

DS-GA.3001

Embodied Learning and Vision

Mengye Ren

NYU

Spring 2025

embodied-learning-vision-course.github.io



Lecture Slides for Note Taking




Logistics

- Next week will be the last lecture.
- Student topic presentation begins.

	Component 1	Component 2
Week 6 (Feb 27)	End-to-End Planning, Continual Learning (1.5 hr)	<u>Tutorial: LLM Agents</u>
Week 7 (Mar 6)	Continual Learning, Few-shot Learning (1 hr)	<u>Deep Learning for Structured Prediction</u> Kangrui Yu, Tanishq Sardana, Qing Mu, Owais Shuja
Week 8 (Mar 13)	Guest Lecture – Prof. Wei-Chiu Ma (1 hr)	<u>3D Vision and Mapping</u> Sihang Li, Kanishkha Jaisankar, Denis Mbey Akola, Zijin Hu
<u>Week 9 (Mar 20)</u>	SSL and Object Discovery Anurup Naskar, Dahye Kim, Sal Yeung, Surbhi (1.5 hr)	World Model 1 Sidhartha Reddy Potu, Andrew Deur
<i>Week 10</i>	<i>Spring Break.</i>	

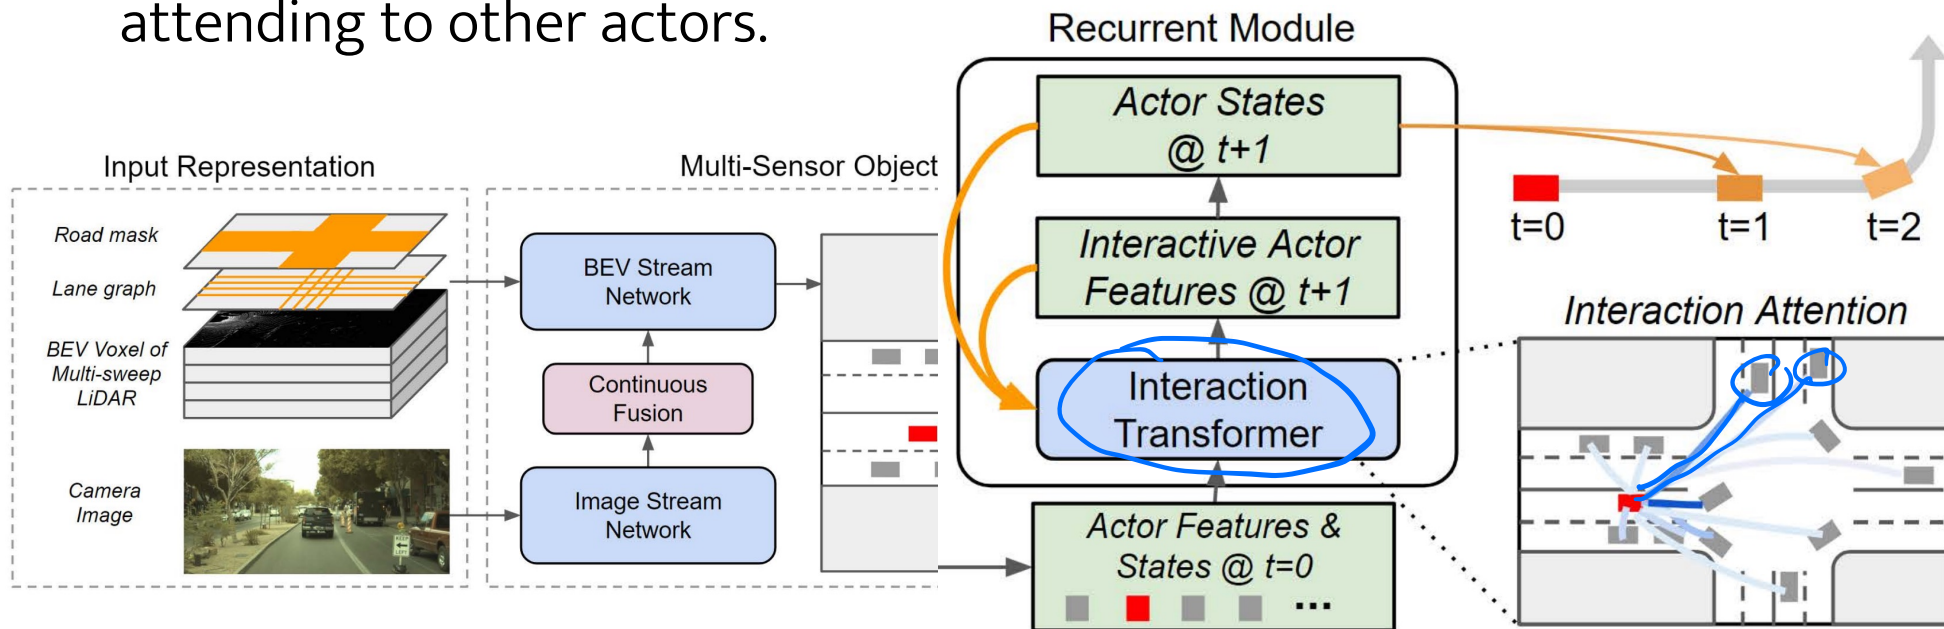
Multi-Agents Joint Prediction

- Joint predict future trajectories by attending to other actors.



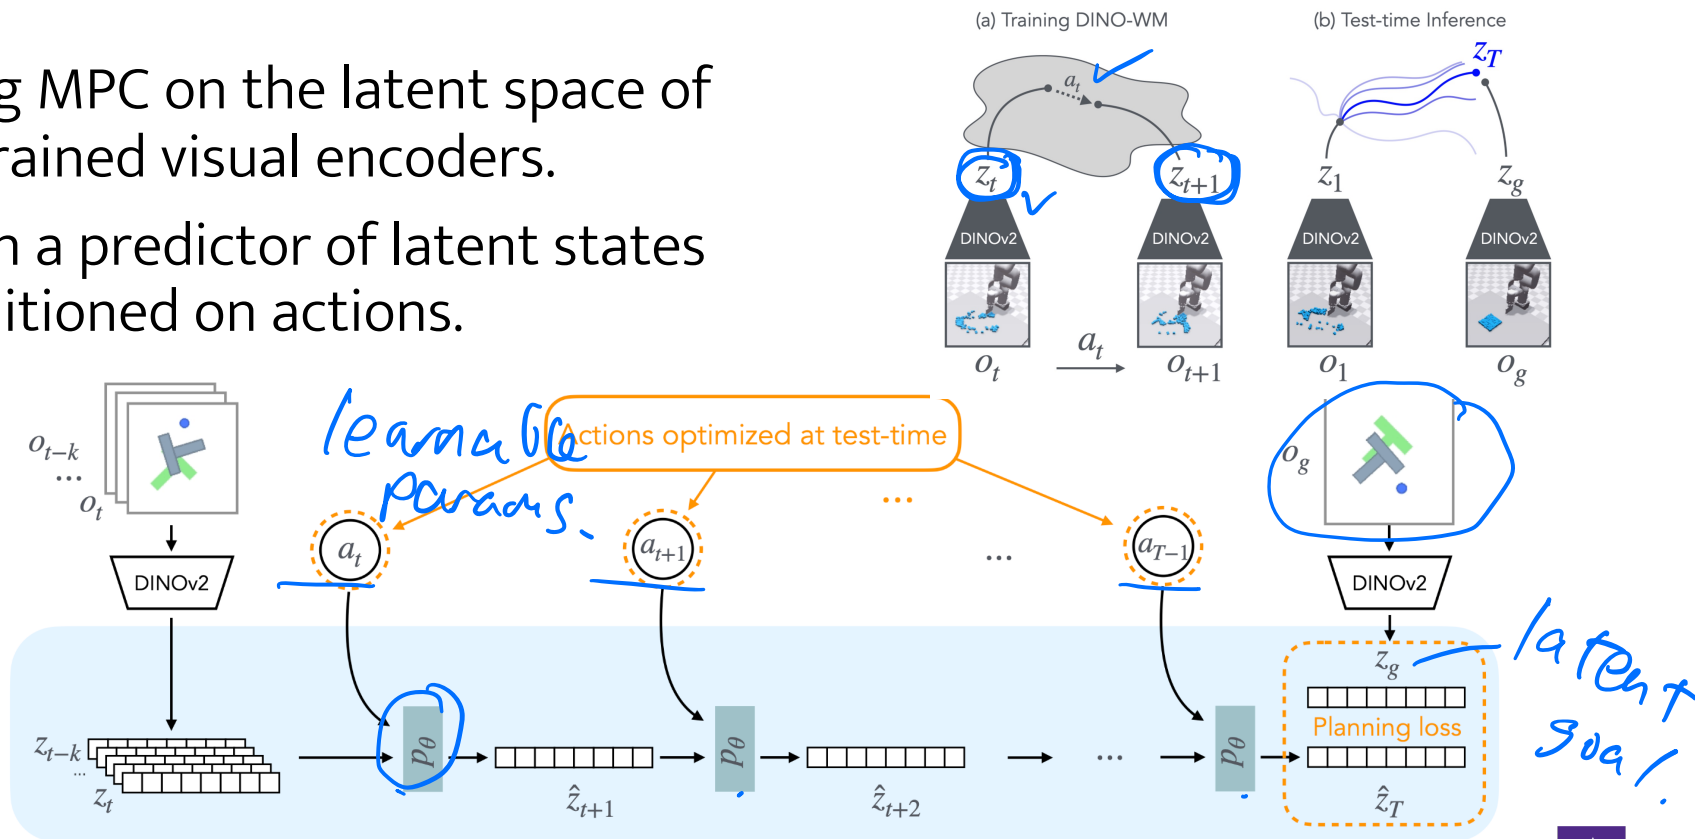
Multi-Agents Joint Prediction

- Joint predict future trajectories by attending to other actors.



Latent Prediction + MPC

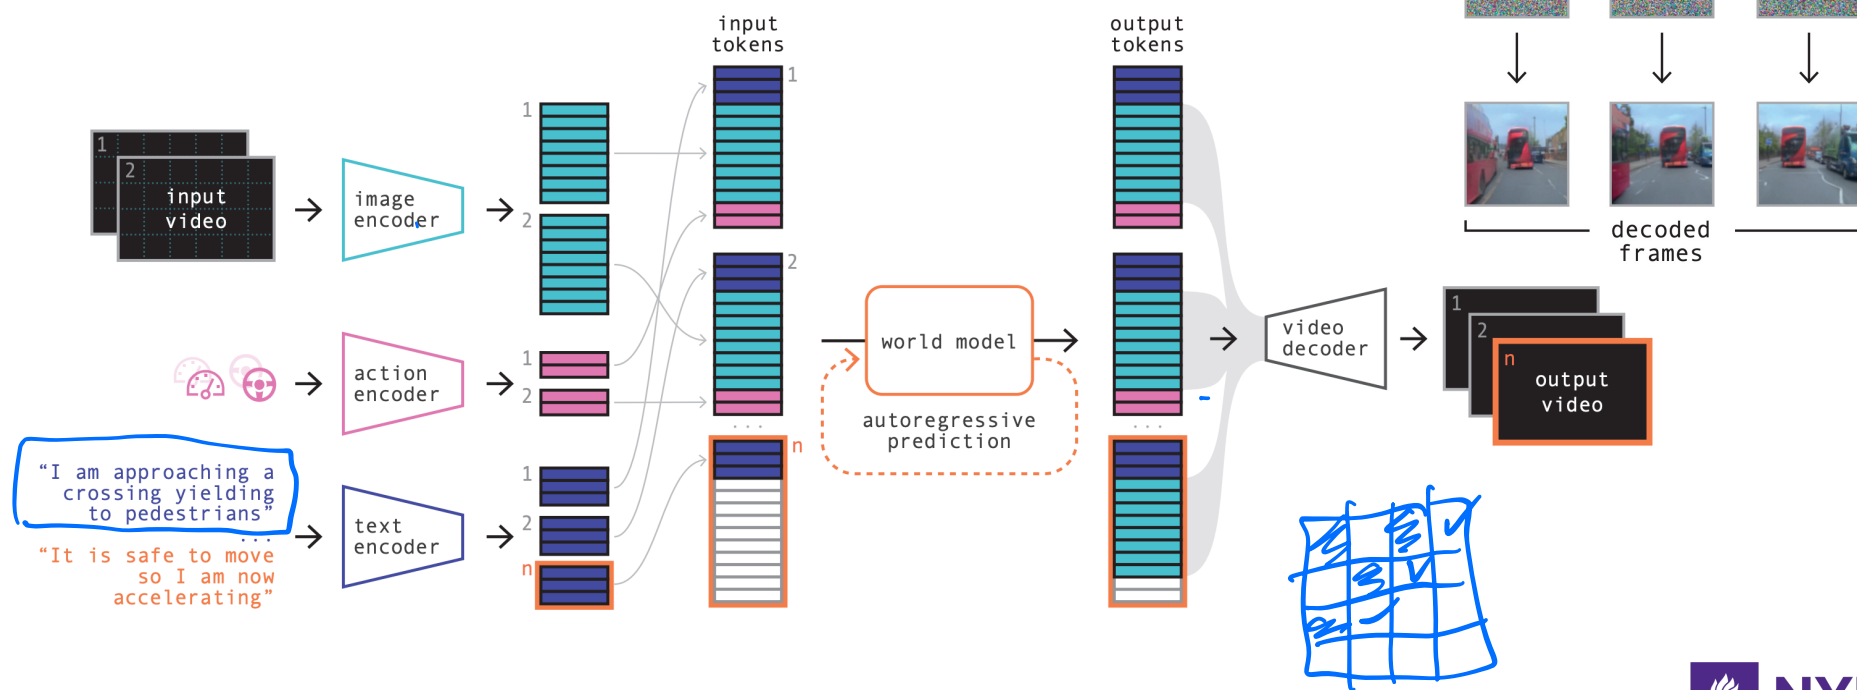
- Using MPC on the latent space of pretrained visual encoders.
- Learn a predictor of latent states conditioned on actions.




World Model in Video Prediction

Mask GPT.

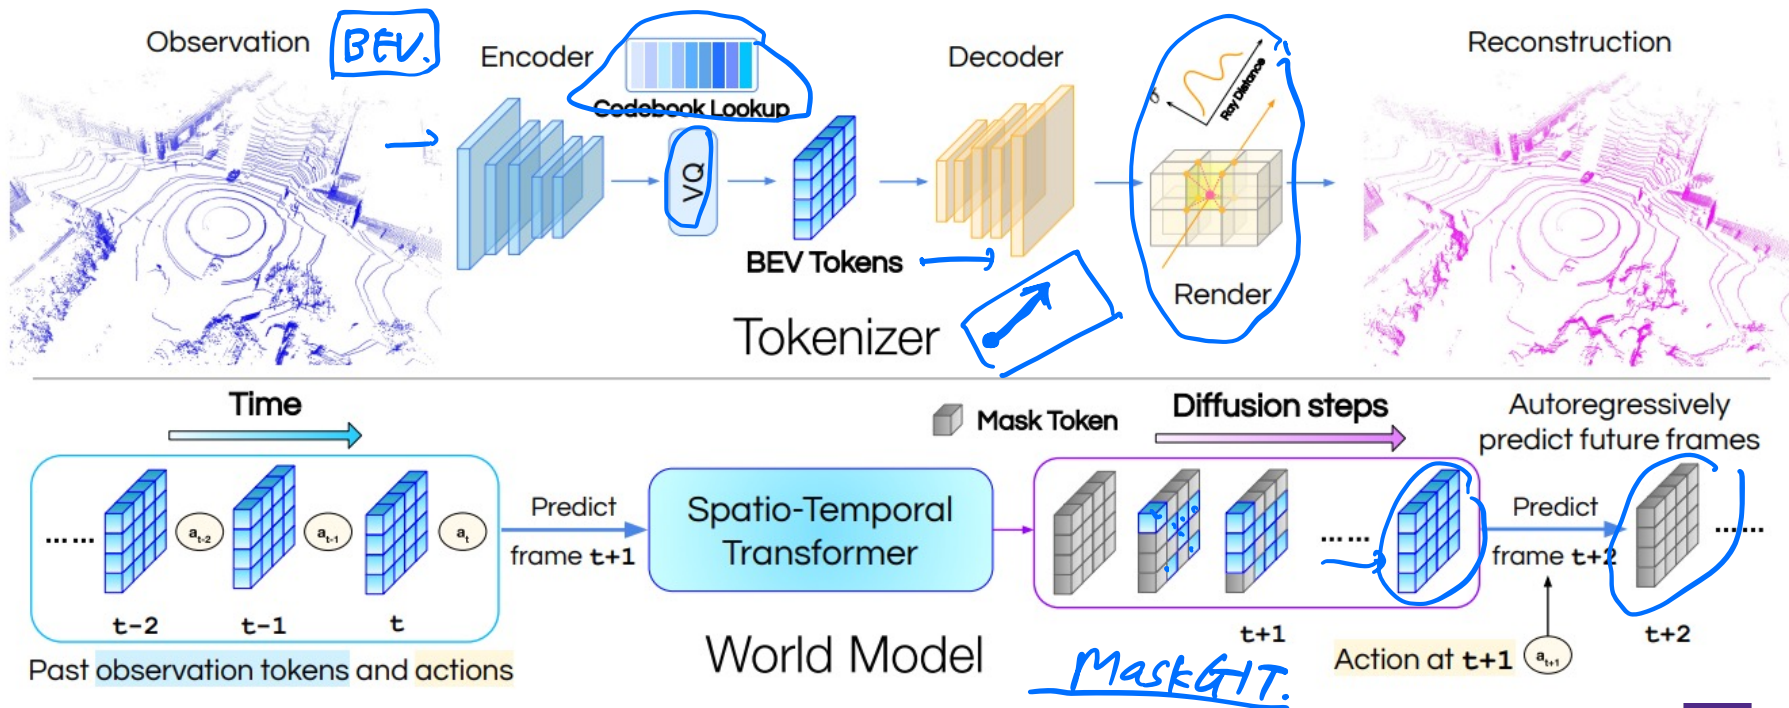
- Text+action conditioned generation. Diffusion decoder.






