Implicit Generation and Generalization with Energy Based Models

Yilun Du and Igor Mordatch

Energy-Based Model

- Distribution defined by energy function

$$p_\theta(\mathbf{x}) = \frac{\exp(-E_\theta(\mathbf{x}))}{Z(\theta)} \quad Z(\theta) = \int \exp(-E_\theta(\mathbf{x}))d\mathbf{x}$$

see [LeCun et al, 2006] for review

Energy-Based Model

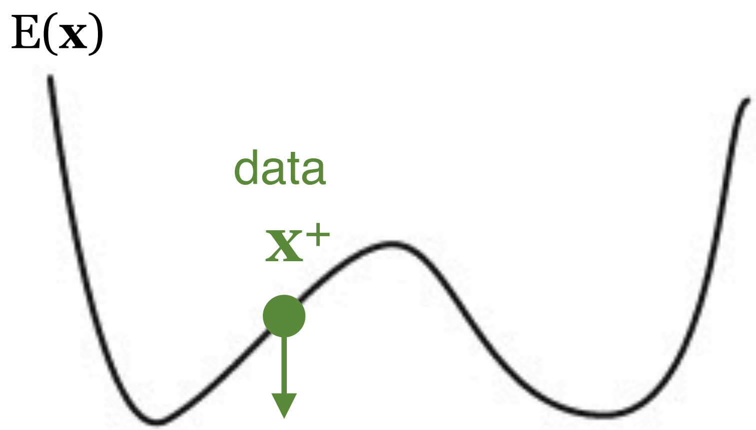

- Distribution defined by energy function

$$p_\theta(\mathbf{x}) = \frac{\exp(-E_\theta(\mathbf{x}))}{Z(\theta)}$$

- Train to maximize data likelihood $\mathcal{L}_{\text{ML}}(\theta) = \mathbb{E}_{\mathbf{x} \sim p_D} [-\log p_\theta(\mathbf{x})]$

Energy-Based Model

- Distribution defined by energy function

$$p_\theta(\mathbf{x}) = \frac{\exp(-E_\theta(\mathbf{x}))}{Z(\theta)}$$

- Train to maximize data likelihood $\mathcal{L}_{\text{ML}}(\theta) = \mathbb{E}_{\mathbf{x}\sim p_D}\left[-\log p_\theta(\mathbf{x})\right]$

  - gradient: $\mathbb{E}_{\mathbf{x}^+\sim p_D}\left[\nabla_\theta E_\theta(\mathbf{x}^+)\right] - \mathbb{E}_{\mathbf{x}^-\sim p_\theta}\left[\nabla_\theta E_\theta(\mathbf{x}^-)\right]$

  See [Turner, 2006] for derivation

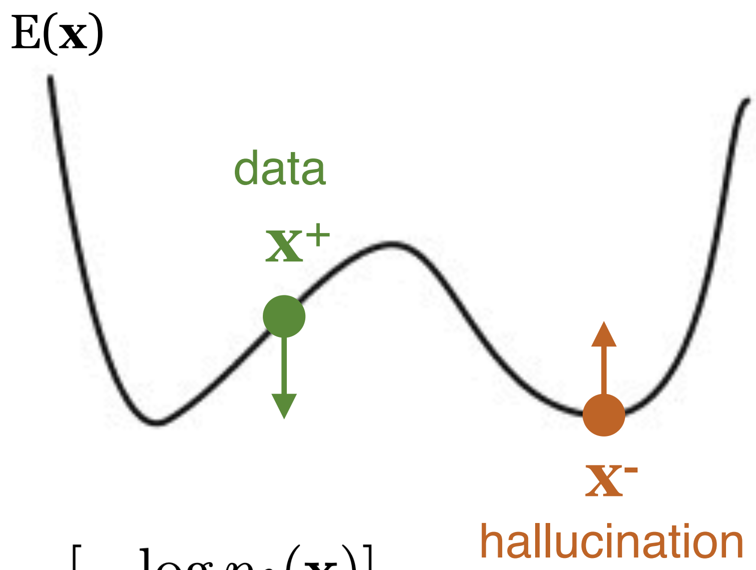

Energy-Based Model

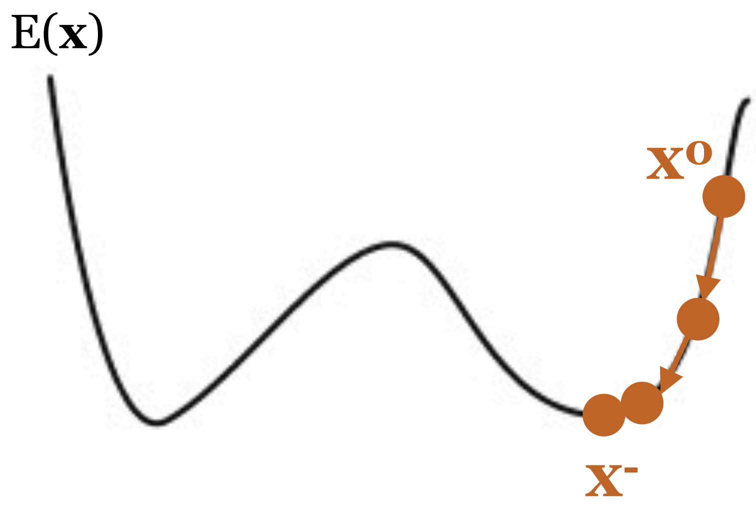

- Distribution defined by energy function

$$p_\theta(\mathbf{x}) = \frac{\exp(-E_\theta(\mathbf{x}))}{Z(\theta)}$$

- Train to maximize data likelihood
  - gradient: $\mathbb{E}_{\mathbf{x}^+ \sim p_D}[\nabla_\theta E_\theta(\mathbf{x}^+)] - \mathbb{E}_{\mathbf{x}^- \sim p_\theta}[\nabla_\theta E_\theta(\mathbf{x}^-)]$
- Generate model samples implicitly via stochastic optimization

