief without lengthening h_t , thus reducing the data required for learning (Del Moral 1996) is a model-based algorithm. It approximates the belief as a set of particles that typically have well-understood meaning from particle filtering: the idea of approximating belief as a set of weighted particles, and the powerful approximation capacity of particle filtering. It approximates the variable and multi-modal distribution of weighted latent vectors $\{h^1, h^2, \dots\}$ with the same distribution. Like standard RNNs, PF-RNNs follow a model-free approach: PF-RNNs’ latent vectors are learned distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian based filters, e.g., Kalman filters, particle filtering is a non-parametric approximator that offers a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2015). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,



POMDP

Further discussion

Particle Filter Recurrent Neural Networks

Xiao Ma*, Peter Karkus*, David Hsu, Wee Sun Lee
National University of Singapore
{xiao-ma, karkus, dyhsu, leews}@comp.nus.edu.sg

Abstract

Recurrent neural networks (RNNs) have been extraordinarily successful for prediction with sequential data. To tackle highly variable and multi-modal real-world data, we introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new RNN family that explicitly models uncertainty in its internal structure: while an RNN relies on a long, deterministic latent state vector, a PF-RNN maintains a latent *state distribution*, approximated as a set of particles. For effective learning, we provide a fully differentiable particle filter algorithm that updates the PF-RNN latent state distribution according to the Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.


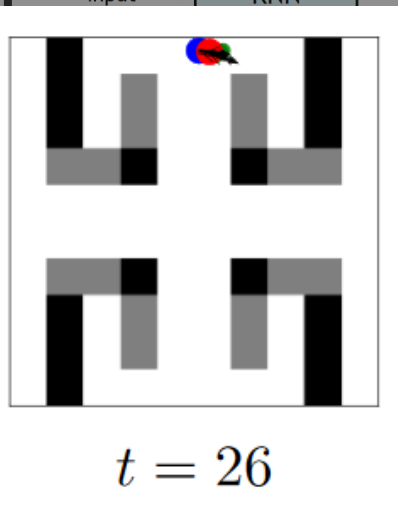
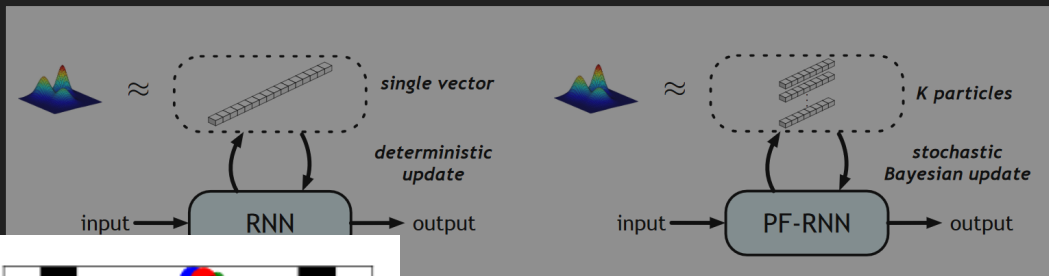
Introduction

Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

thus increasing the number of network parameters and amount of data required for training.

We introduce *Particle Filter Recurrent Neural Networks* (PF-RNNs), a new family of RNNs that approximate belief without lengthening h_t , thus reducing the data required for learning (Del Moral 1996) is a model-based filtering algorithm. It approximates the belief as a set of particles that typically have well-understood meaning from particle filtering: the idea of approximating belief as a set of weighted particles, and the powerful approximation capacity of approximating the variable and multi-modal distribution of weighted latent vectors $\{h^1, h^2, \dots\}$ with the same distribution. Like standard RNNs, PF-RNNs follow a model-free approach: PF-RNNs’ latent vectors are learned distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian based filters, e.g., Kalman filters, particle filtering is a non-parametric approximator that offers a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2015). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,






POMDP

Further discussion

Classic

With Particles



Bayes rule. Experiments demonstrate that the proposed PF-RNNs outperform the corresponding standard gated RNNs on a synthetic robot localization dataset and 10 real-world sequence prediction datasets for text classification, stock price prediction, etc.

Introduction

Prediction with sequential data is a long-standing challenge in machine learning. It has many applications, e.g., object tracking (Blake and Isard 1997), speech recognition (Xiong et al. 2018), and decision making under uncertainty (Somani et al. 2013). For effective prediction, predictors require “memory”, which summarizes and tracks information in the input sequence. The memory state is generally not observable, hence the need for a *belief*, i.e., a posterior state distribution that captures the sufficient statistic of the input for making predictions. Modeling the belief manually is often difficult. Consider the task of classifying news text—treated

as a set of weighted particles, and combine with the powerful approximation capacity of RNNs. PF-RNN approximates the variable and multi-modal belief as a set of weighted latent vectors $\{h^1, h^2, \dots\}$ sampled from the same distribution. Like standard RNNs, PF-RNNs follow a model-free approach: PF-RNNs’ latent vectors are learned distributed representations, which are not necessarily interpretable. As an alternative to the Gaussian based filters, e.g., Kalman filters, particle filtering is a non-parametric approximator that offers a more flexible belief representation (Del Moral 1996); it is also proven to give a tighter *evidence lower bound* (ELBO) in the data generation domain (Burda, Grosse, and Salakhutdinov 2015). In our case, the approximate representation is trained from data to optimize the prediction performance. For effective training with gradient methods, we employ a fully differentiable particle filter algorithm that maintains the latent belief. See Fig. 1 for a comparison of RNN and PF-RNN.

We apply the underlying idea of PF-RNN to gated RNNs,

Thank you for your attention

Discussion

- 1) Are POMDPs with a deterministic state transition a field worth more research?

VRM is a model based RL algorithm. However the results where not always very disincetive to RNNs.

- 2) How can we be confident that an algorithm truly learned a model?