World Model in 3D volume prediction

- Autoregressively predict future 3D point clouds.



Summary: World Models

Summary: World Models




- Explicit object representation
 - Traditional, light weight, instance-specific, hard to learn jointly

Summary: World Models

- Explicit object representation
 - Traditional, light weight, instance-specific, hard to learn jointly
- Differentiable occupancy, motion field
 - Relatively heavy, spatially grounded, end-to-end learnable

Summary: World Models

- Explicit object representation
 - Traditional, light weight, instance-specific, hard to learn jointly
- Differentiable occupancy, motion field
 - Relatively heavy, spatially grounded, end-to-end learnable
- Global latent, RNNs, graph landmarks
 - General-purpose, unstructured



Summary: World Models

- Explicit object representation
 - Traditional, light weight, instance-specific, hard to learn jointly
- Differentiable occupancy, motion field
 - Relatively heavy, spatially grounded, end-to-end learnable
- Global latent, RNNs, graph landmarks
 - General-purpose, unstructured
- Raw video/3D prediction
 - Expensive, good for simulation

Imitation Learning

- The explicit policy model, supervised learning (behavior cloning)

$$\hat{a} = f_{\theta}(\underline{x})$$

$$\mathcal{L} = \min_i \underbrace{\|a_i - \hat{a}\|_2^2}_{MSE}$$

$$\mathcal{L} = \underbrace{-\log \hat{a}_j}$$

Imitation Learning

- The explicit policy model, supervised learning (behavior cloning)

$$\hat{a} = f_{\theta}(x) \quad \mathcal{L} = \min_i \|a_i - \hat{a}\|_2^2 \quad \mathcal{L} = -\log \hat{a}_j$$

- Energy-based (cost-based) approach

$$\tau^* = \operatorname{argmin}_{\tau} E(x, \tau)$$
$$p(\tau \mid x) = \frac{\exp(E(x, \tau))}{\int_{\tau} \exp(E(x, \tau))}$$

Imitation Learning

- The explicit policy model, supervised learning (behavior cloning)

$$\hat{a} = f_{\theta}(x) \quad \mathcal{L} = \min_i \|a_i - \hat{a}\|_2^2 \quad \mathcal{L} = -\log \hat{a}_j$$

- Energy-based (cost-based) approach

$$\tau^* = \operatorname{argmin}_{\tau} E(x, \tau)$$
$$p(\tau \mid x) = \frac{\exp(E(x, \tau))}{\int_{\tau} \exp(E(x, \tau))}$$

- Dataset Aggregation (DAgger)
 - Learned policy may deviate from experts
 - Need to collect more groundtruths

```
Initialize  $\mathcal{D} \leftarrow \emptyset$ .
Initialize  $\hat{\pi}_1$  to any policy in  $\Pi$ .
for  $i = 1$  to  $N$  do
  Let  $\pi_i = \beta_i \pi^* + (1 - \beta_i) \hat{\pi}_i$ .
  Sample  $T$ -step trajectories using  $\pi_i$ .
  Get dataset  $\mathcal{D}_i = \{(s, \pi^*(s))\}$  of visited states by  $\pi_i$ 
  and actions given by expert.
  Aggregate datasets:  $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_i$ .
  Train classifier  $\hat{\pi}_{i+1}$  on  $\mathcal{D}$ .
end for
Return best  $\hat{\pi}_i$  on validation.
```

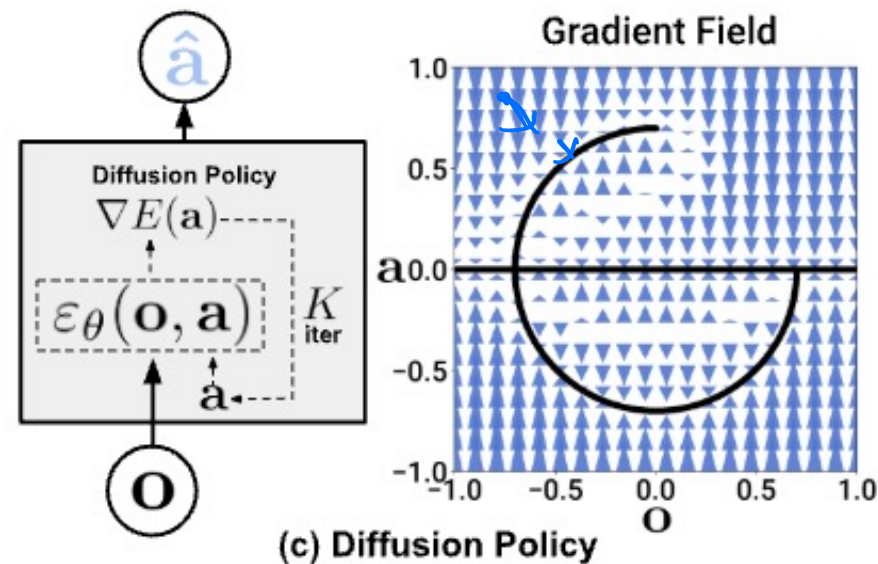
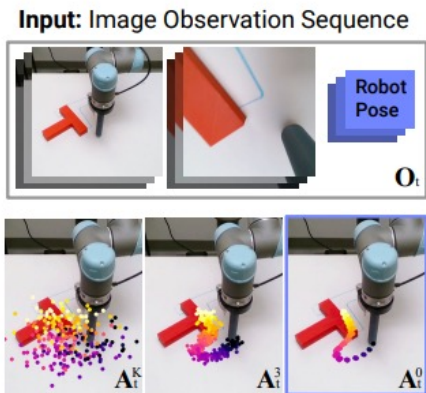

Algorithm 3.1: DAGGER Algorithm.

Direct Policy Learning from Diffusion

- Error prediction network is conditioned on observation features.


$$A_t^{k-1} = \alpha(A_t^k - \gamma \epsilon_{\theta}(O_t, A_t^k, k) + \mathcal{N}(0, \sigma^2 I)).$$

$$\mathcal{L} = \text{MSE}(\epsilon_{\theta}^k, \epsilon_{\theta}(O_t, A_t + \epsilon^k, k)).$$




Cost/Value Volume Reasoning

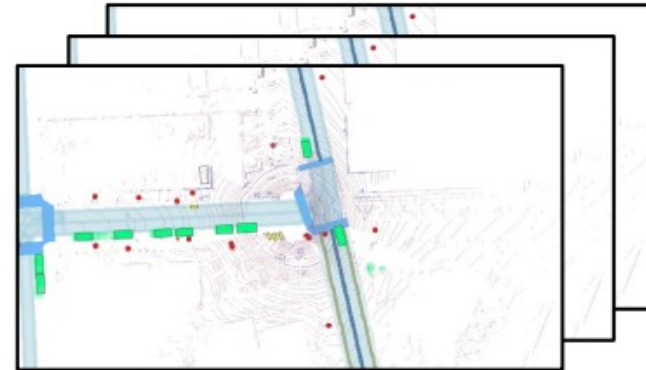
- Interpretability (both costs and planner inputs)




Rasterization



Semantic Occupancy




Cost/Value Volume Reasoning



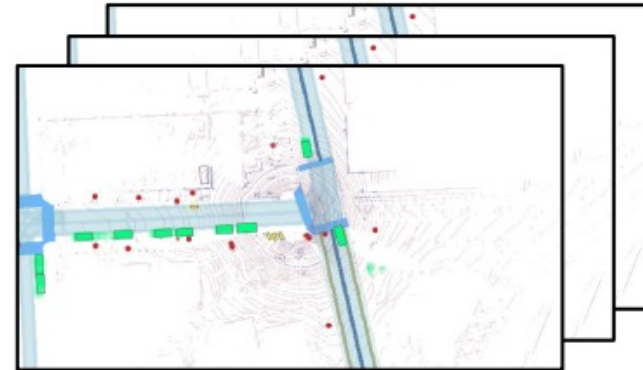
- Interpretability (both costs and planner inputs)
- Use spatial geometry to form cost from explicit objects

obstacles.

Rasterization

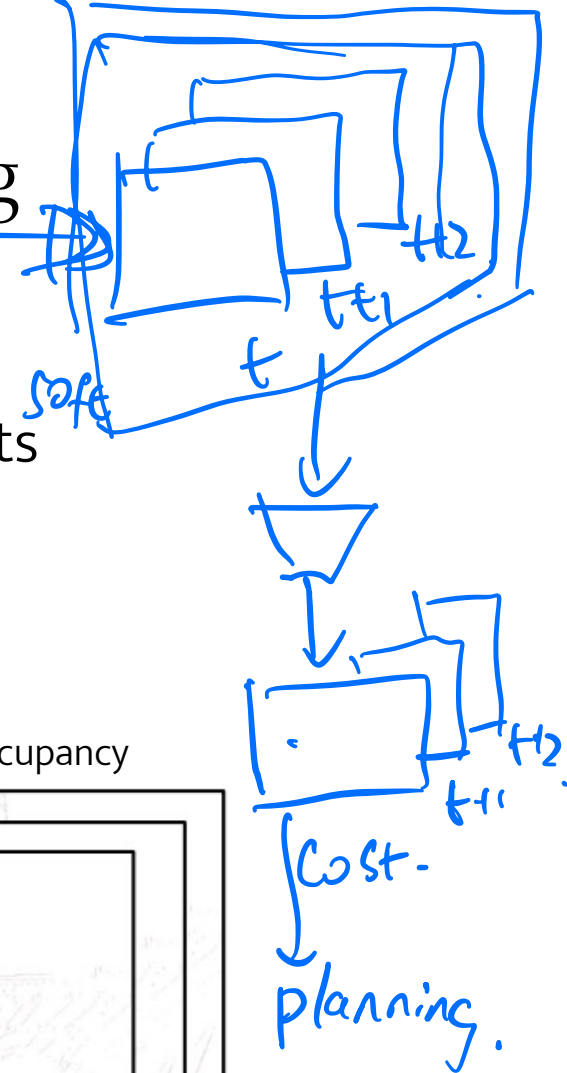
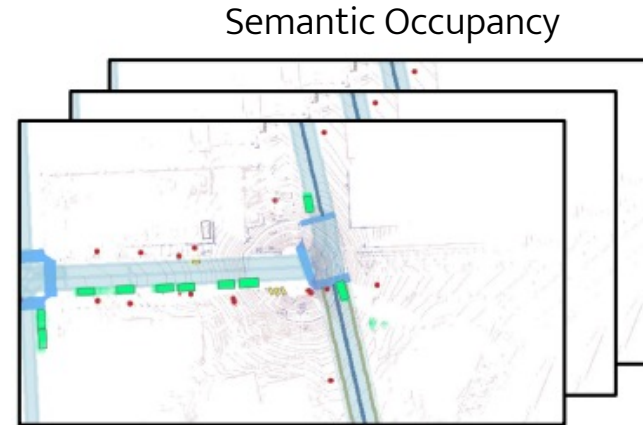
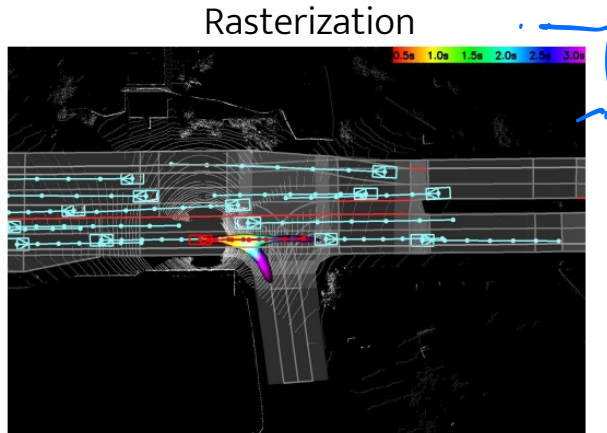


Semantic Occupancy




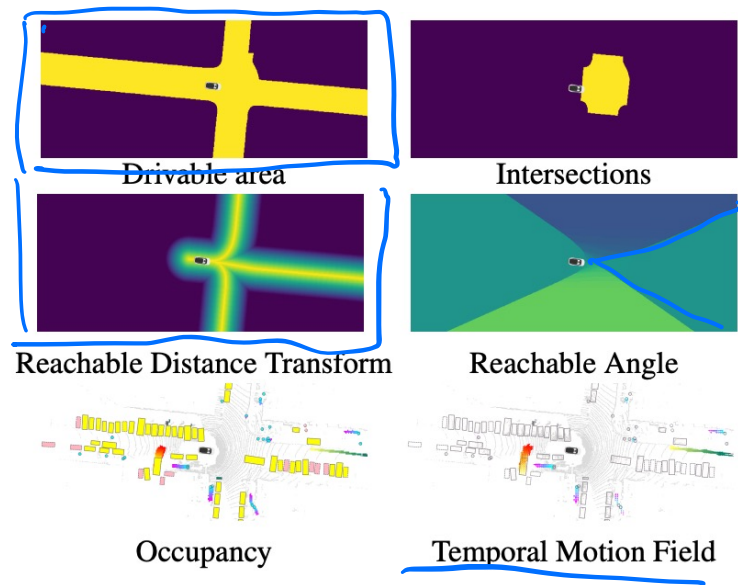

Cost/Value Volume Reasoning

- Interpretability (both costs and planner inputs)
- Use spatial geometry to form cost from explicit objects
- Predict spatial cost volume
 - Rasterize the scene for spatial inputs
 - Predict soft occupancy volumes (present and future)




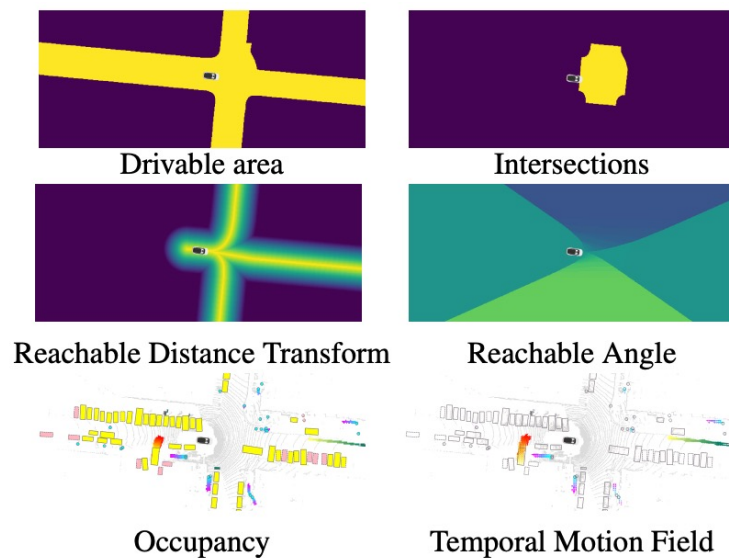
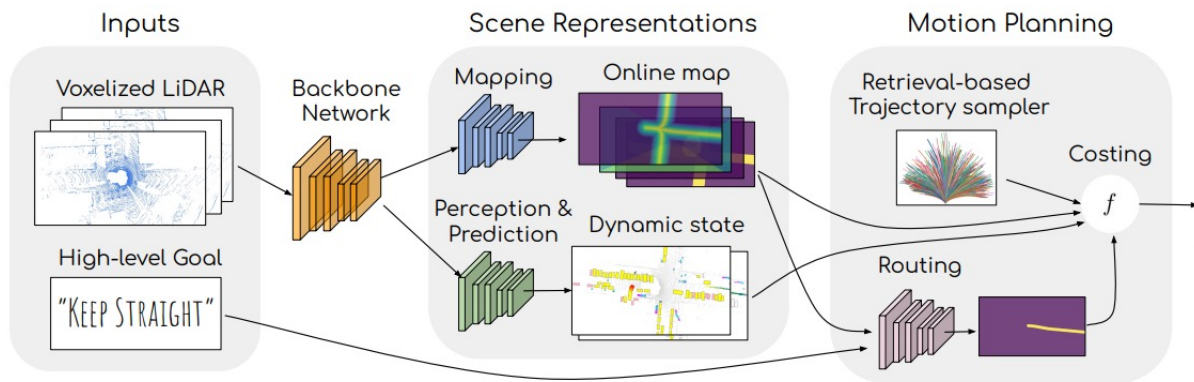
Learning Through Interpretable Predictions

- Semantic occupancy, motion field, mapping, etc. as intermediate predictions.




Learning Through Interpretable Predictions

- Semantic occupancy, motion field, mapping, etc. as intermediate predictions.
- Differentiable, supports end-to-end interpretable learning from perception to planning.



Max-Margin Planning with Explicit Cost Volume

- If we have an explicit cost volume, the cost of a trajectory can be directly queried.