$$\tilde{\mathbf{x}}^k = \tilde{\mathbf{x}}^{k-1} - \frac{\lambda}{2}\nabla_{\mathbf{x}} E_\theta(\tilde{\mathbf{x}}^{k-1}) + \omega^k, \ \omega^k \sim \mathcal{N}(0, \lambda)$$

Langevin Dynamics
[Welling and Teh, 2011]

Why Energy-Based Generative Models?

1. Implicit Generation
   - Flexibility
   - One Object to Learn
   - Compositionalitly
   - Generic Initialization and Computation Time
2. Intriguing Properties
   - Robustness
   - Online Learning

Why Do EBMs Work Now?

More compute and modern deep learning practices

Faster Sampling

- Continuous gradient based sampling using Langevin Dynamics
- Replay buffer of past samples (similar to persistent CD)

Stability improvements

- Constrain Lipschitz constant of energy function (spectral norm)
- Smoother activations (swish)
- And others ...

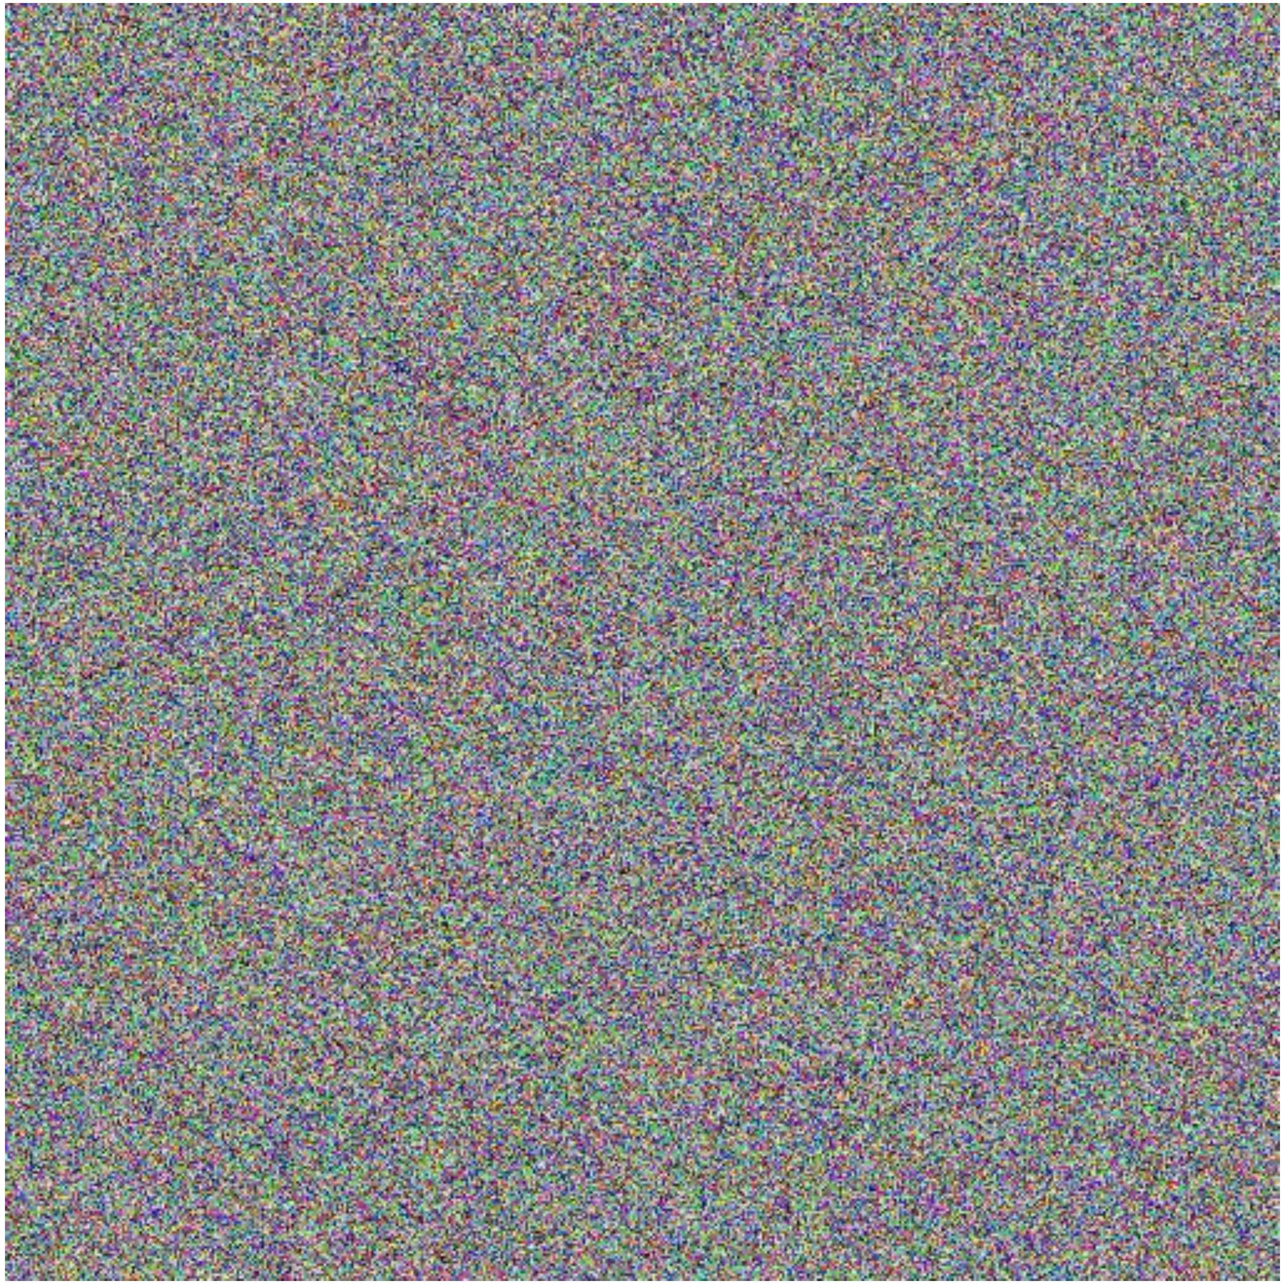

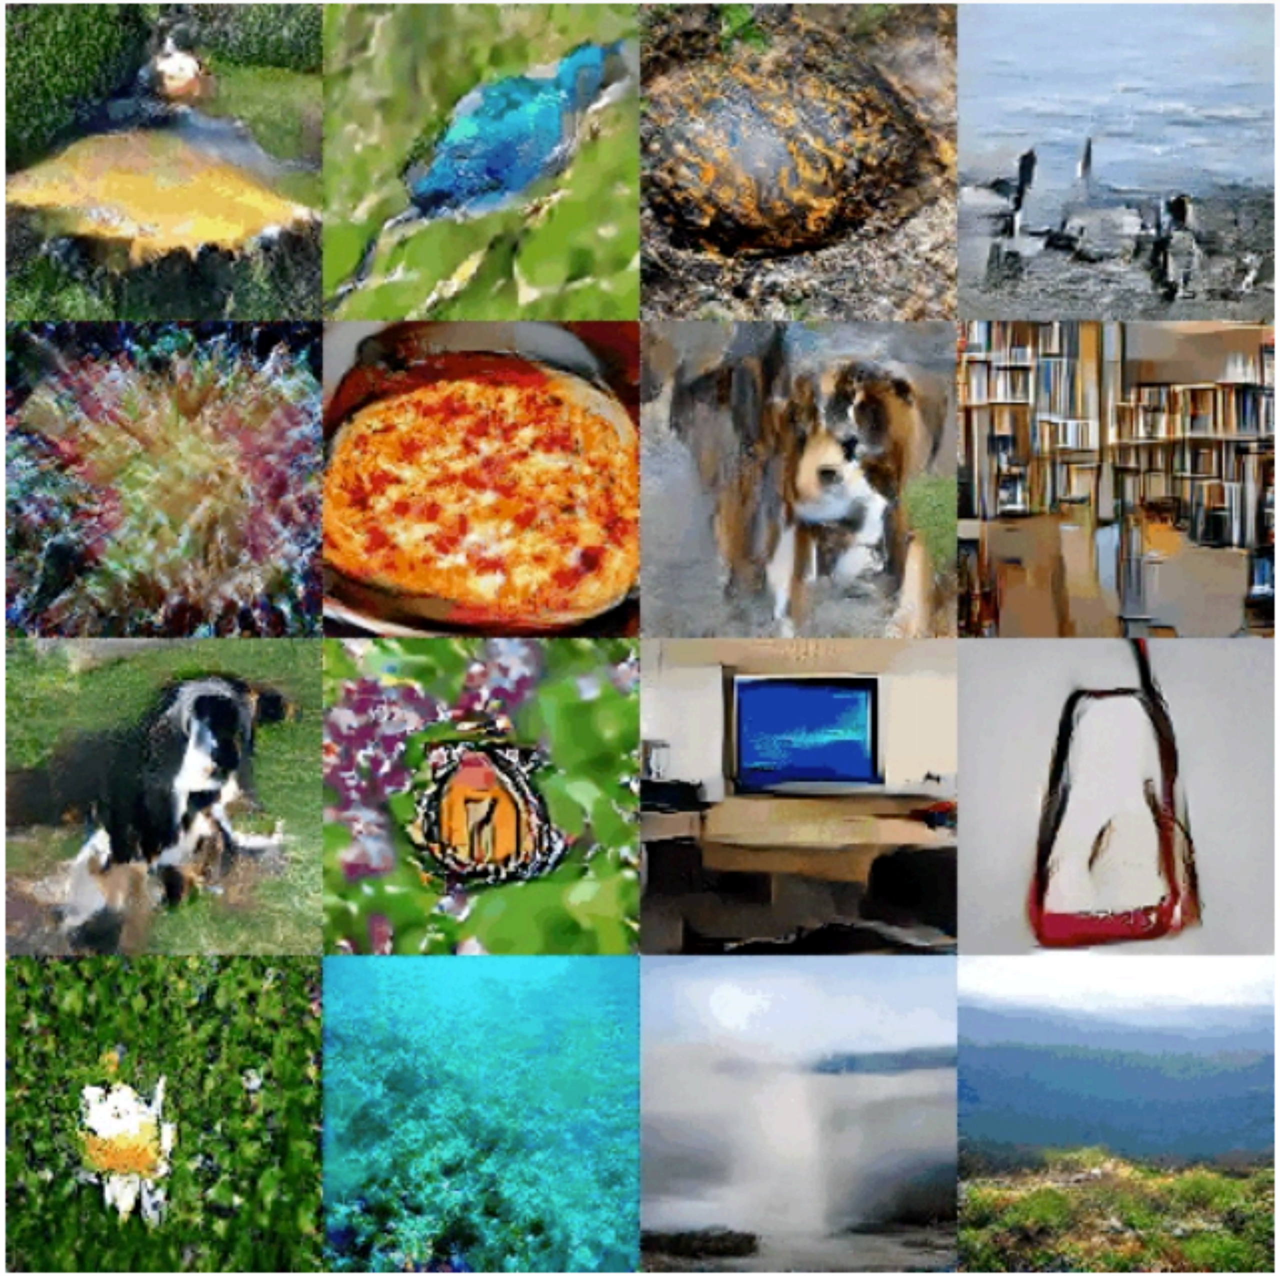

Comparison to Other Generative Models

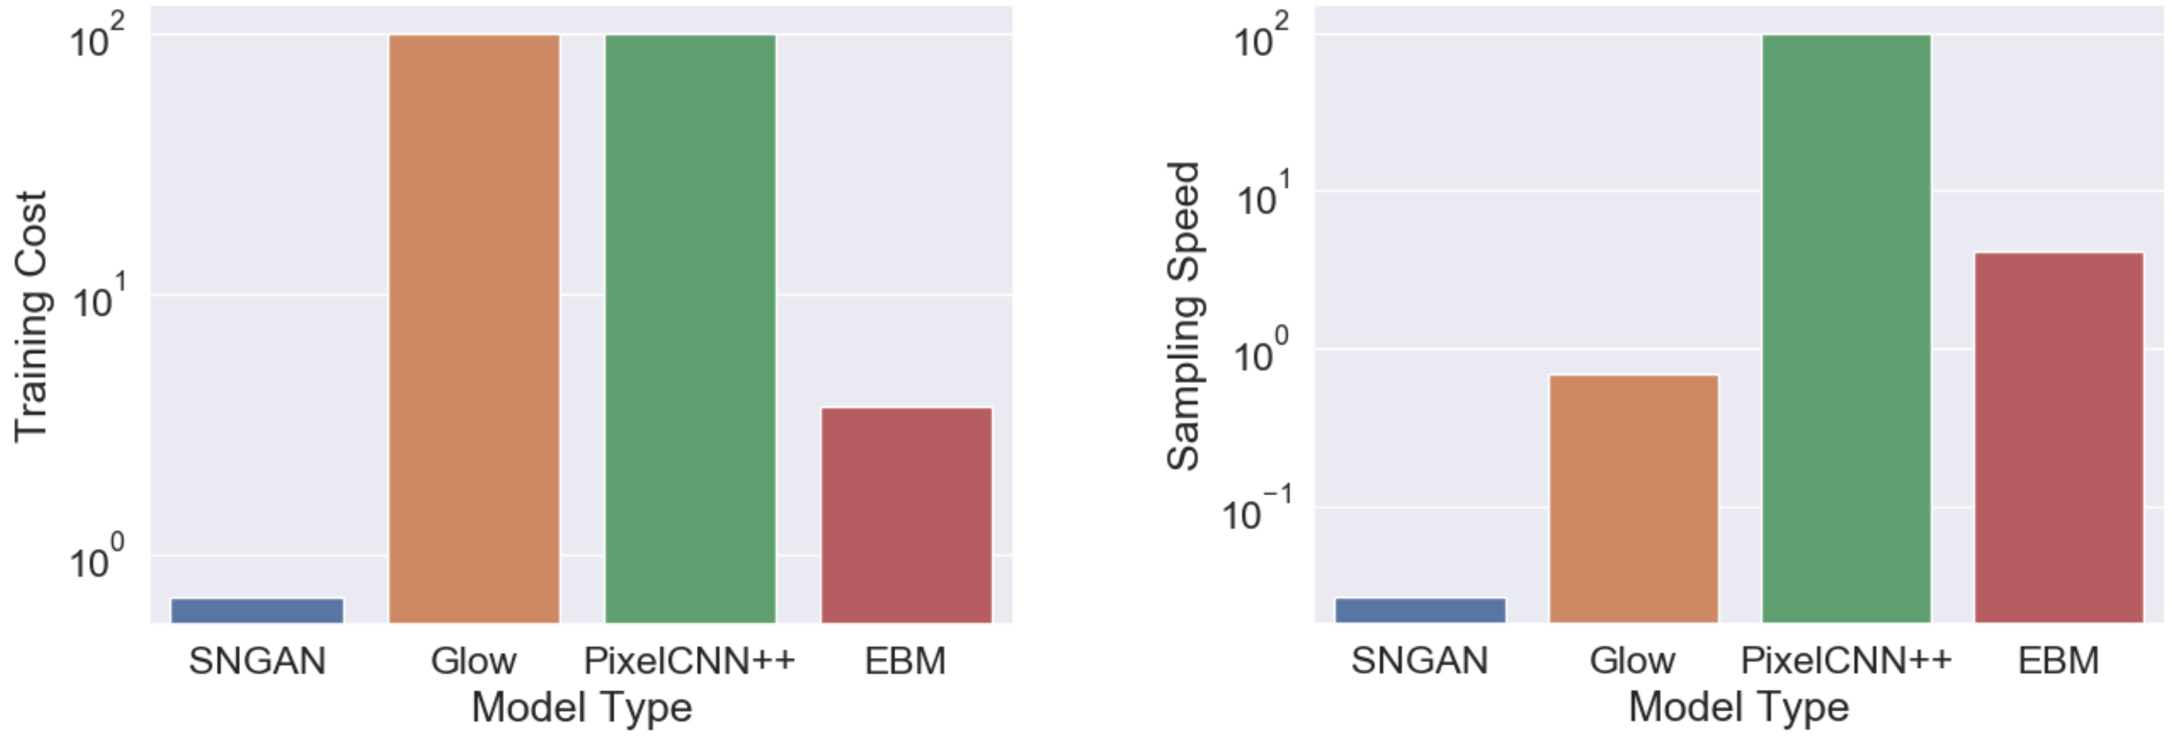

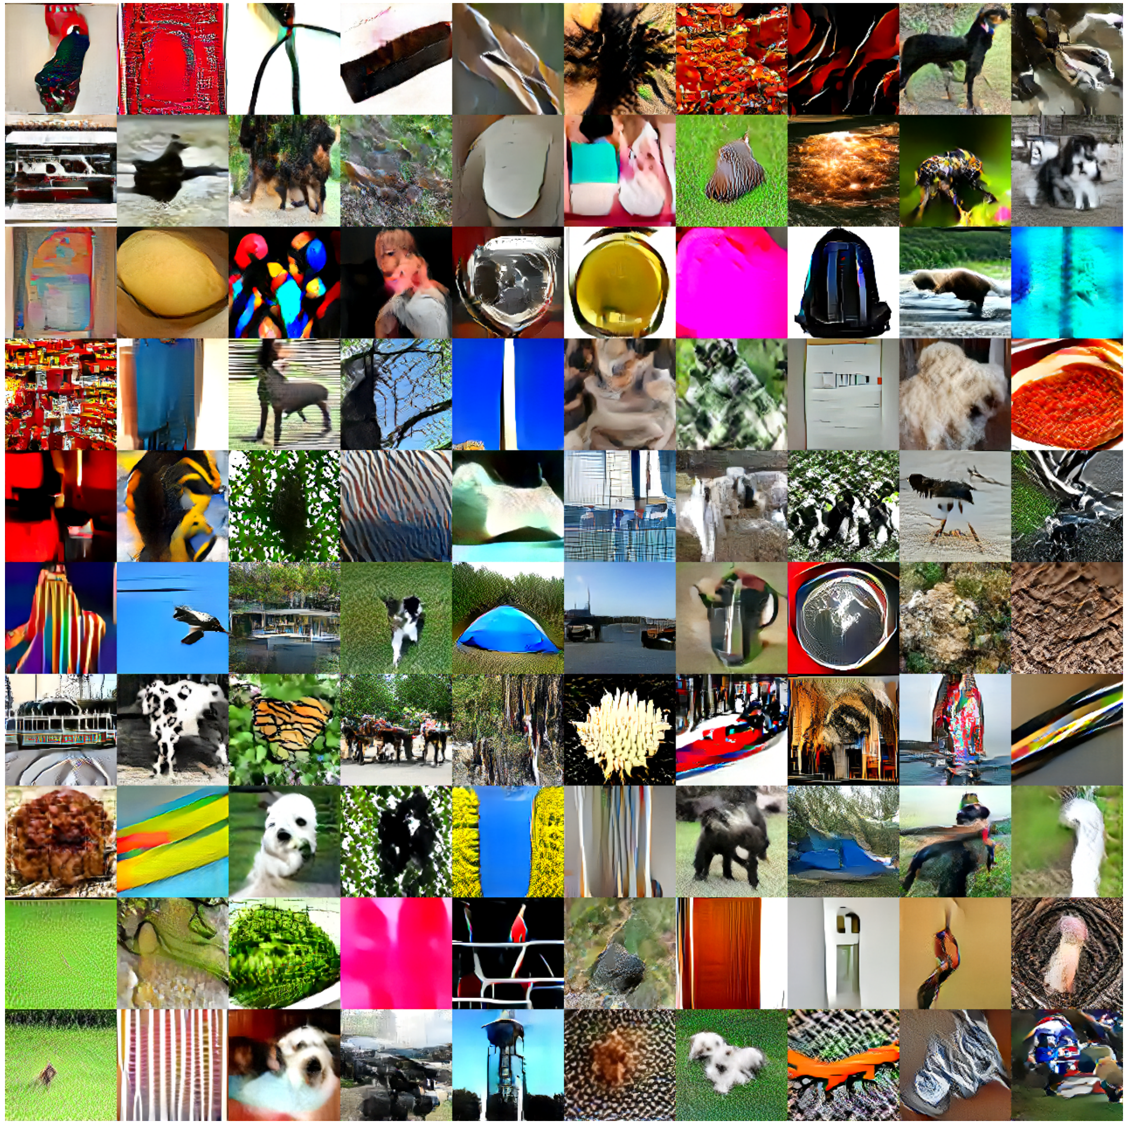

ImageNet 128x128

Model	Inception	FID
CIFAR-10 Unconditional		
PixelCNN (Van Oord et al., 2016)	4.60	65.93
PixelIQN (Ostrovski et al., 2018)	5.29	49.46
EBM (single)	6.02	40.58
DCGAN (Radford et al., 2016)	6.40	37.11
WGAN + GP (Gulrajani et al., 2017)	6.50	36.4
EBM (10 historical ensemble)	6.78	38.2
SNGAN (Miyato et al., 2018)	**8.22**	21.7
CIFAR-10 Conditional		
Improved GAN	8.09	-
EBM (single)	8.30	37.9
Spectral Normalization GAN	**8.59**	25.5
ImageNet 32x32 Conditional		
PixelCNN	8.33	33.27
PixelIQN	10.18	22.99
EBM (single)	**18.22**	**14.31**
ImageNet 128x128 Conditional		
ACGAN (Odena et al., 2017)	28.5	-
EBM* (single)	28.6	43.7
SNGAN	**36.8**	**27.62**