Max-Margin Planning with Explicit Cost Volume

- If we have an explicit cost volume, the cost of a trajectory can be directly queried.
- We can use the max-margin objective to make the groundtruth trajectory have lower costs.

$$\underset{\theta}{\operatorname{argmin}} \sum_{\{(\hat{x}_i^t, \hat{y}_i^t)\}_{i=1 \dots N}} \max_i \sum_{t=1}^T \underbrace{C_{\theta}^t[x_t, y_t]}_{\text{expert}} - \underbrace{C_{\theta}^t[\hat{x}_i^t, \hat{y}_i^t]}_{\text{model}} + d_i^t$$

Max-Margin Planning with Explicit Cost Volume

- If we have an explicit cost volume, the cost of a trajectory can be directly queried.
- We can use the max-margin objective to make the groundtruth trajectory have lower costs.
- Find the lowest cost trajectory among a batch of samples.



$$\underset{\theta}{\operatorname{argmin}} \sum_{\{(\hat{x}_i^t, \hat{y}_i^t)\}_{i=1 \dots N}} \max_i \sum_{t=1}^T C_{\theta}^t[x_t, y_t] - \underline{C_{\theta}^t[\hat{x}_i^t, \hat{y}_i^t]} + d_i^t$$

Max-Margin Planning with Explicit Cost Volume

- If we have an explicit cost volume, the cost of a trajectory can be directly queried.
- We can use the max-margin objective to make the groundtruth trajectory have lower costs.
- Find the lowest cost trajectory among a batch of samples.
- Low-dimensional/known dynamics problems: External samplers

$$\operatorname{argmin}_{\theta} \sum_{\{(\hat{x}_i^t, \hat{y}_i^t)\}_{i=1 \dots N}} \max_i \sum_{t=1}^T C_{\theta}^t[x_t, y_t] - C_{\theta}^t[\hat{x}_i^t, \hat{y}_i^t] + d_i^t$$

Max-Margin Planning with Explicit Cost Volume

- If we have an explicit cost volume, the cost of a trajectory can be directly queried.
- We can use the max-margin objective to make the groundtruth trajectory have lower costs.
- Find the lowest cost trajectory among a batch of samples.
- Low-dimensional/known dynamics problems: External samplers
- In general, needs to perform optimization (e.g. DP)

$$\operatorname{argmin}_{\theta} \sum_{\{(\hat{x}_i^t, \hat{y}_i^t)\}_{i=1 \dots N}} \max_i \sum_{t=1}^T C_{\theta}^t[x_t, y_t] - C_{\theta}^t[\hat{x}_i^t, \hat{y}_i^t] + d_i^t$$

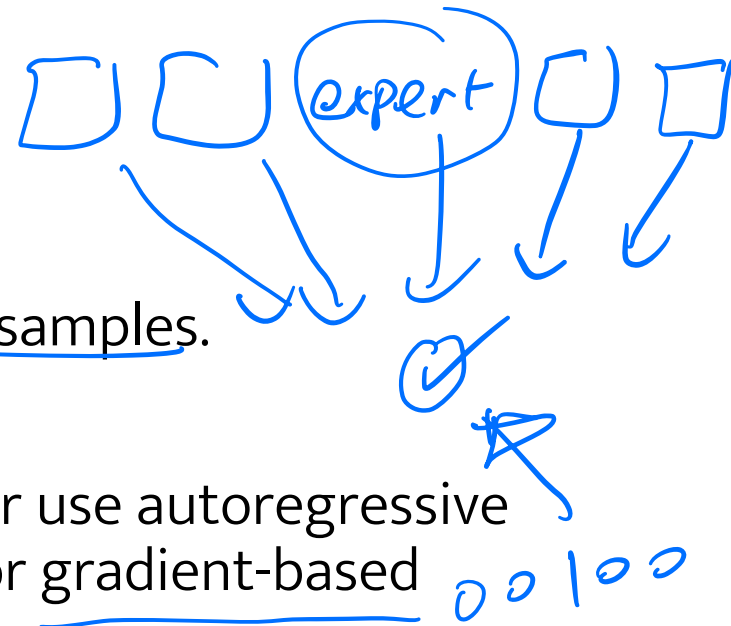
EBM Planning

- Energy-based framework also needs negative samples.

EBM Planning

- Energy-based framework also needs negative samples.
- “Pick” the groundtruth sample among others.

EBM Planning



- Energy-based framework also needs negative samples.
- “Pick” the groundtruth sample among others.
- If there isn't an external sampler, we can either use autoregressive energy of sampling one dimension at a time, or gradient-based Langevin MCMC.

EBM Planning



- Energy-based framework also needs negative samples.
- “Pick” the groundtruth sample among others.
- If there isn’t an external sampler, we can either use autoregressive energy of sampling one dimension at a time, or gradient-based Langevin MCMC.

Langevin MCMC: $\tilde{\mathbf{y}}_i^k = \tilde{\mathbf{y}}_i^{k-1} - \lambda \left(\frac{1}{2} \nabla_{\mathbf{y}} E_{\theta}(\mathbf{x}_i, \mathbf{y}_i^{k-1}) + \omega^k \right), \omega^k \sim \mathcal{N}(0, \sigma).$


EBM Planning

- Energy-based framework also needs negative samples.
- “Pick” the groundtruth sample among others.
- If there isn’t an external sampler, we can either use autoregressive energy of sampling one dimension at a time, or gradient-based Langevin MCMC.

Langevin MCMC: $\tilde{\mathbf{y}}_i^k = \tilde{\mathbf{y}}_i^{k-1} - \lambda \left(\frac{1}{2} \nabla_{\mathbf{y}} E_{\theta}(\mathbf{x}_i, \mathbf{y}_i^{k-1}) + \omega^k \right), \omega^k \sim \mathcal{N}(0, \sigma).$

Loss: $\mathcal{L} = \sum_i -\log(p_{\theta}(\mathbf{y}_i | \mathbf{x}, \{\tilde{\mathbf{y}}_i\}_j))$ $p_{\theta}(\mathbf{y}_i | \mathbf{x}, \{\tilde{\mathbf{y}}_i\}_j) = \frac{e^{-E_{\theta}(\mathbf{x}_i, \mathbf{y}_i)}}{e^{-E_{\theta}(\mathbf{x}_i, \mathbf{y}_i)} + \sum_j e^{-E_{\theta}(\mathbf{x}_i, \tilde{\mathbf{y}}_{i,j})}}$

Value Iteration Networks



- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.

Value Iteration Networks

- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.
- Classic VI algorithm:

Value Iteration Networks

- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.
- Classic VI algorithm:

$$V^*(s) = \max_{\pi} V^{\pi}(s)$$

bad policy.

Value Iteration Networks

- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.
- Classic VI algorithm:

$$V^*(s) = \max_{\pi} V^{\pi}(s)$$

$$\boxed{V^{\pi}(s)} = \mathbb{E}^{\pi} \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$$

Value Iteration Networks

- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.
- Classic VI algorithm:

$$V^*(s) = \max_{\pi} V^{\pi}(s)$$

$$V^{\pi}(s) = \mathbb{E}^{\pi} \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$$

$$Q_n(s, a) = R(s, a) + \gamma \sum_{s'} P(s'|s, a) V_n(s')$$

Value Iteration Networks

- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.
- Classic VI algorithm:

$$V^*(s) = \max_{\pi} V^{\pi}(s)$$

$$V^{\pi}(s) = \mathbb{E}^{\pi} \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$$

$$Q_n(s, a) = R(s, a) + \gamma \sum_{s'} P(s'|s, a) V_n(s')$$

$$V_{n+1}(s) = \max_a Q_n(s, a)$$

Value Iteration Networks

- A network design for predicting cost volumes that are grounded from the classic value iteration algorithm.
- Classic VI algorithm:

$$V^*(s) = \max_{\pi} V^{\pi}(s)$$


$$V^{\pi}(s) = \mathbb{E}^{\pi} \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$$

$$Q_n(s, a) = \underbrace{R(s, a)}_{\text{reward}} + \gamma \sum_{s'} P(s'|s, a) V_n(s')$$

$$V_{n+1}(s) = \max_a Q_n(s, a)$$

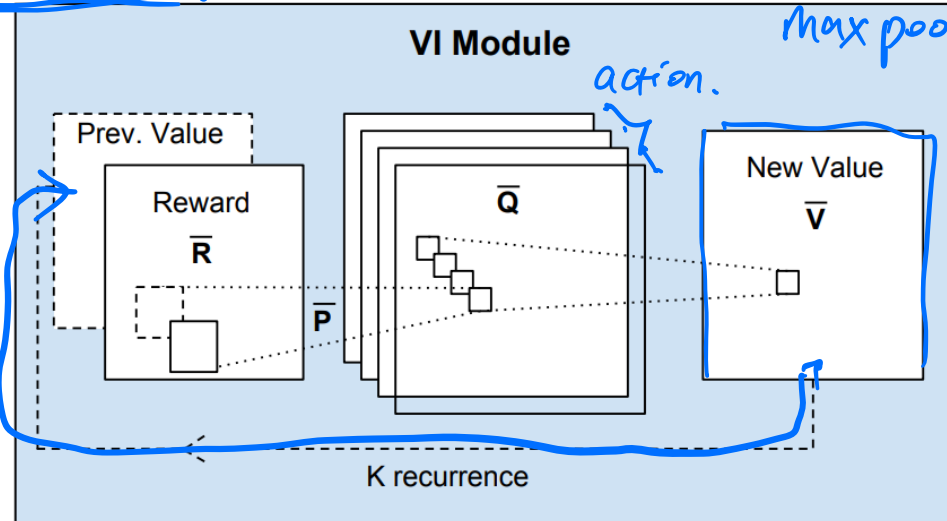
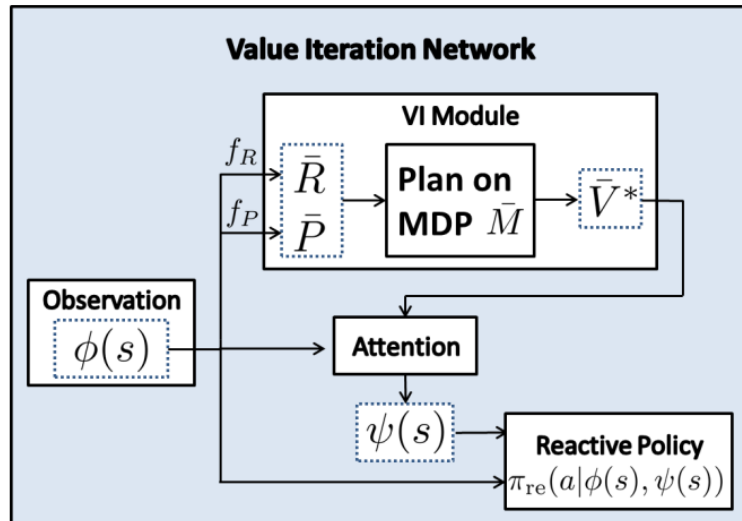
$$\pi^*(s) = \operatorname{argmax}_a Q_{\infty}(s, a)$$

Value Iteration Networks



- Reward and previous value are fed into a CNN to generate Q of A channels. Transition matrix is convolutional kernel. Then Max-Pooling.

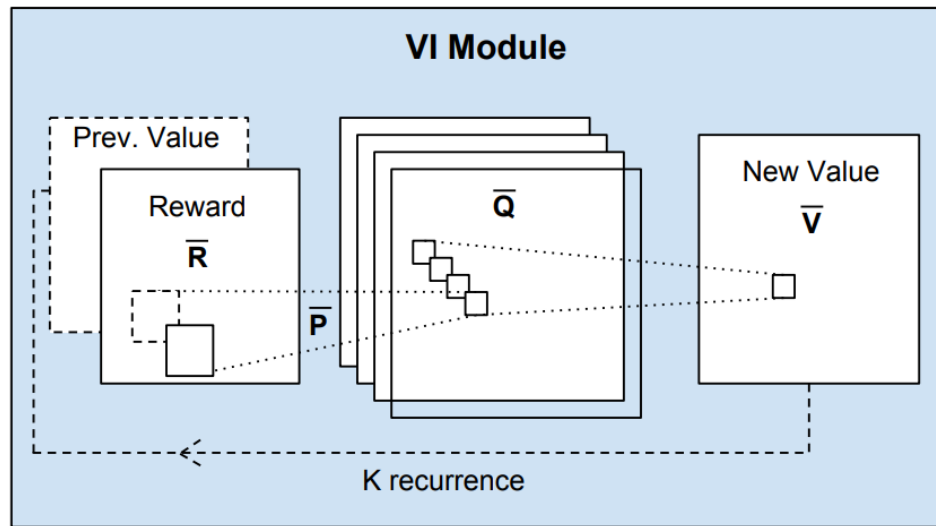
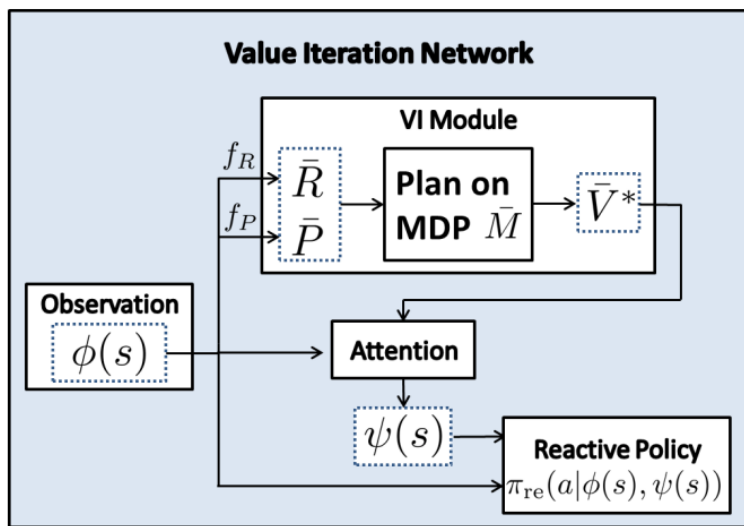
$$Q_n(s, a) = R(s, a) + \gamma \sum_{s'} P(s'|s, a) V_n(s') \quad V_{n+1}(s) = \max_a Q_n(s, a)$$



Value Iteration Networks

- Select the current state and choose an action from softmax.

$$\hat{a} \sim \text{softmax}_a(Q(s, a))$$



Value Iteration Networks


- A baseline would be to untie the weights through iterations, more like a feedforward CNN.
- Achieve more training data efficiency by imposing the structure.

Training data	VIN			VIN Untied Weights		
	Pred. loss	Succ. rate	Traj. diff.	Pred. loss	Succ. rate	Traj. diff.
20%	0.06	98.2%	0.106	0.09	91.9%	0.094
50%	0.05	99.4%	0.018	0.07	95.2%	0.078
100%	0.05	99.3%	0.089	0.05	95.6%	0.068

VIN \rightarrow RNN,
k steps

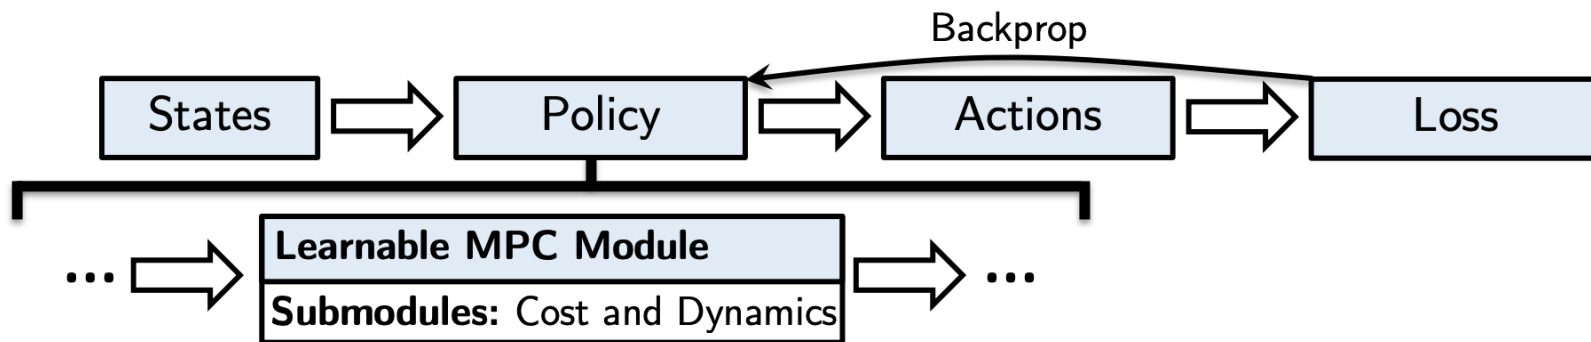
Backprop through Planning

- Treat planning as an end-to-end layer. Can be used for RL/Imitation.