Cross Class Mapping

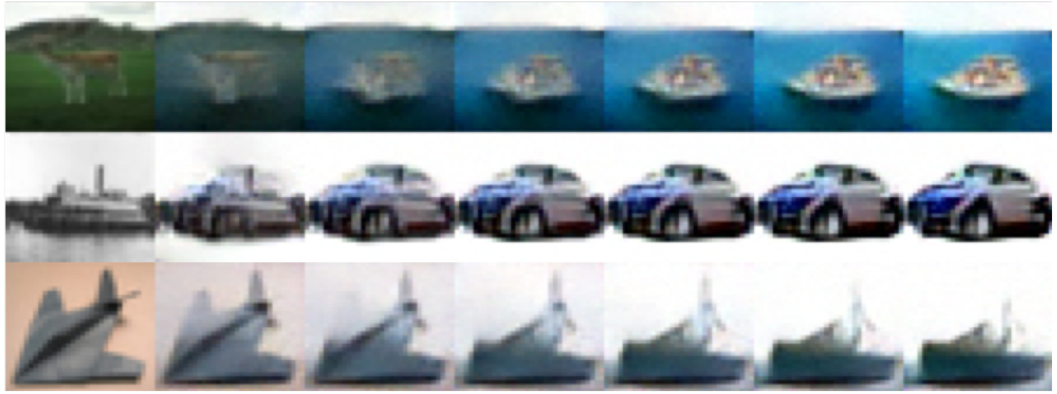

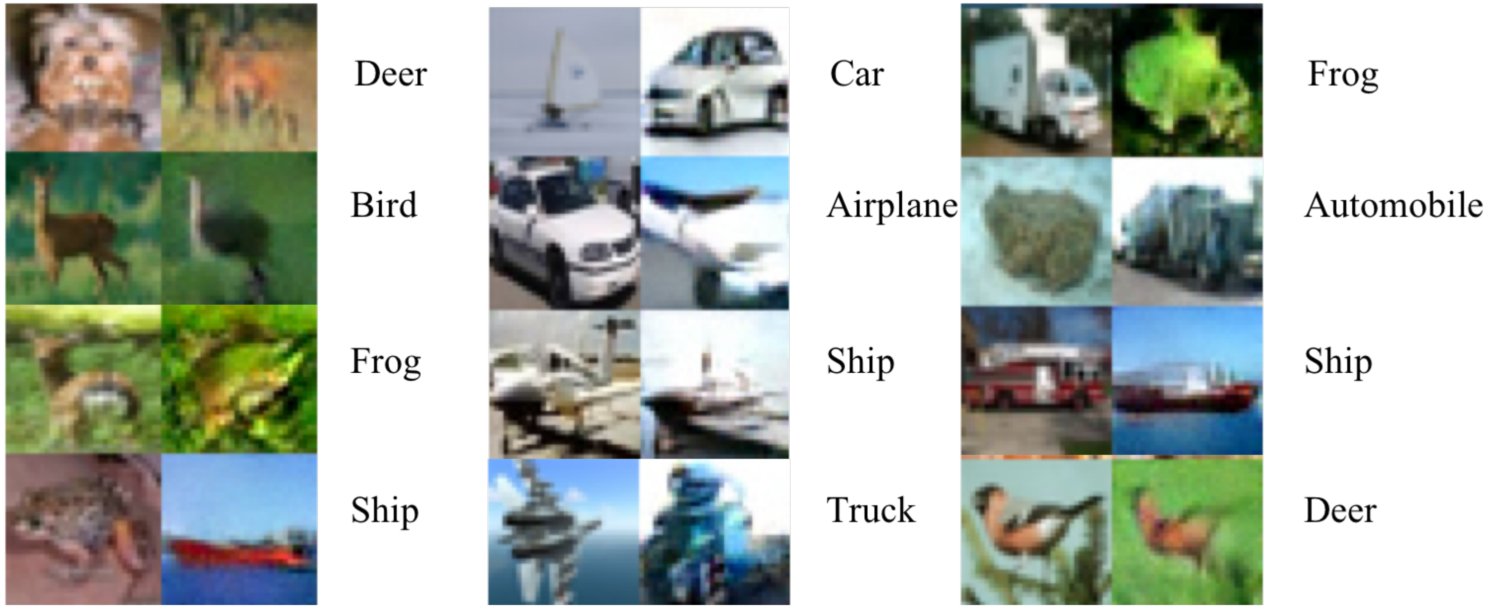

Cross Class Mapping

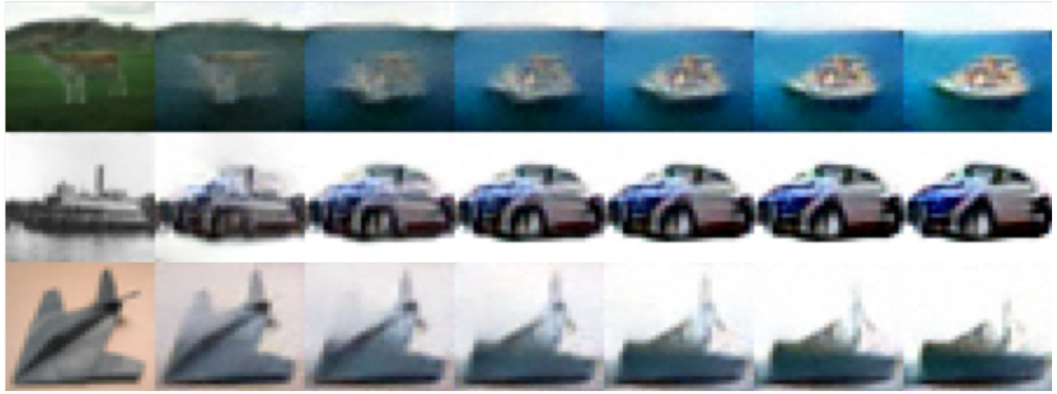

Surprising Benefits of Energy-Based Models

- Robustness
- Continual Learning
- Compositionality
- Trajectory Modeling

Surprising Benefits of Energy-Based Models

- Robustness
- Continual Learning
- Compositionality
- Trajectory Modeling

Out-of-Distribution Relative Likelihoods

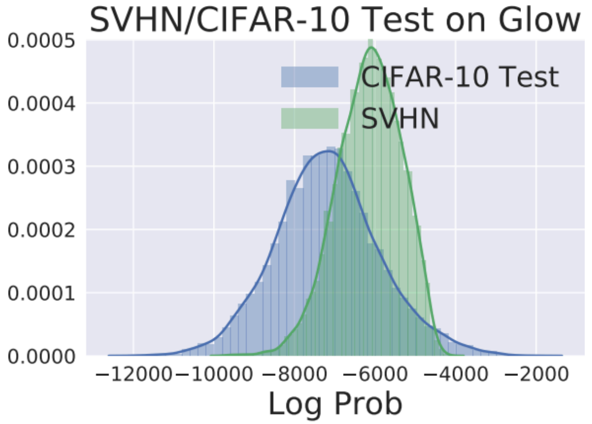

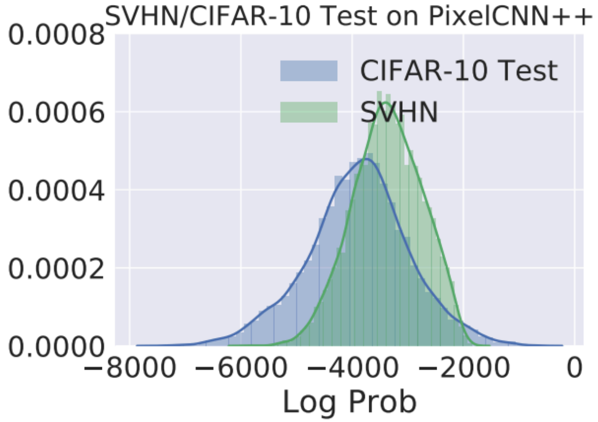

Also observed by [Hendrycks et al 2018]
and [Nalisnick et al 2019]