Backprop through Planning

- Treat planning as an end-to-end layer. Can be used for RL/Imitation.
- Option 1: Unrolling a finite number of steps



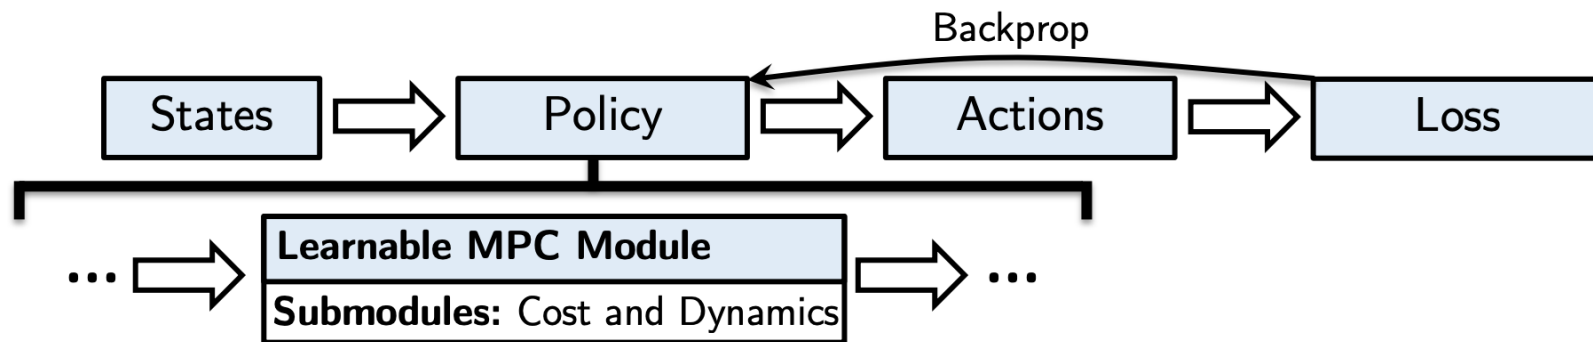
Backprop through Planning

- Treat planning as an end-to-end layer. Can be used for RL/Imitation.
- Option 1: Unrolling a finite number of steps
- Option 2: Solve till convergence, backprop for a finite step

κ


a^*

T-BPTT



Backprop through Planning

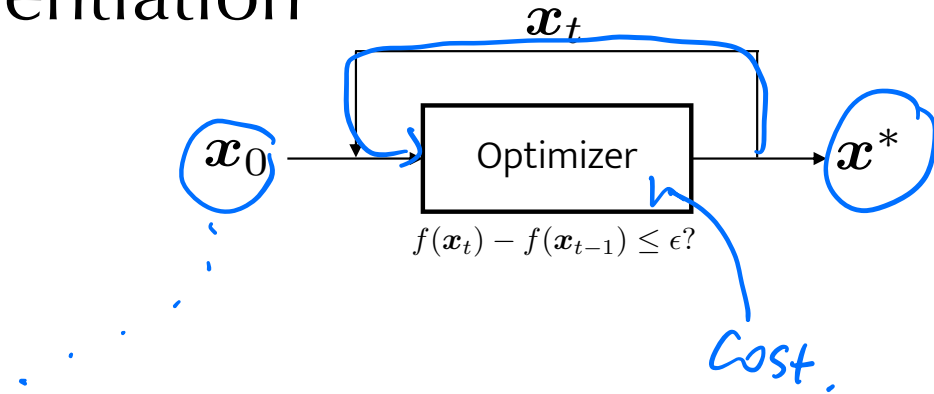
- Treat planning as an end-to-end layer. Can be used for RL/Imitation.
- Option 1: Unrolling a finite number of steps
- Option 2: Solve till convergence, backprop for a finite step
- Option 3: Converged at fixed point: Implicit differentiation



Implicit Differentiation

- Unconstrained case

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} f(\mathbf{x}; \boldsymbol{\theta}).$$



Implicit Differentiation

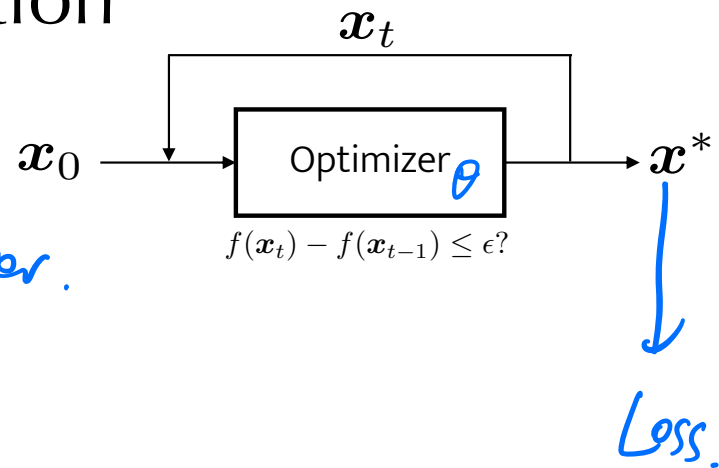
- Unconstrained case

$$x^* = \underset{x}{\operatorname{argmin}} f(x; \theta).$$

$$\frac{d}{d\theta} 0 = \frac{d}{d\theta} J_{f, x^*}(x^*; \theta)$$

cost

parameter.




Implicit Differentiation

- Unconstrained case

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} f(\mathbf{x}; \boldsymbol{\theta}).$$

$$0 = \frac{d}{d\boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$

$$0 = \frac{\partial}{\partial \mathbf{x}^*} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}) \frac{\partial \mathbf{x}^*}{\partial \boldsymbol{\theta}} + \frac{\partial}{\partial \boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$



Implicit Differentiation

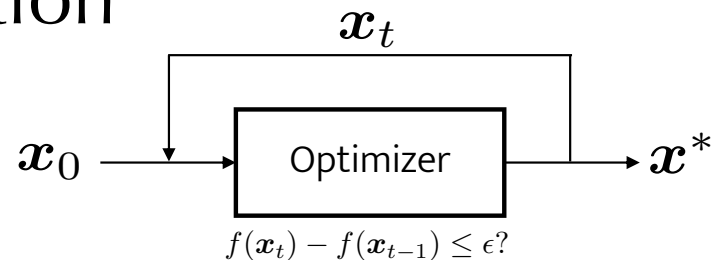
- Unconstrained case

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} f(\mathbf{x}; \boldsymbol{\theta}).$$

$$\mathbf{0} = \frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$

$$\mathbf{0} = \frac{\partial}{\partial \mathbf{x}^*} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}) \frac{\partial \mathbf{x}^*}{\partial \boldsymbol{\theta}} + \frac{\partial}{\partial \boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$

$$\mathbf{0} = H_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}) \frac{\partial \mathbf{x}^*}{\partial \boldsymbol{\theta}} + \frac{\partial}{\partial \boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$



Implicit Differentiation

- Unconstrained case


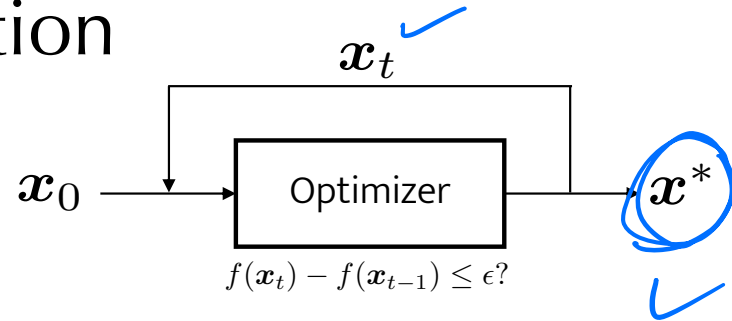
$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} f(\mathbf{x}; \boldsymbol{\theta}).$$

$$\mathbf{0} = \frac{d}{d\boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$

$$\mathbf{0} = \frac{\partial}{\partial \mathbf{x}^*} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}) \frac{\partial \mathbf{x}^*}{\partial \boldsymbol{\theta}} + \frac{\partial}{\partial \boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$

$$\mathbf{0} = H_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}) \frac{\partial \mathbf{x}^*}{\partial \boldsymbol{\theta}} + \frac{\partial}{\partial \boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta})$$

$$\frac{\partial \mathbf{x}^*}{\partial \boldsymbol{\theta}} = - \left[H_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}) \right]^{-1} \frac{\partial}{\partial \boldsymbol{\theta}} J_{f, \mathbf{x}^*}(\mathbf{x}^*; \boldsymbol{\theta}).$$



$N \times N$

→ vector.

Implicit Differentiation

- How to compute Hessian inverse vector product?

Implicit Differentiation

- How to compute Hessian inverse vector product?
- Conjugate gradient, solve

$$\boxed{Ax = b}$$

Solve x

$$x = \boxed{A^{-1}b}$$

Conjugate Gradient Method

$$\mathbf{r}_0 := \mathbf{b} - \mathbf{A}\mathbf{x}_0$$

if \mathbf{r}_0 is sufficiently small, then return \mathbf{x}_0 as the result

$$\mathbf{p}_0 := \mathbf{r}_0$$

$$k := 0$$

repeat

$$\alpha_k := \frac{\mathbf{r}_k^\top \mathbf{r}_k}{\mathbf{p}_k^\top \mathbf{A} \mathbf{p}_k}$$

$$\mathbf{x}_{k+1} := \mathbf{x}_k + \alpha_k \mathbf{p}_k$$

$$\mathbf{r}_{k+1} := \mathbf{r}_k - \alpha_k \mathbf{A} \mathbf{p}_k$$

if \mathbf{r}_{k+1} is sufficiently small, then exit loop

$$\beta_k := \frac{\mathbf{r}_{k+1}^\top \mathbf{r}_{k+1}}{\mathbf{r}_k^\top \mathbf{r}_k}$$


$$\mathbf{p}_{k+1} := \mathbf{r}_{k+1} + \beta_k \mathbf{p}_k$$

$$k := k + 1$$

end repeat

return \mathbf{x}_{k+1} as the result

Implicit Differentiation



- How to compute Hessian inverse vector product?
- Conjugate gradient, solve $Ax = b$
- Neumann series (finite truncation)

$$(I - A)^{-1} = \sum_{k=0}^{\infty} A^k$$

- Same as backprop the last K steps (Option 2).
- Memory savings.

Conjugate Gradient Method

$$\mathbf{r}_0 := \mathbf{b} - \mathbf{A}\mathbf{x}_0$$

if \mathbf{r}_0 is sufficiently small, then return \mathbf{x}_0 as the result

$$\mathbf{p}_0 := \mathbf{r}_0$$

$$k := 0$$

repeat

$$\alpha_k := \frac{\mathbf{r}_k^\top \mathbf{r}_k}{\mathbf{p}_k^\top \mathbf{A} \mathbf{p}_k}$$

$$\mathbf{x}_{k+1} := \mathbf{x}_k + \alpha_k \mathbf{p}_k$$

$$\mathbf{r}_{k+1} := \mathbf{r}_k - \alpha_k \mathbf{A} \mathbf{p}_k$$

if \mathbf{r}_{k+1} is sufficiently small, then exit loop

$$\beta_k := \frac{\mathbf{r}_{k+1}^\top \mathbf{r}_{k+1}}{\mathbf{r}_k^\top \mathbf{r}_k}$$

$$\mathbf{p}_{k+1} := \mathbf{r}_{k+1} + \beta_k \mathbf{p}_k$$


$$k := k + 1$$

end repeat

return \mathbf{x}_{k+1} as the result

Hv

Differentiable LQR



- Now add linear equality constraints on the dynamics and initialization.

$$\tau_{1:T} = \{x_t, u_t\}_{1:T}$$

$$\begin{aligned} & \underset{\tau_{1:T}}{\operatorname{argmin}} \sum_{t=1}^T \frac{1}{2} \tau_t^\top C_t \tau_t && \text{quadratic function} \leftarrow \theta. \\ & \text{subject to } \underline{x_{t+1} = F_t \tau_t + f_t, x_1 = x_{\text{init}}.} && \text{linear} \end{aligned}$$

Differentiable LQR

- Now add linear equality constraints on the dynamics and initialization.

$$\tau_{1:T} = \{x_t, u_t\}_{1:T}$$

- Chain rule: $\frac{\partial \ell}{\partial \theta} = \frac{\partial \ell}{\partial \tau_{1:T}^*} \frac{\partial \tau_{1:T}^*}{\partial \theta}$.

$$\operatorname{argmin}_{\tau_{1:T}} \sum_{t=1}^T \frac{1}{2} \tau_t^\top C_t \tau_t$$

$$\text{subject to } x_{t+1} = F_t \tau_t + f_t, x_1 = x_{\text{init}}.$$

Differentiable LQR

- Now add linear equality constraints on the dynamics and initialization.

$$\tau_{1:T} = \{x_t, u_t\}_{1:T}$$

- Chain rule: $\frac{\partial \ell}{\partial \theta} = \frac{\partial \ell}{\partial \tau_{1:T}^*} \frac{\partial \tau_{1:T}^*}{\partial \theta}$.

$$\operatorname{argmin}_{\tau_{1:T}} \sum_{t=1}^T \frac{1}{2} \tau_t^\top C_t \tau_t$$

subject to $x_{t+1} = F_t \tau_t + f_t, x_1 = x_{\text{init}}$.

General QP:

$$\mathbf{x}^* = \operatorname{argmin} \underbrace{\frac{1}{2} \mathbf{x}^\top Q \mathbf{x} + \mathbf{c}^\top \mathbf{x}}$$

subject to $A \mathbf{x} = \mathbf{b}$.

Differentiable LQR

- Now add linear equality constraints on the dynamics and initialization.

$$\tau_{1:T} = \{x_t, u_t\}_{1:T}$$

- Chain rule: $\frac{\partial \ell}{\partial \theta} = \frac{\partial \ell}{\partial \tau_{1:T}^*} \frac{\partial \tau_{1:T}^*}{\partial \theta}.$
- KKT:

$$\operatorname{argmin}_{\tau_{1:T}} \sum_{t=1}^T \frac{1}{2} \tau_t^\top C_t \tau_t$$

$$\text{subject to } x_{t+1} = F_t \tau_t + f_t, x_1 = x_{\text{init}}.$$

General QP:

$$\mathbf{x}^* = \operatorname{argmin} \frac{1}{2} \mathbf{x}^\top Q \mathbf{x} + \mathbf{c}^\top \mathbf{x}$$

$$\text{subject to } A\mathbf{x} = \mathbf{b}.$$

Differentiable LQR

- Now add linear equality constraints on the dynamics and initialization.

$$\tau_{1:T} = \{x_t, u_t\}_{1:T}$$

- Chain rule: $\frac{\partial \ell}{\partial \theta} = \frac{\partial \ell}{\partial \tau_{1:T}^*} \frac{\partial \tau_{1:T}^*}{\partial \theta}$.
- KKT:

$$\begin{bmatrix} Q & A^\top \\ A & \mathbf{0} \end{bmatrix} \begin{bmatrix} x^* \\ \lambda^* \end{bmatrix} = \begin{bmatrix} -c \\ b \end{bmatrix}$$

$$\underbrace{\begin{bmatrix} Q & A^\top \\ A & \mathbf{0} \end{bmatrix}}_{K} \begin{bmatrix} x^* \\ \lambda^* \end{bmatrix} = v.$$

$$\operatorname{argmin}_{\tau_{1:T}} \sum_{t=1}^T \frac{1}{2} \tau_t^\top C_t \tau_t$$

subject to $x_{t+1} = F_t \tau_t + f_t, x_1 = x_{\text{init}}$.

General QP:

$$x^* = \operatorname{argmin} \frac{1}{2} x^\top Q x + c^\top x$$

subject to $Ax = b$.

In classic LQR solver, the Riccati recursion solves this linear system.

Differentiable LQR

- Apply differentiation

$$\frac{d}{d\theta} \left(K \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} \right) = \frac{dv}{d\theta}.$$
$$\frac{dK}{d\theta} \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} + K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = \frac{dv}{d\theta}$$

Differentiable LQR

- Apply differentiation

$$\frac{d}{d\theta} \left(K \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} \right) = \frac{dv}{d\theta}.$$

$$\frac{dK}{d\theta} \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} + K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = \frac{dv}{d\theta}$$

$$K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} & \frac{d\mathbf{x}^*}{d\mathbf{b}} & \frac{d\mathbf{x}^*}{dQ} & \frac{d\mathbf{x}^*}{dA} \\ \frac{d\lambda^*}{d\mathbf{c}} & \frac{d\lambda^*}{d\mathbf{b}} & \frac{d\lambda^*}{dQ} & \frac{d\lambda^*}{dA} \end{bmatrix} = \begin{bmatrix} -I & 0 & -\mathbf{x}^* & -\lambda^* \\ 0 & I & 0 & -\mathbf{x}^* \end{bmatrix}$$

Differentiable LQR

- Apply differentiation

$$\frac{d}{d\theta} \left(K \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} \right) = \frac{dv}{d\theta}.$$

$$\frac{dK}{d\theta} \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} + K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = \frac{dv}{d\theta}$$

$$K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} & \frac{d\mathbf{x}^*}{d\mathbf{b}} & \frac{d\mathbf{x}^*}{dQ} & \frac{d\mathbf{x}^*}{dA} \\ \frac{d\lambda^*}{d\mathbf{c}} & \frac{d\lambda^*}{d\mathbf{b}} & \frac{d\lambda^*}{dQ} & \frac{d\lambda^*}{dA} \end{bmatrix} = \begin{bmatrix} -I & 0 & -\mathbf{x}^* & -\lambda^* \\ 0 & I & 0 & -\mathbf{x}^* \end{bmatrix}$$

$$K \frac{\partial \ell}{\partial \mathbf{z}^*} \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} & \frac{d\mathbf{x}^*}{d\mathbf{b}} \\ \frac{d\lambda^*}{d\mathbf{c}} & \frac{d\lambda^*}{d\mathbf{b}} \end{bmatrix} = \begin{bmatrix} -\frac{\partial \ell}{\partial \mathbf{x}^*} \\ \mathbf{0} \end{bmatrix}$$

$$K \mathbf{d}^* = \begin{bmatrix} -\frac{\partial \ell}{\partial \mathbf{x}^*} \\ \mathbf{0} \end{bmatrix}$$

Differentiable LQR

- Apply differentiation

$$\frac{d}{d\theta} \left(K \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} \right) = \frac{dv}{d\theta}.$$

$$\frac{dK}{d\theta} \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} + K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = \frac{dv}{d\theta}$$

$$K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} & \frac{d\mathbf{x}^*}{d\mathbf{b}} & \frac{d\mathbf{x}^*}{dQ} & \frac{d\mathbf{x}^*}{dA} \\ \frac{d\lambda^*}{d\mathbf{c}} & \frac{d\lambda^*}{d\mathbf{b}} & \frac{d\lambda^*}{dQ} & \frac{d\lambda^*}{dA} \end{bmatrix} = \begin{bmatrix} -I & 0 & -\mathbf{x}^* & -\lambda^* \\ 0 & I & 0 & -\mathbf{x}^* \end{bmatrix} K \frac{\partial \ell}{\partial \mathbf{z}^*} \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} & \frac{d\mathbf{x}^*}{d\mathbf{b}} \\ \frac{d\lambda^*}{d\mathbf{c}} & \frac{d\lambda^*}{d\mathbf{b}} \end{bmatrix} = \begin{bmatrix} -\frac{\partial \ell}{\partial \mathbf{x}^*} \\ \mathbf{0} \end{bmatrix}$$

$$K d^* = \begin{bmatrix} -\frac{\partial \ell}{\partial \mathbf{x}^*} \\ \mathbf{0} \end{bmatrix}$$

- Equivalent **QP**:

$$d^* = \underset{d}{\operatorname{argmin}} \frac{1}{2} d^\top Q d + \frac{\partial \ell}{\partial \mathbf{x}^*}^\top d, \quad \text{subject to } A d = \mathbf{0}.$$

LQR

Differentiable LQR

- Apply differentiation

$$\frac{d}{d\theta} \left(K \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} \right) = \frac{dv}{d\theta}.$$

$$\frac{dK}{d\theta} \begin{bmatrix} \mathbf{x}^* \\ \lambda^* \end{bmatrix} + K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = \frac{dv}{d\theta}$$

$$K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\lambda^*}{d\theta} \end{bmatrix} = K \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} & \frac{d\mathbf{x}^*}{d\mathbf{b}} & \frac{d\mathbf{x}^*}{dQ} & \frac{d\mathbf{x}^*}{dA} \\ \frac{d\lambda^*}{d\mathbf{c}} & \frac{d\lambda^*}{d\mathbf{b}} & \frac{d\lambda^*}{dQ} & \frac{d\lambda^*}{dA} \end{bmatrix} = \begin{bmatrix} -I & 0 & -\mathbf{x}^* & -\lambda^* \\ 0 & I & 0 & -\mathbf{x}^* \end{bmatrix} K \frac{\partial \ell}{\partial \mathbf{z}^*} \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{c}} \\ \frac{d\lambda^*}{d\mathbf{c}} \end{bmatrix} \begin{bmatrix} \frac{d\mathbf{x}^*}{d\mathbf{b}} \\ \frac{d\lambda^*}{d\mathbf{b}} \end{bmatrix} = \begin{bmatrix} -\frac{\partial \ell}{\partial \mathbf{x}^*} \\ \mathbf{0} \end{bmatrix}$$

$$K d^* = \begin{bmatrix} -\frac{\partial \ell}{\partial \mathbf{x}^*} \\ \mathbf{0} \end{bmatrix}$$

- Equivalent QP:

$$d^* = \underset{d}{\operatorname{argmin}} \frac{1}{2} d^\top Q d + \frac{\partial \ell}{\partial \mathbf{x}^*}^\top d,$$

$$\frac{\partial \ell}{\partial Q} = \frac{1}{2} (\underline{d_x^*} \otimes \mathbf{x}^* + \mathbf{x}^* \otimes \underline{d_x^*})$$

$$\frac{\partial \ell}{\partial A} = d_\lambda^* \otimes \underline{\mathbf{x}^*} + \underline{\lambda^*} \otimes d_x^*.$$

subject to $Ad = \mathbf{0}$.

Differentiable LQR

- The backward pass can also be formulated as a LQR problem.
- Swap c to $\nabla_{\tau^*} \ell$ and f to 0.

Module 1 Differentiable LQR

(The LQR algorithm is defined in [appendix A](#))

Input: Initial state x_{init}

Parameters: $\theta = \{C, c, F, f\}$

Forward Pass:

1: $\tau_{1:T}^* = \text{LQR}_T(x_{\text{init}}, C, c, F, f)$

▷ Solve (2)

2: Compute $\lambda_{1:T}^*$ with (7)

Backward Pass:

1: $d_{\tau_{1:T}}^* = \text{LQR}_T(0; C, \nabla_{\tau^*} \ell, F, 0)$ ▷ Solve (9), ideally reusing the factorizations from the forward pass

2: Compute $d_{\lambda_{1:T}}^*$ with (7)

3: Compute the derivatives of ℓ with respect to C, c, F, f , and x_{init} with (8)

Differentiable MPC

- What about general MPC?

$$\begin{aligned} & \underset{x_{1:T} \in \mathcal{X}, u_{1:T} \in \mathcal{U}}{\operatorname{argmin}} \sum_{t=1}^T C_t(x_t, u_t) \\ & \text{subject to } x_{t+1} = f(x_t, u_t), x_1 = x_{\text{init}}. \end{aligned}$$

Differentiable MPC

- What about general MPC?

$$\operatorname{argmin}_{x_{1:T} \in \mathcal{X}, u_{1:T} \in \mathcal{U}} \sum_{t=1}^T C_t(x_t, u_t)$$

subject to $x_{t+1} = f(x_t, u_t), x_1 = x_{\text{init}}$.

- Use Taylor expansion to approximate.

$$\tilde{C}_{\theta,t}^i = C_{\theta,t}(\tau_t^i) + p_t^{i\top} (\tau_t - \tau_t^i) + \frac{1}{2} (\tau_t - \tau_t^i)^\top H_t^i (\tau_t - \tau_t^i).$$

Differentiable MPC

- What about general MPC?

$$\operatorname{argmin}_{x_{1:T} \in \mathcal{X}, u_{1:T} \in \mathcal{U}} \sum_{t=1}^T C_t(x_t, u_t)$$

subject to $x_{t+1} = f(x_t, u_t), x_1 = x_{\text{init}}$.

- Use Taylor expansion to approximate.

$$\tilde{C}_{\theta,t}^i = C_{\theta,t}(\tau_t^i) + p_t^{i\top} (\tau_t - \tau_t^i) + \frac{1}{2} (\tau_t - \tau_t^i)^\top H_t^i (\tau_t - \tau_t^i).$$

- Fixed point iteration.

$$\tau^{i+1} = \operatorname{argmin}_{\tau} \sum_t^T \tilde{C}_t^i(\tau_t^i).$$

Differentiable MPC

- What about general MPC?

$$\begin{aligned} & \underset{x_{1:T} \in \mathcal{X}, u_{1:T} \in \mathcal{U}}{\operatorname{argmin}} \sum_{t=1}^T C_t(x_t, u_t) \\ & \text{subject to } x_{t+1} = f(x_t, u_t), x_1 = x_{\text{init}}. \end{aligned}$$


- Use Taylor expansion to approximate.

$$\tilde{C}_{\theta,t}^i = C_{\theta,t}(\tau_t^i) + p_t^{i\top} (\tau_t - \tau_t^i) + \frac{1}{2} (\tau_t - \tau_t^i)^\top H_t^i (\tau_t - \tau_t^i).$$

- Fixed point iteration.


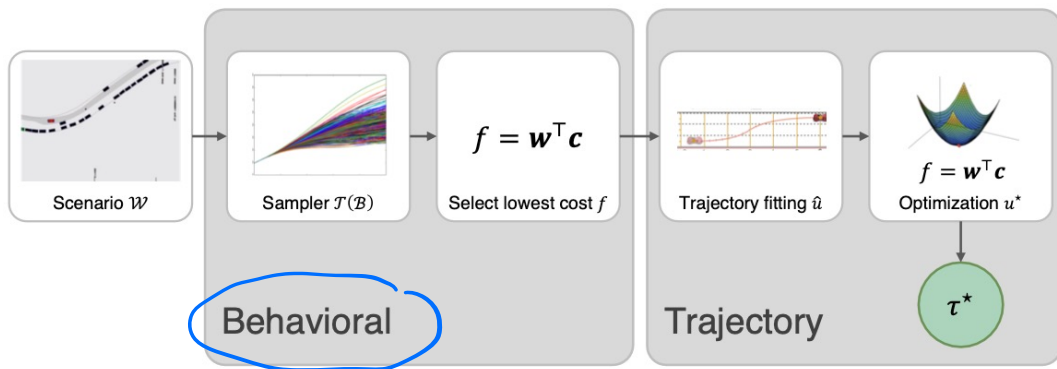
$$\tau^{i+1} = \underset{\tau}{\operatorname{argmin}} \sum_t^T \tilde{C}_t^i(\tau_t^i).$$

- Backward only depends on final quadratic approximation.






Behavioral vs. Trajectory Planning

- Gradient-based optimization provides a locally optimized trajectory.



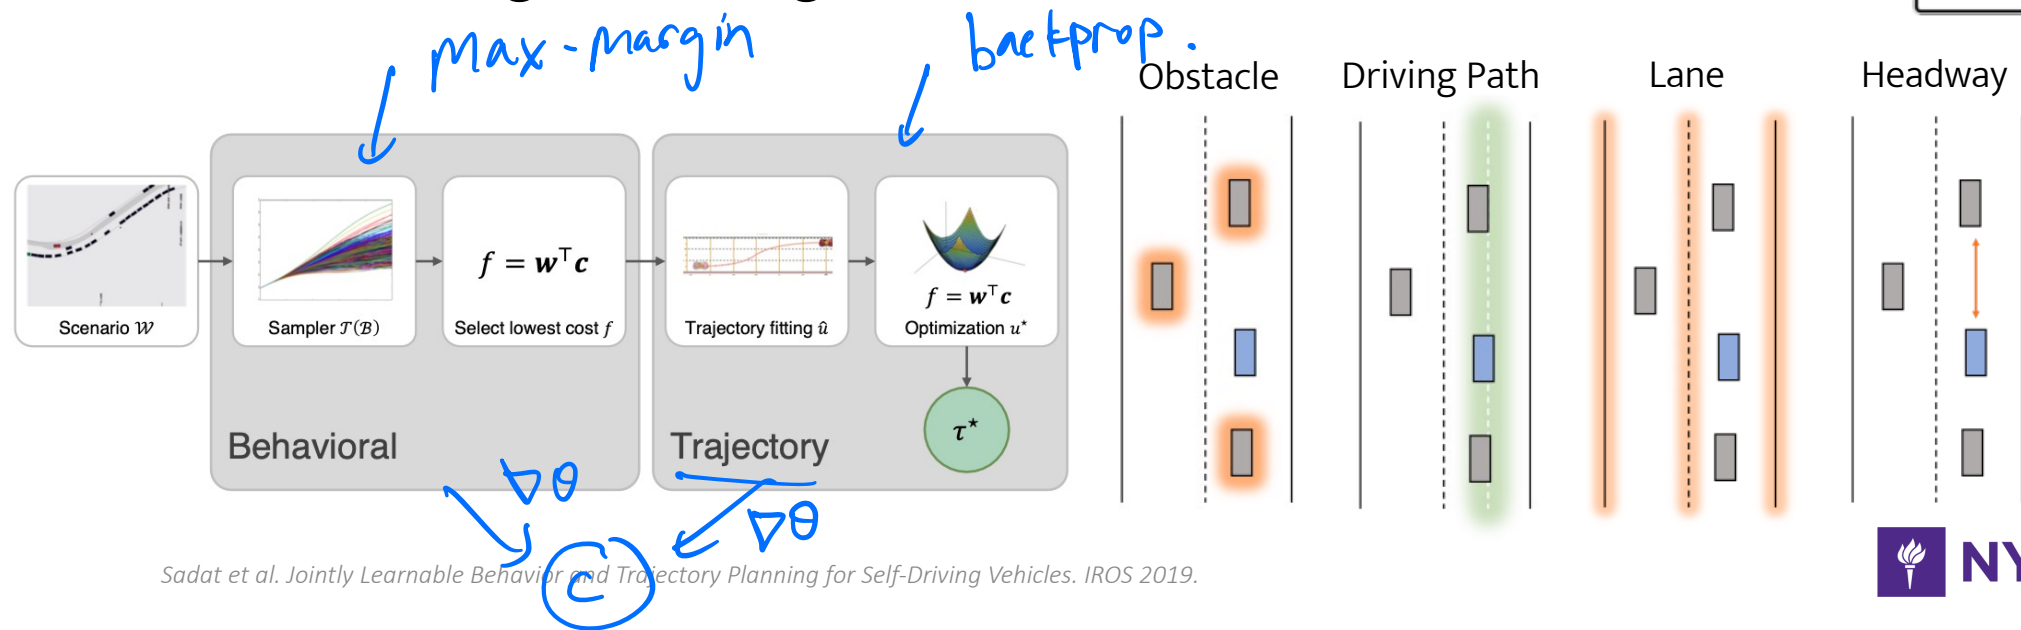
Behavioral vs. Trajectory Planning

- Gradient-based optimization provides a locally optimized trajectory.
- Samples may be needed for reasoning global structure.



Behavioral vs. Trajectory Planning

- Gradient-based optimization provides a locally optimized trajectory.
- Samples may be needed for reasoning global structure.
- Can learn together using the same learned costs.




Planning with Social Reasoning



- Jointly reason the future trajectories of multiple agents as an energy-based graphical model.

$$p(\mathbf{s}_1, \dots, \mathbf{s}_N \mid \mathbf{X}) = \frac{1}{Z} \exp(-E_\theta(\mathbf{s}_1, \dots, \mathbf{s}_N \mid \mathbf{X}))$$



Planning with Social Reasoning


- Jointly reason the future trajectories of multiple agents as an energy-based graphical model.

$$p(\mathbf{s}_1, \dots, \mathbf{s}_N \mid \mathbf{X}) = \frac{1}{Z} \exp(-E_\theta(\mathbf{s}_1, \dots, \mathbf{s}_N) \mid \mathbf{X})$$

- Trajectory Goodness + Collision.

$$\sum_i E_\theta(\mathbf{s}_i \mid \mathbf{X}) + \sum_{i \neq j} E(\mathbf{s}_i, \mathbf{s}_j)$$

$$E(\mathbf{s}_i, \mathbf{s}_j) = \gamma \text{ if } \mathbf{s}_i \text{ collides } \mathbf{s}_j$$



Planning with Social Reasoning


- Jointly reason the future trajectories of multiple agents as an energy-based graphical model.

$$p(\mathbf{s}_1, \dots, \mathbf{s}_N \mid \mathbf{X}) = \frac{1}{Z} \exp(-E_\theta(\mathbf{s}_1, \dots, \mathbf{s}_N) \mid \mathbf{X})$$

- Trajectory Goodness + Collision.
- Batch of trajectory samples.

$$\sum_i E_\theta(s_i \mid X) + \sum_{i \neq j} E(s_i, s_j)$$

$$E(s_i, s_j) = \gamma \text{ if } s_i \text{ collides } s_j$$



Planning with Social Reasoning

- Jointly reason the future trajectories of multiple agents as an energy-based graphical model.

$$p(\mathbf{s}_1, \dots, \mathbf{s}_N \mid \mathbf{X}) = \frac{1}{Z} \exp(-E_\theta(\mathbf{s}_1, \dots, \mathbf{s}_N) \mid \mathbf{X})$$


- Trajectory Goodness + Collision.

$$\sum_i E_\theta(s_i \mid X) + \sum_{i \neq j} E(s_i, s_j)$$

- Batch of trajectory samples.

$$E(s_i, s_j) = \gamma \text{ if } s_i \text{ collides } s_j$$

- Classification of groundtruth trajectory.



Summary: End-to-End Planning

- Direct Policy Prediction
 - Condition perception features into the model
 - Use of diffusion models

Summary: End-to-End Planning

- Direct Policy Prediction
 - Condition perception features into the model
 - Use of diffusion models
- Cost Learning (IRL) from Experts
 - Max-margin, max-entropy/EBM
 - Need negative samples
 - Can be combined with efficient external samplers
 - Cost volume prediction: parametric + non-parametric

Summary: End-to-End Planning

- Direct Policy Prediction
 - Condition perception features into the model
 - Use of diffusion models
- Cost Learning (IRL) from Experts
 - Max-margin, max-entropy/EBM
 - Need negative samples
 - Can be combined with efficient external samplers
 - Cost volume prediction: parametric + non-parametric
- Differentiable Planner
 - Backprop through local optimization
 - Can be memory efficient, implicit differentiation

Module 5:
Continual Learning,
Few-Shot Learning, Meta-Learning

Why Continual Learning?

- The world is not a dataset that allows you to get IID samples.

Why Continual Learning?

- The world is not a dataset that allows you to get IID samples.
- The world keeps changing and evolving.

Why Continual Learning?

- The world is not a dataset that allows you to get IID samples.
- The world keeps changing and evolving.
- Online vs. Continual
 - Online means that samples arrive in a streaming / temporal partial order, but they may still come from a static distribution.

$$\theta_t = f_{\theta}(x_t, \theta_{t-1}) \quad x_{1:T} \sim \mathcal{X}$$

- Example: Online reinforcement learning, trajectory roll out is online, but the environment is the same.
- Continual learning means that there will be distribution shift.

What is Continual Learning?

- Distribution shift: Forgetting
 - Learning on A and then B, results in worse performance on A.

*backward
transfer.*

What is Continual Learning?

- Distribution shift: Forgetting
 - Learning on A and then B, results in worse performance on A.
- Multi-task learning: Forward transfer
 - Learning Task A + B results in better learning in Task C compared to learning C alone.
 - Leverage the similarity between tasks.

What is Continual Learning?



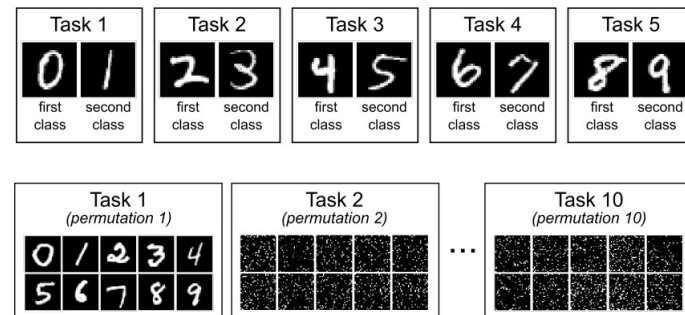
- Distribution shift: Forgetting
 - Learning on A and then B, results in worse performance on A.
- Multi-task learning: Forward transfer
 - Learning Task A + B results in better learning in Task C compared to learning C alone.
 - Leverage the similarity between tasks.
- Compositionality
 - Learning A and B first, and then learning tasks with composed A+B.

What is Continual Learning?


- ✓ Distribution shift: Forgetting ✱
 - Learning on A and then B, results in worse performance on A.
- ✓ Multi-task learning: Forward transfer
 - Learning Task A + B results in better learning in Task C compared to learning C alone.
 - Leverage the similarity between tasks.
- ✓ Compositionality
 - Learning A and B first, and then learning tasks with composed A+B.
- ✓ Incremental/curriculum Learning
 - Learning A->B->C is easier than at random order.

Continual Learning


- Learning a sequence of tasks without looking back.




(a) Pong variants



(b) Labyrinth games








(c) Atari games




Continual Learning

increment


- Learning a sequence of tasks without looking back.
- Goal is to do well on all of the tasks at the end.




(a) Pong variants



(b) Labyrinth games



(c) Atari games



button press



door open



drawer close








drawer open




peg insert side

Continual Learning


- Learning a sequence of tasks without looking back.
- Goal is to do well on all of the tasks at the end.
- Task boundary



(a) Pong variants



(b) Labyrinth games



(c) Atari games



button press



door open



drawer close




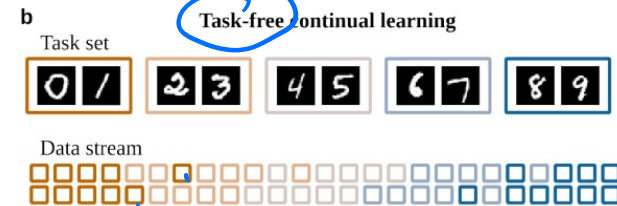



drawer open




peg insert side

Continual Learning


- Learning a sequence of tasks without looking back.
- Goal is to do well on all of the tasks at the end.
- Task boundary
- Memory constraints



(a) Pong variants



(b) Labyrinth games



(c) Atari games



button press



door open



drawer close



drawer open



peg insert side

Parameter Regularization

- Over-completeness Assumption. A multitude of models can reach equivalent performance.

$$\mathcal{S}_A = \{\theta \mid \ell_A(\theta) < \epsilon\}$$
$$\mathcal{S}_A \cap \mathcal{S}_B \neq \emptyset$$

Parameter Regularization

- Over-completeness Assumption. A multitude of models can reach equivalent performance.
- What is left is to efficiently find the intersection between A and B.

$$p(\theta \mid \mathcal{D}_A) = \mathcal{N}(\theta; \theta^*, \Sigma)$$

$$\mathcal{S}_A = \{\theta \mid \ell_A(\theta) < \epsilon\}$$

$$\mathcal{S}_A \cap \mathcal{S}_B \neq \emptyset$$

Parameter Regularization

- Over-completeness Assumption. A multitude of models can reach equivalent performance.

- What is left is to efficiently find the intersection between A and B.

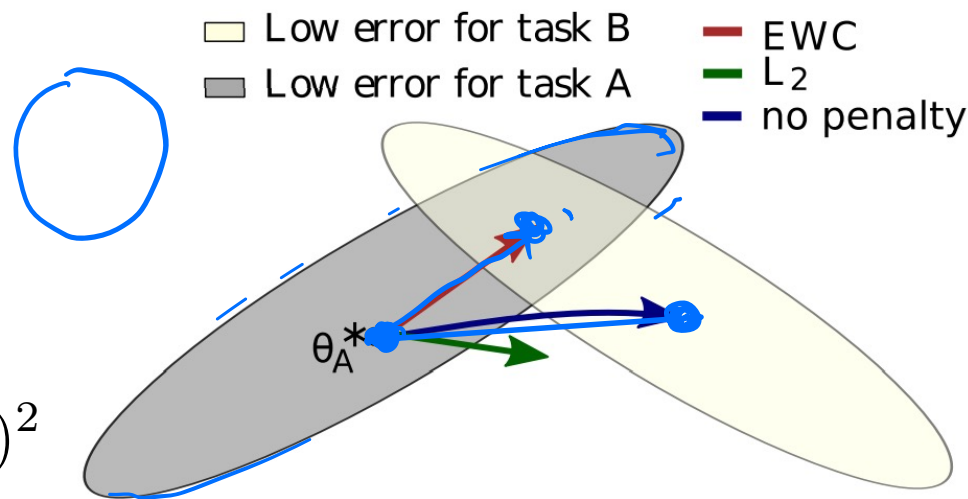
$$p(\theta \mid \mathcal{D}_A) = \mathcal{N}(\theta; \theta^*, \Sigma)$$

- Elastic Weight Consolidation (EWC):

$$\mathcal{L}(\theta) = \mathcal{L}_B(\theta) + \sum_i \frac{\lambda}{2} (F_i(\theta_i) - \theta_{A,i}^*)^2$$

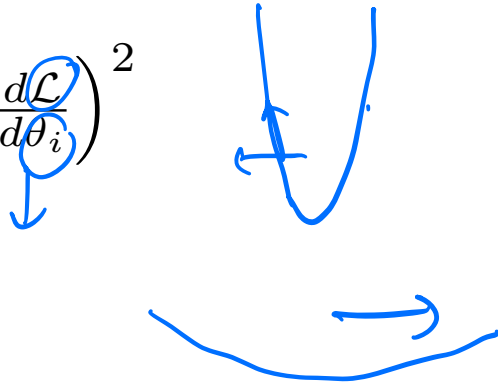
$$\mathcal{S}_A = \{\theta \mid \ell_A(\theta) < \epsilon\}$$


$$\mathcal{S}_A \cap \mathcal{S}_B \neq \emptyset$$



Computing Fisher

- At the end of each epoch, compute the gradient squared:

$$F_i = \left(\frac{d\mathcal{L}}{d\theta_i} \right)^2$$





Computing Fisher

- At the end of each epoch, compute the gradient squared:

$$F_i = \left(\frac{d\mathcal{L}}{d\theta_i} \right)^2$$

- Measures the sensitivity on each parameter dimension.




Computing Fisher

plasticity --
stability. ++

- At the end of each epoch, compute the gradient squared:

$$F_i = \left(\frac{d\mathcal{L}}{d\theta_i} \right)^2$$

- Measures the sensitivity on each parameter dimension.
- You can also accumulate an online estimate.



Variational Continual Learning (VCL)

- Bayesian formulation:

$$p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T}) \propto p(\boldsymbol{\theta}) \prod_{t=1}^T p(\mathcal{D}_t \mid \boldsymbol{\theta}) \propto p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T-1})p(\mathcal{D}_T \mid \boldsymbol{\theta}).$$

Variational Continual Learning (VCL)

- Bayesian formulation:

$$p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T}) \propto p(\boldsymbol{\theta}) \prod_{t=1}^T p(\mathcal{D}_t \mid \boldsymbol{\theta}) \propto p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T-1})p(\mathcal{D}_T \mid \boldsymbol{\theta}).$$

- Variational approach:

$$q_t(\boldsymbol{\theta}) = \operatorname{argmin}_{q \in \mathcal{Q}} \operatorname{KL} \left(q(\boldsymbol{\theta}) \parallel \frac{1}{Z_t} q_{t-1}(\boldsymbol{\theta}) p(\mathcal{D}_t \mid \boldsymbol{\theta}) \right).$$

Variational Continual Learning (VCL)

- Bayesian formulation:

$$p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T}) \propto p(\boldsymbol{\theta}) \prod_{t=1}^T p(\mathcal{D}_t \mid \boldsymbol{\theta}) \propto p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T-1})p(\mathcal{D}_T \mid \boldsymbol{\theta}).$$

- Variational approach:

$$q_t(\boldsymbol{\theta}) = \operatorname{argmin}_{q \in \mathcal{Q}} \operatorname{KL} \left(q(\boldsymbol{\theta}) \parallel \frac{1}{Z_t} q_{t-1}(\boldsymbol{\theta}) p(\mathcal{D}_t \mid \boldsymbol{\theta}) \right).$$

- Loss: $\mathcal{L}(q_t(\boldsymbol{\theta})) = \mathbb{E}_{\boldsymbol{\theta} \sim q_t(\boldsymbol{\theta})} [-\log p(\mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta})] + \operatorname{KL}(q_t(\boldsymbol{\theta}) \parallel q_{t-1}(\boldsymbol{\theta})).$

Variational Continual Learning (VCL)

- Bayesian formulation:

$$p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T}) \propto p(\boldsymbol{\theta}) \prod_{t=1}^T p(\mathcal{D}_t \mid \boldsymbol{\theta}) \propto p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T-1})p(\mathcal{D}_T \mid \boldsymbol{\theta}).$$

- Variational approach:

$$q_t(\boldsymbol{\theta}) = \operatorname{argmin}_{q \in \mathcal{Q}} \operatorname{KL} \left(q(\boldsymbol{\theta}) \parallel \frac{1}{Z_t} q_{t-1}(\boldsymbol{\theta}) p(\mathcal{D}_t \mid \boldsymbol{\theta}) \right).$$

- Loss: $\mathcal{L}(q_t(\boldsymbol{\theta})) = \mathbb{E}_{\boldsymbol{\theta} \sim q_t(\boldsymbol{\theta})} [-\log p(\mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta})] + \operatorname{KL}(q_t(\boldsymbol{\theta}) \parallel q_{t-1}(\boldsymbol{\theta})).$
 $q_t(\boldsymbol{\theta}) = \prod_{d=1}^D \mathcal{N}(\theta_{t,d}; \mu_{t,d}, \sigma_{t,d}^2).$

Variational Continual Learning (VCL)

- Bayesian formulation:

$$p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T}) \propto p(\boldsymbol{\theta}) \prod_{t=1}^T p(\mathcal{D}_t \mid \boldsymbol{\theta}) \propto p(\boldsymbol{\theta} \mid \mathcal{D}_{1:T-1})p(\mathcal{D}_T \mid \boldsymbol{\theta}).$$

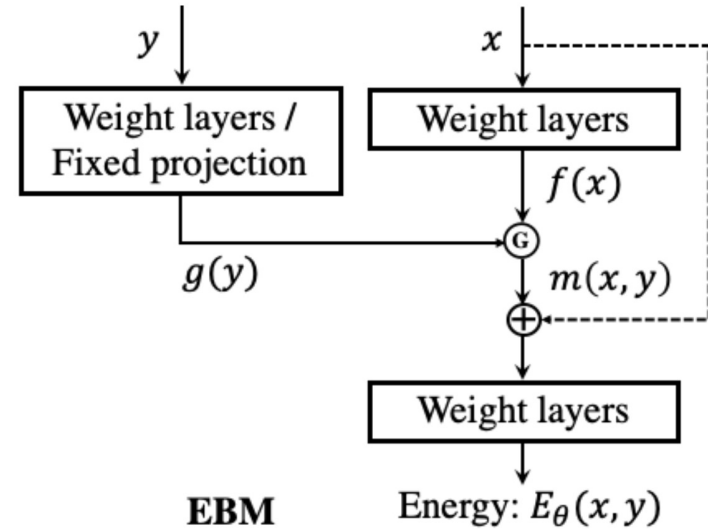

- Variational approach:

$$q_t(\boldsymbol{\theta}) = \operatorname{argmin}_{q \in \mathcal{Q}} \operatorname{KL} \left(q(\boldsymbol{\theta}) \parallel \frac{1}{Z_t} q_{t-1}(\boldsymbol{\theta}) p(\mathcal{D}_t \mid \boldsymbol{\theta}) \right).$$

- Loss: $\mathcal{L}(q_t(\boldsymbol{\theta})) = \mathbb{E}_{\boldsymbol{\theta} \sim q_t(\boldsymbol{\theta})} [-\log p(\mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta})] + \operatorname{KL}(q_t(\boldsymbol{\theta}) \parallel q_{t-1}(\boldsymbol{\theta})).$
 $q_t(\boldsymbol{\theta}) = \prod_{d=1}^D \mathcal{N}(\theta_{t,d}; \mu_{t,d}, \sigma_{t,d}^2).$
- Compare to EWC: Maintains uncertainty throughout training.



EBM for Continual Learning

- Softmax layer is known to be sensitive to distribution shift.



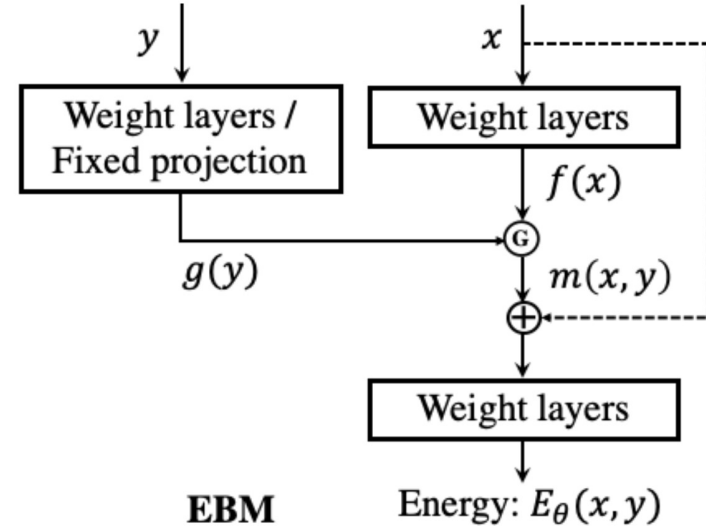

EBM for Continual Learning

- Softmax layer is known to be sensitive to distribution shift.
- A common approach is to use nearest mean classifier.



EBM for Continual Learning

- Softmax layer is known to be sensitive to distribution shift.
- A common approach is to use nearest mean classifier.
- Can be generalized to EBMs

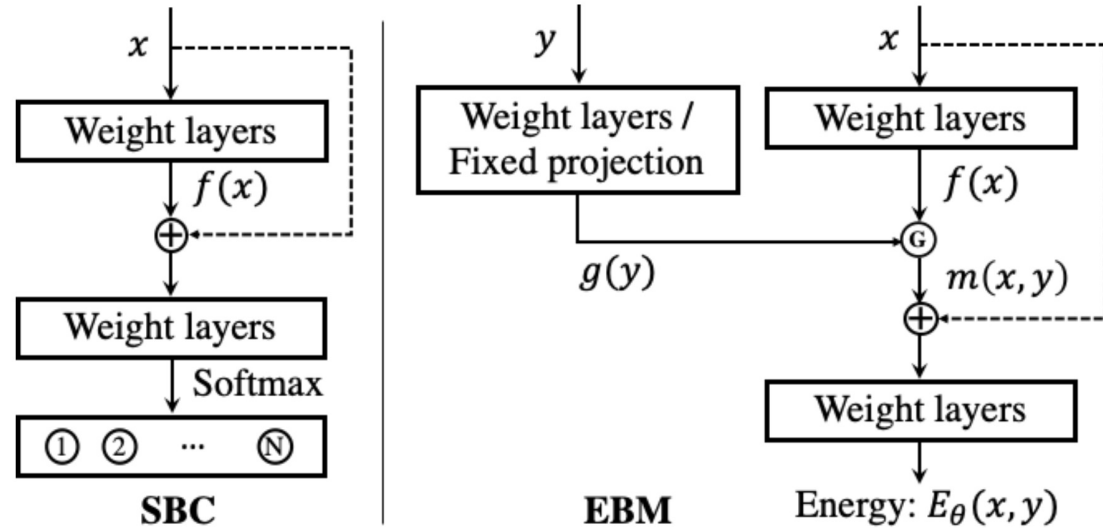


EBM for Continual Learning

- Softmax layer is known to be sensitive to distribution shift.
- A common approach is to use nearest mean classifier.
- Can be generalized to EBMs
- Energy between inputs and labels:


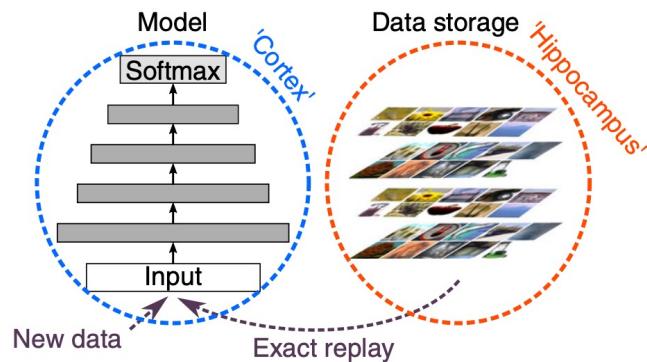
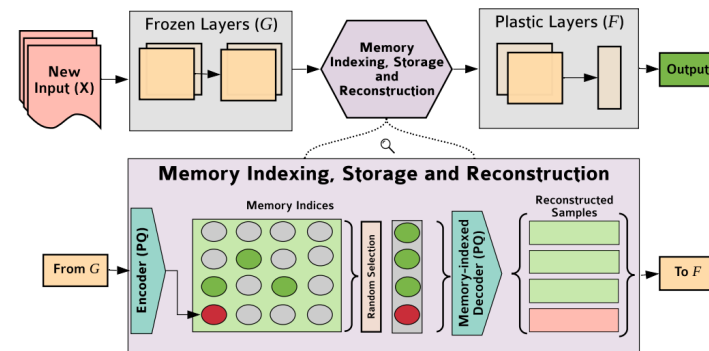
$$m(\mathbf{x}, y) = G(f(\mathbf{x}), g(y))$$

$$L_{CD}(\theta) = \mathbb{E}_{(\mathbf{x}, y, y^-) \sim p_D} [E_{\theta}(\mathbf{x}, y) - E_{\theta}(\mathbf{x}, y^-)].$$

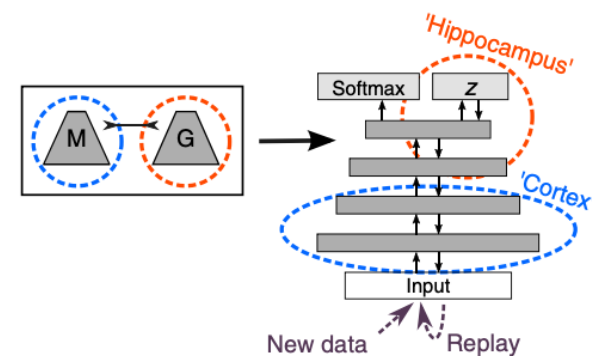


Replay

- Store raw data, representations, or train a generative model



Feedback Connections



Rebuffi et al. iCaRL: Incremental Classifier and Representation Learning. CVPR 2017.

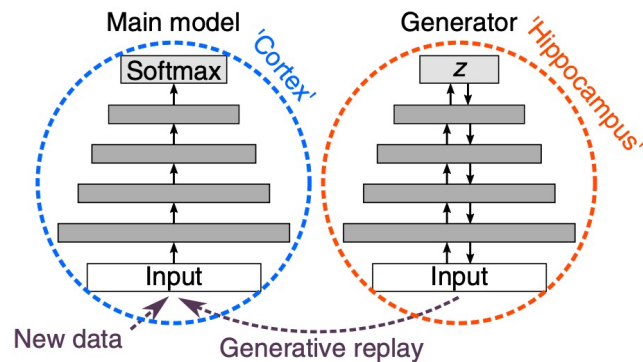


van de Ven et al. Brain-inspired replay for continual learning with artificial neural networks. Nature communications 2020.

Hayes et al. REMIND Your Neural Network to Prevent Catastrophic Forgetting. ECCV 2020.

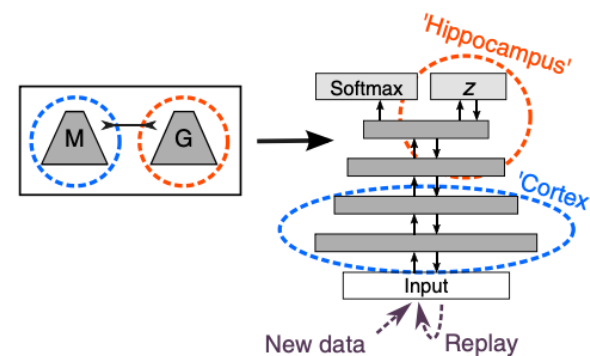
Replay

- Store raw data, representations, or train a generative model
- Coreset selection $\mu \leftarrow \frac{1}{n} \sum_{x \in X} \varphi(x)$

$$p_k \leftarrow \operatorname{argmin}_{x \in X} \left\| \mu - \frac{1}{k} [\varphi(x) + \sum_{j=1}^{k-1} \varphi(p_j)] \right\|.$$



Feedback Connections



Knowledge Distillation

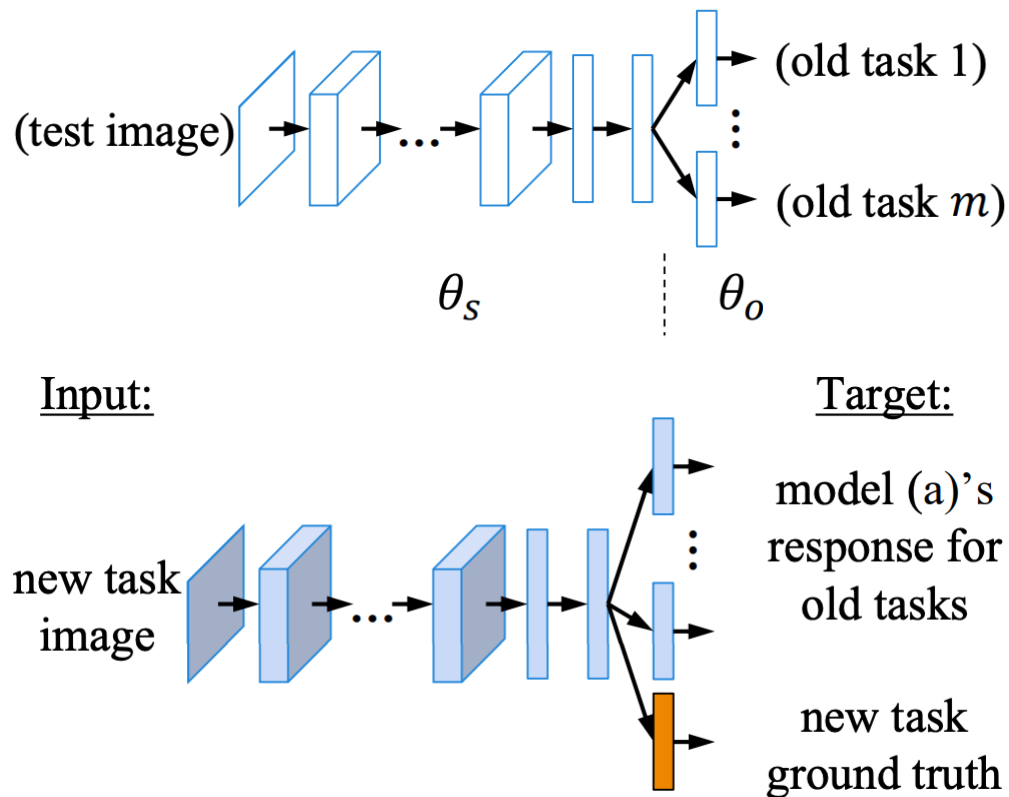
- Instead of saving the data points, we can also save the previous model checkpoint.

$$y_o = f(x_n; \theta_o).$$

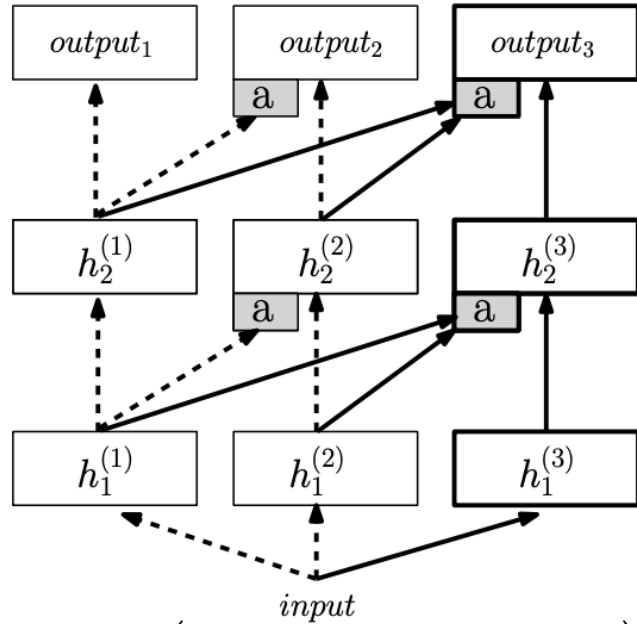
$$\hat{y}_o = f(x_n; \theta_n).$$

$$\mathcal{L}(y_o, \hat{y}_o)$$

- Use new data points and old weights to “distill”





Architecture Expansion



$$h_i^{(k)} = f \left(W_i^{(k)} h_{i-1}^{(k)} + \sum_{j < k} U_i^{(k:j)} h_{i-1}^{(j)} \right).$$

Architecture Expansion




$$h_i^{(k)} = f \left(W_i^{(k)} h_{i-1}^{(k)} + \sum_{j < k} U_i^{(k:j)} h_{i-1}^{(j)} \right).$$

Rusu et al. *Progressive Neural Networks*. NIPS 2016 Deep Learning Symposium.


PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. CVPR 2018.

Yoon et al. *Lifelong Learning with Dynamically Expandable Networks*. ICLR 2018.


Architecture Expansion



$$h_i^{(k)} = f \left(W_i^{(k)} h_{i-1}^{(k)} + \sum_{j < k} U_i^{(k:j)} h_{i-1}^{(j)} \right).$$




Selective Training



Selective Training

$$\mathcal{L}(\mathbf{W}_L^t, \mathbf{W}_{1:L-1}^{t-1}, \mathcal{D}_t) + \mu \|\mathbf{W}_L^t\|_1$$

Dynamic Expansion




Dynamic Expansion

$$\mathcal{L}(\mathbf{W}_l^N, \mathbf{W}_l^{t-1}, \mathcal{D}_t) + \mu \|\mathbf{W}_l^N\|_1 + \gamma \sum_g \|\mathbf{W}_{l,g}^N\|_2$$

Group sparsity

Network Split



Rusu et al. Progressive Neural Networks. NIPS 2016 Deep Learning Symposium.


PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. CVPR 2018.

Yoon et al. Lifelong Learning with Dynamically Expandable Networks. ICLR 2018.


Adapting Pretrained Models


- Pretrained models have general knowledge that can be adapted to a continual stream of tasks.

Rehearsal-based methods:
Fine-tuning



Our method:
Prompt selection + tuning


 : Trainable




Adapting Pretrained Models


- Pretrained models have general knowledge that can be adapted to a continual stream of tasks.
- Learn adaptation parameters for each task and store these as “task embeddings.”

Rehearsal-based methods:
Fine-tuning



Our method:
Prompt selection + tuning


 : Trainable




Adapting Pretrained Models


- Pretrained models have general knowledge that can be adapted to a continual stream of tasks.
- Learn adaptation parameters for each task and store these as “task embeddings.”
- Main model is frozen.

Rehearsal-based methods:
Fine-tuning



Our method:
Prompt selection + tuning


 : Trainable



Learning to Prompt

Prompt pool (slot memory)

$$\{(k_1, p_1), \dots, (k_M, p_M)\}$$

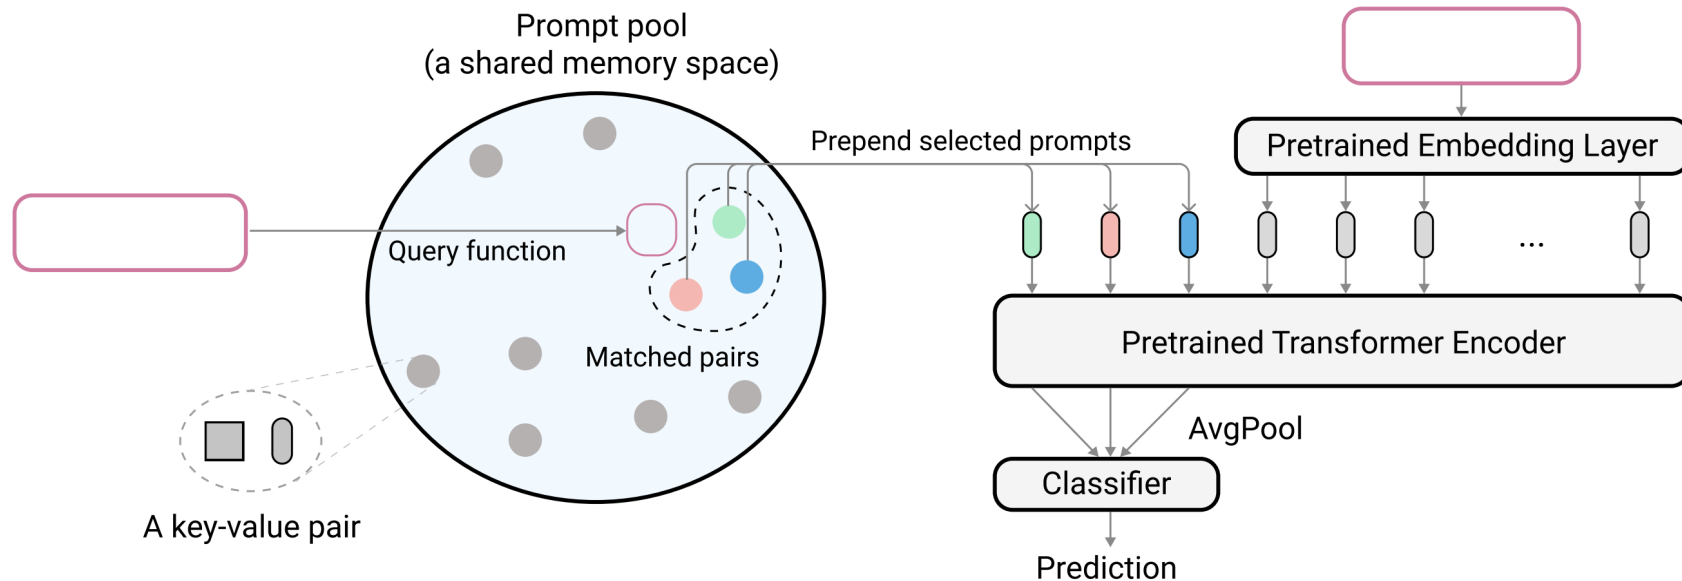


Learning to Prompt

Prompt pool (slot memory)

$\{(k_1, p_1), \dots, (k_M, p_M)\}$

$$K_s = \operatorname{argmin}_{\mathcal{S}} \sum_{i \in \mathcal{S}} \gamma(q(x), k_i)$$



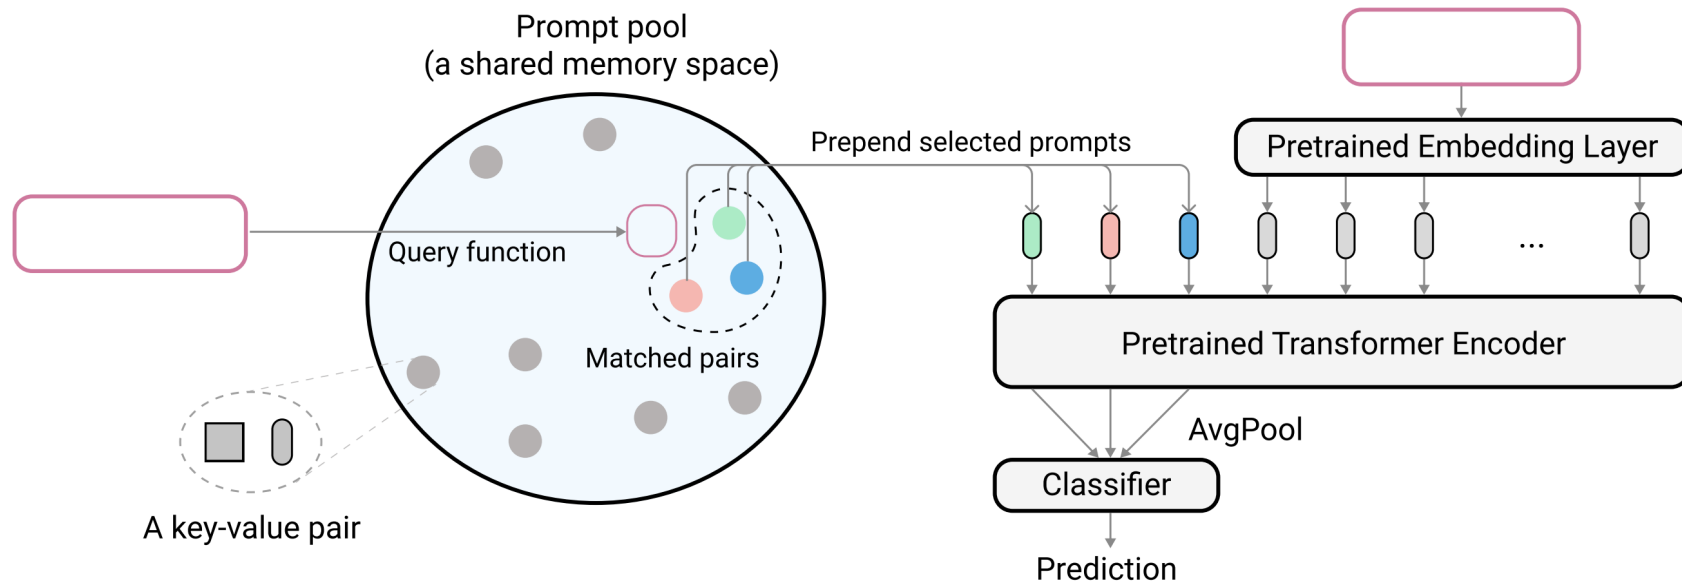
Learning to Prompt

Prompt pool (slot memory)

$\{(k_1, p_1), \dots, (k_M, p_M)\}$

$$K_s = \operatorname{argmin}_{\mathcal{S}} \sum_{i \in \mathcal{S}} \gamma(q(x), k_i)$$

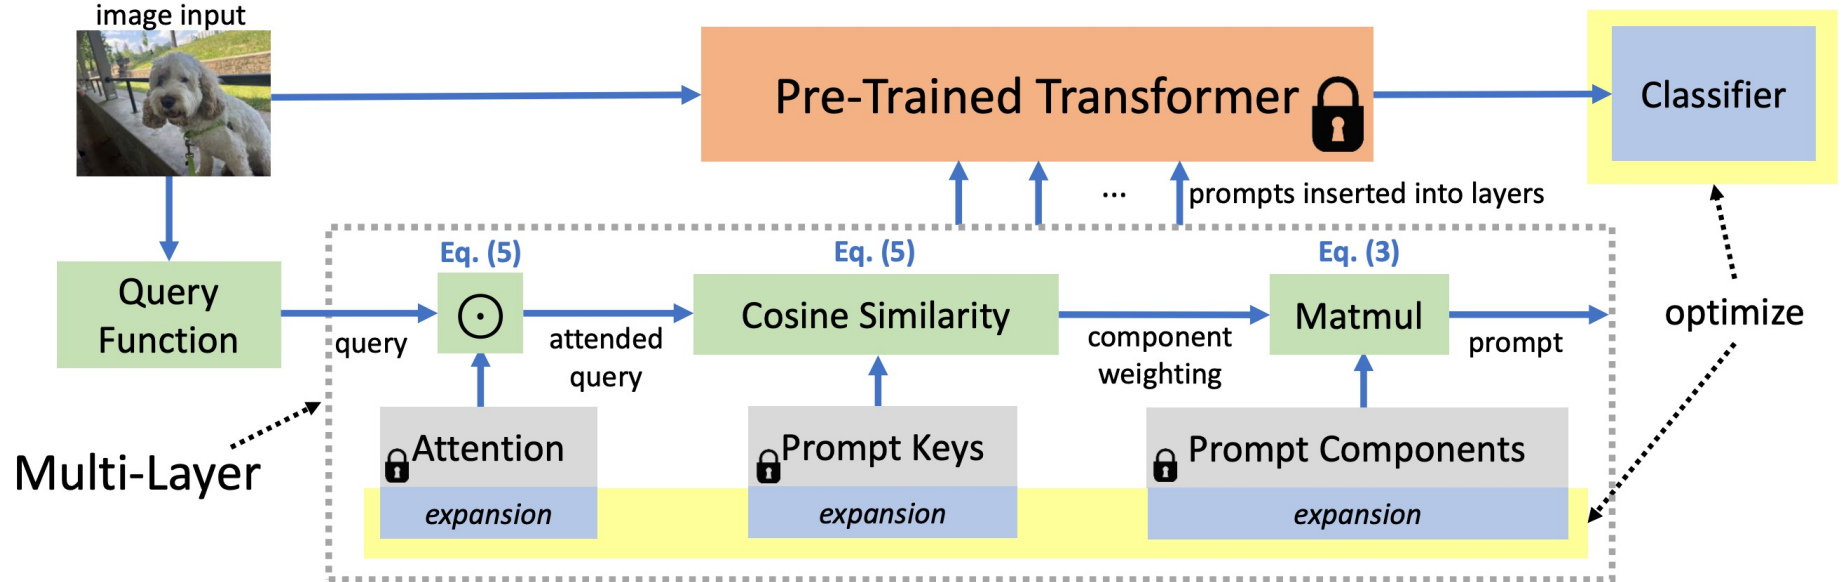
$$\min_{P, K, \phi} \mathcal{L}(g_{\phi}(x), y) + \lambda \sum_{i \in K_s} \gamma(q(x), k_i)$$



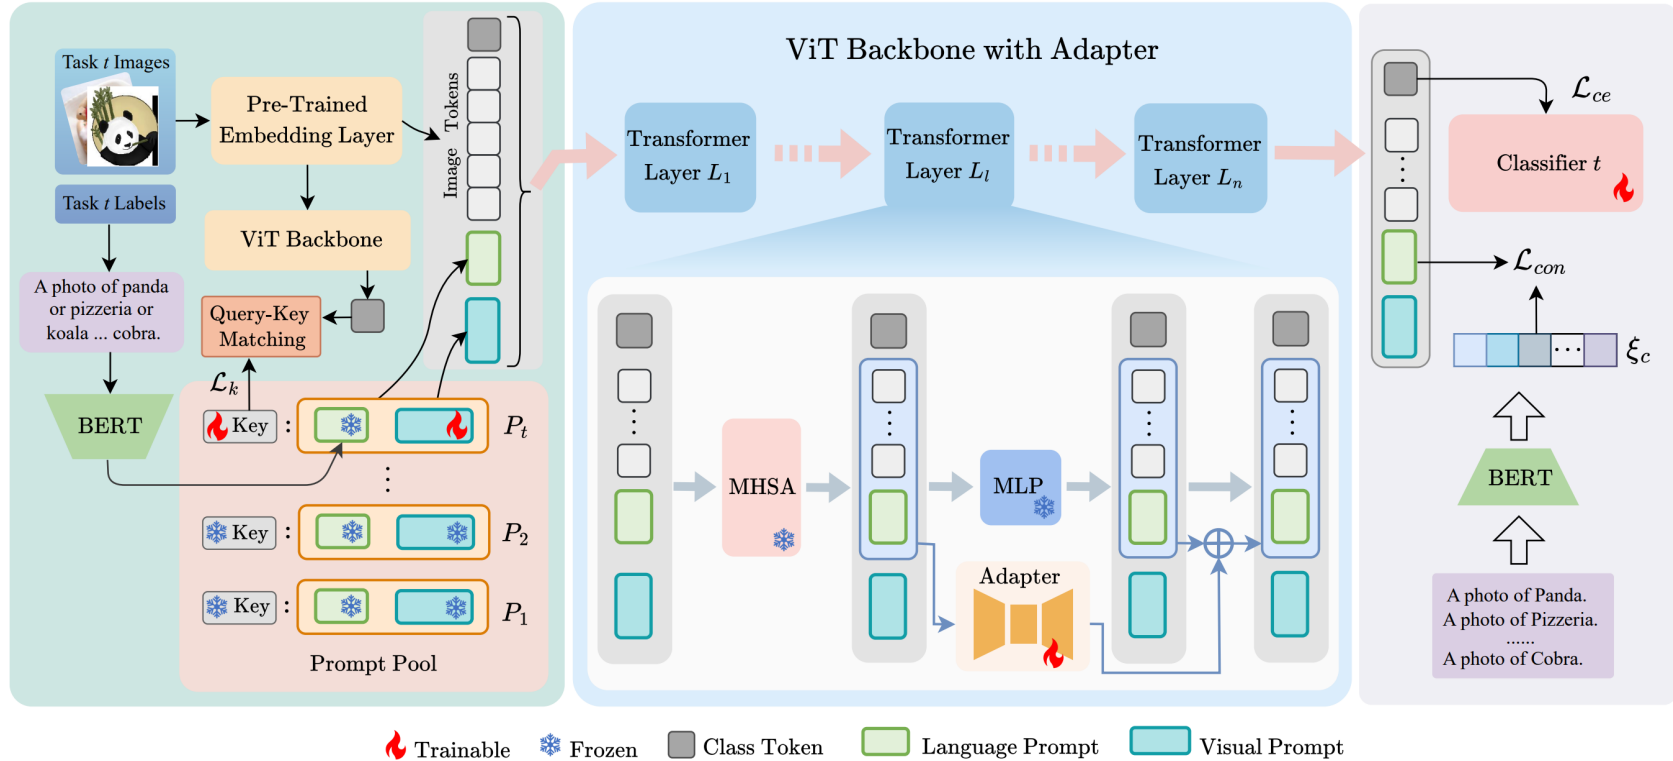
Learned Prompt Query

- The query function can be end-to-end learned.

$$\alpha = \{\gamma(q(x) \odot A_1, K_1), \dots, \gamma(q(x) \odot A_M, K_M)\}$$



Multimodal Semantic Prompts



Continual Learning and Memory

- Fragility of feedforward gradient descent of the entire networks

Continual Learning and Memory

- Fragility of feedforward gradient descent of the entire networks
- If we have representations ready, continual learning is just memorizing a sequence of new tasks.

Continual Learning and Memory

- Fragility of feedforward gradient descent of the entire networks
- If we have representations ready, continual learning is just memorizing a sequence of new tasks.
- In prompting approaches:
 - prompt pool = memory
 - pretrained network = representations

Continual Learning and Memory


- Fragility of feedforward gradient descent of the entire networks
- If we have representations ready, continual learning is just memorizing a sequence of new tasks.
- In prompting approaches:
 - prompt pool = memory
 - pretrained network = representations
- But what if representations also need to be built sequentially?

Continual Learning and Memory

- Fragility of feedforward gradient descent of the entire networks
- If we have representations ready, continual learning is just memorizing a sequence of new tasks.
- In prompting approaches:
 - prompt pool = memory
 - pretrained network = representations
- But what if representations also need to be built sequentially?
- It's also plausible that representations are just “deeper memory.”

Associative Memory

- Memory aims to store content for easy retrieval



Associative Memory

- Memory aims to store content for easy retrieval
 - Associative memories (Hopfield Networks) can be viewed as energy-based models



Associative Memory

- Memory aims to store content for easy retrieval
 - Associative memories (Hopfield Networks) can be viewed as energy-based models

$$E = -\frac{1}{2} \sum_{i,j=1}^N s_i W_{ij} s_j, \quad W_{ij} = \sum_{k=1}^K \xi_i^k \xi_j^k.$$

"Superposition" of k slots
Hebbian learning



Associative Memory

- Memory aims to store content for easy retrieval
 - Associative memories (Hopfield Networks) can be viewed as energy-based models
$$E = -\frac{1}{2} \sum_{i,j=1}^N s_i W_{ij} s_j, \quad W_{ij} = \sum_{k=1}^K \xi_i^k \xi_j^k.$$

"Superposition" of k slots
Hebbian learning
 - When presented with a new pattern the network should respond with a stored memory which most closely resembles the input.



Associative Memory


- Memory aims to store content for easy retrieval

- Associative memories (Hopfield Networks) can be viewed as energy-based models

$$E = -\frac{1}{2} \sum_{i,j=1}^N s_i W_{ij} s_j, \quad W_{ij} = \sum_{k=1}^K \xi_i^k \xi_j^k.$$

"Superposition" of k slots
Hebbian learning



- When presented with a new pattern the network should respond with a stored memory which most closely resembles the input.
- Retrieval: $s_i = \text{sign}(\sum_j W_{ij} s_j)$ Storage: $C \approx \frac{d}{2 \log(d)} = 0.14d$.



Memory Duality

- Duality with a feedforward network.

$$E = - \sum_k F(\sum_i \xi_i^k s_i)$$





Memory Duality

- Duality with a feedforward network.

$$E = - \sum_k F(\sum_i \xi_i^k s_i)$$

- Non-linearity allows us to store more patterns.





Memory Duality


- Duality with a feedforward network.

$$E = - \sum_k F(\sum_i \xi_i^k s_i)$$

- Non-linearity allows us to store more patterns.
- Deep Boltzmann machines



Deep Boltzmann Machine




Memory Duality




- Duality with a feedforward network.

$$E = - \sum_k F(\sum_i \xi_i^k s_i)$$

- Non-linearity allows us to store more patterns.
- Deep Boltzmann machines
- Hierarchical associative memory



Deep Boltzmann Machine



Relation to Transformers

- General form:


$$E = - \sum_k F\left(\sum_i \xi_i^k s_i\right).$$

Relation to Transformers

- General form:

$$E = - \sum_k F\left(\sum_i \xi_i^k s_i\right).$$

- When $F(z) = z^2$ it gives the classic HN.



$$\nabla_{s_i} E = - \sum_j W_{ij} s_j$$
$$s_i \leftarrow \text{sign}\left(\sum_j W_{ij} s_j\right)$$


Relation to Transformers

- General form:

$$E = - \sum_k F\left(\sum_i \xi_i^k s_i\right).$$

- When $F(z) = z^2$ it gives the classic HN.
- Transformer-like attention operation:

$$\mathbf{Z} \leftarrow \text{softmax}(\beta \mathbf{X} \mathbf{W}_q \mathbf{W}_k^\top \mathbf{Y}^\top) \mathbf{Y} \mathbf{W}_v.$$



$$\nabla_{s_i} E = - \sum_j W_{ij} s_j$$
$$s_i \leftarrow \text{sign}\left(\sum_j W_{ij} s_j\right)$$

Relation to Transformers


- General form:

$$E = - \sum_k F\left(\sum_i \xi_i^k s_i\right).$$

- When $F(z) = z^2$ it gives the classic HN.
- Transformer-like attention operation:

$$\mathbf{Z} \leftarrow \text{softmax}(\beta \mathbf{X} \mathbf{W}_q \mathbf{W}_k^\top \mathbf{Y}^\top) \mathbf{Y}_i \mathbf{W}_v.$$

$$\mathbf{S} \leftarrow \text{softmax}(\beta \mathbf{S} \mathbf{\Xi}^\top) \mathbf{\Xi}.$$



$$\nabla_{s_i} E = - \sum_j W_{ij} s_j$$
$$s_i \leftarrow \text{sign}\left(\sum_j W_{ij} s_j\right)$$

Relation to Transformers

- General form:


$$E = - \sum_k F\left(\sum_i \xi_i^k s_i\right).$$

- When $F(z) = z^2$ it gives the classic HN.
- Transformer-like attention operation:

$$\mathbf{Z} \leftarrow \text{softmax}(\beta \mathbf{X} \mathbf{W}_q \mathbf{W}_k^\top \mathbf{Y}^\top) \mathbf{Y}_i \mathbf{W}_v.$$

$$\mathbf{S} \leftarrow \text{softmax}(\beta \mathbf{S} \mathbf{\Xi}^\top) \mathbf{\Xi}.$$

$$E = -\text{logsumexp}(\beta, \mathbf{\Xi}^\top \mathbf{s}) + \frac{1}{2} \mathbf{s}^\top \mathbf{s} + \beta^{-1} \log N + \frac{1}{2} M^2.$$




$$\nabla_{s_i} E = - \sum_j W_{ij} s_j$$


$$s_i \leftarrow \text{sign}\left(\sum_j W_{ij} s_j\right)$$

Continual Self-Supervised Learning

- Learning from a stream of unlabeled inputs.




SUPERVISED CONTINUAL LEARNING (SCL)




UNSUPERVISED CONTINUAL LEARNING (UCL)

Continual Self-Supervised Learning

- Learning from a stream of unlabeled inputs.
- Bring SSL to the dynamic world.




SUPERVISED CONTINUAL LEARNING (SCL)




UNSUPERVISED CONTINUAL LEARNING (UCL)

Continual Self-Supervised Learning

- Learning from a stream of unlabeled inputs.
- Bring SSL to the dynamic world.
- SSL can still suffer from distributional shifts.



SUPERVISED CONTINUAL LEARNING (SCL)




UNSUPERVISED CONTINUAL LEARNING (UCL)

Continual Self-Supervised Learning



- Integrating learning objectives of both past and present.

$$\begin{array}{cccc} \text{Current} & \text{Past (Distillation)} & \text{Past (Replay)} & \text{Cross-Consolidation} \\ \mathcal{L}(X; g \circ f_t) + \mathcal{L}(X, X; h \circ f_{t-1}, h \circ f_t) + \mathcal{L}(Y; h \circ f_t) + \mathcal{L}(X, Y; h \circ f_t) \end{array}$$



Outlooks





- Understand continual learning at scale



[Mayo et al. 2023]

Outlooks



- Understand continual learning at scale
- Unified learning architecture, objective and replay, role of sleep





[Mayo et al. 2023]

Outlooks


- Understand continual learning at scale
- Unified learning architecture, objective and replay, role of sleep
- Continual learning with real world structure



[van de Ven et al. 2020]



Typical setting in continual, few-shot, transfer, and meta-learning



Typical setting of human learning [Mayo et al. 2023]



Tasks are interspersed and recur
No opportunity to master one before confronted with another

Summary: Continual Learning

- Regularization, Distillation, Architecture Expansion/Isolation

Summary: Continual Learning

- Regularization, Distillation, Architecture Expansion/Isolation
- Frozen representation: prompt learning

Summary: Continual Learning

- Regularization, Distillation, Architecture Expansion/Isolation
- Frozen representation: prompt learning
- Integration of memory and representations

Summary: Continual Learning

- Regularization, Distillation, Architecture Expansion/Isolation
- Frozen representation: prompt learning
- Integration of memory and representations
- Combination with self-supervised learning

Summary: Continual Learning

- Regularization, Distillation, Architecture Expansion/Isolation
- Frozen representation: prompt learning
- Integration of memory and representations
- Combination with self-supervised learning
- Exploration of multimodal continual learning from embodied environments