Out-of-Distribution Relative Likelihoods

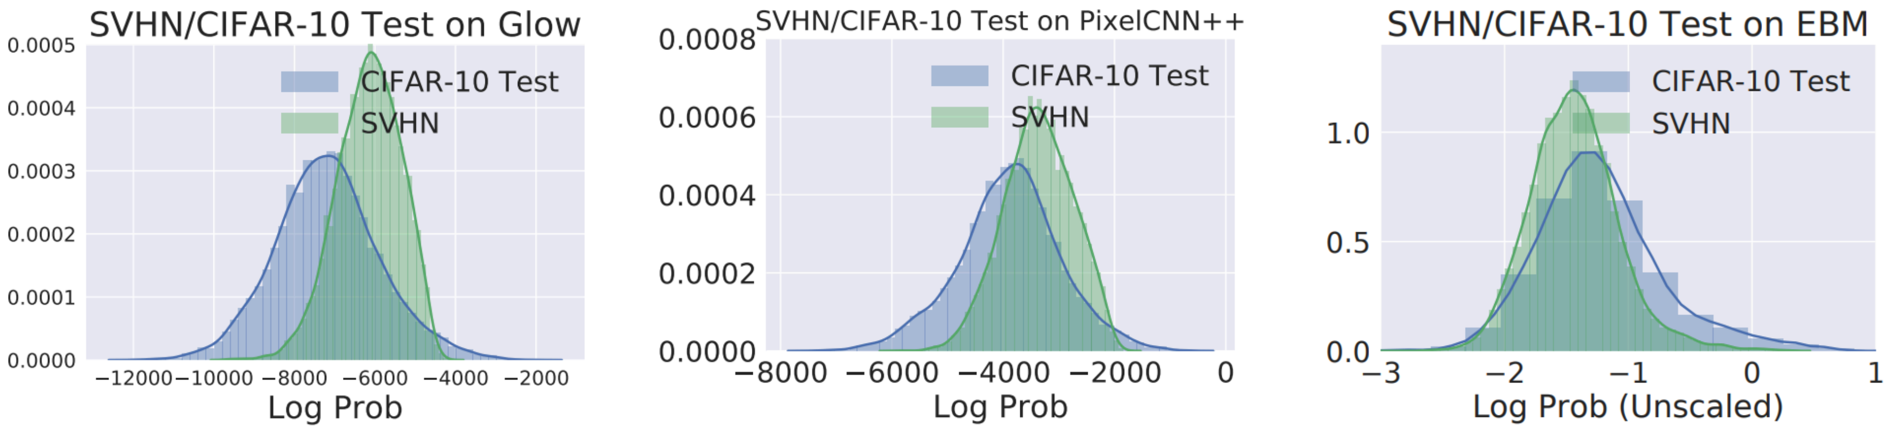

Also observed by [Hendrycks et al 2018] and [Nalisnick et al 2019]

Out-of-Distribution Relative Likelihoods

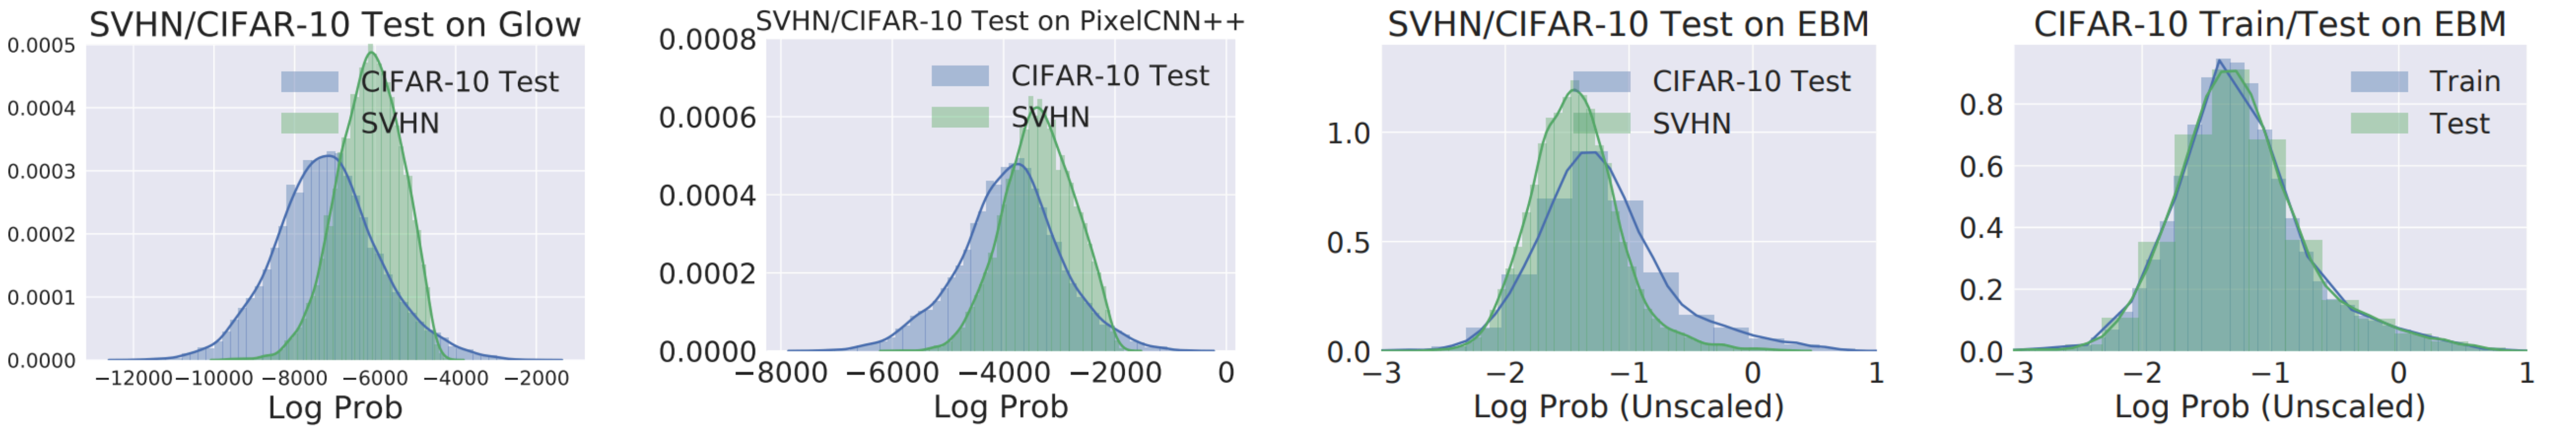

Also observed by [Hendrycks et al 2018] and [Nalisnick et al 2019]

Out-of-Distribution Generalization

- Following [Hendrycks and Gimpel, 2016]

Textures SVHN Constant Uniform Uniform CIFAR10 Mix CIFAR10

Model	SVHN	Textures	Monochrome Uniform	Uniform	CIFAR10 Interpolation	Average
PixelCNN++	0.32	0.33	0.0	1.0	**0.71**	0.47
Glow	0.24	0.27	0.0	1.0	0.59	0.42
EBM (ours)	**0.63**	**0.48**	**0.30**	**1.0**	0.70	**0.62**

Robust Classification

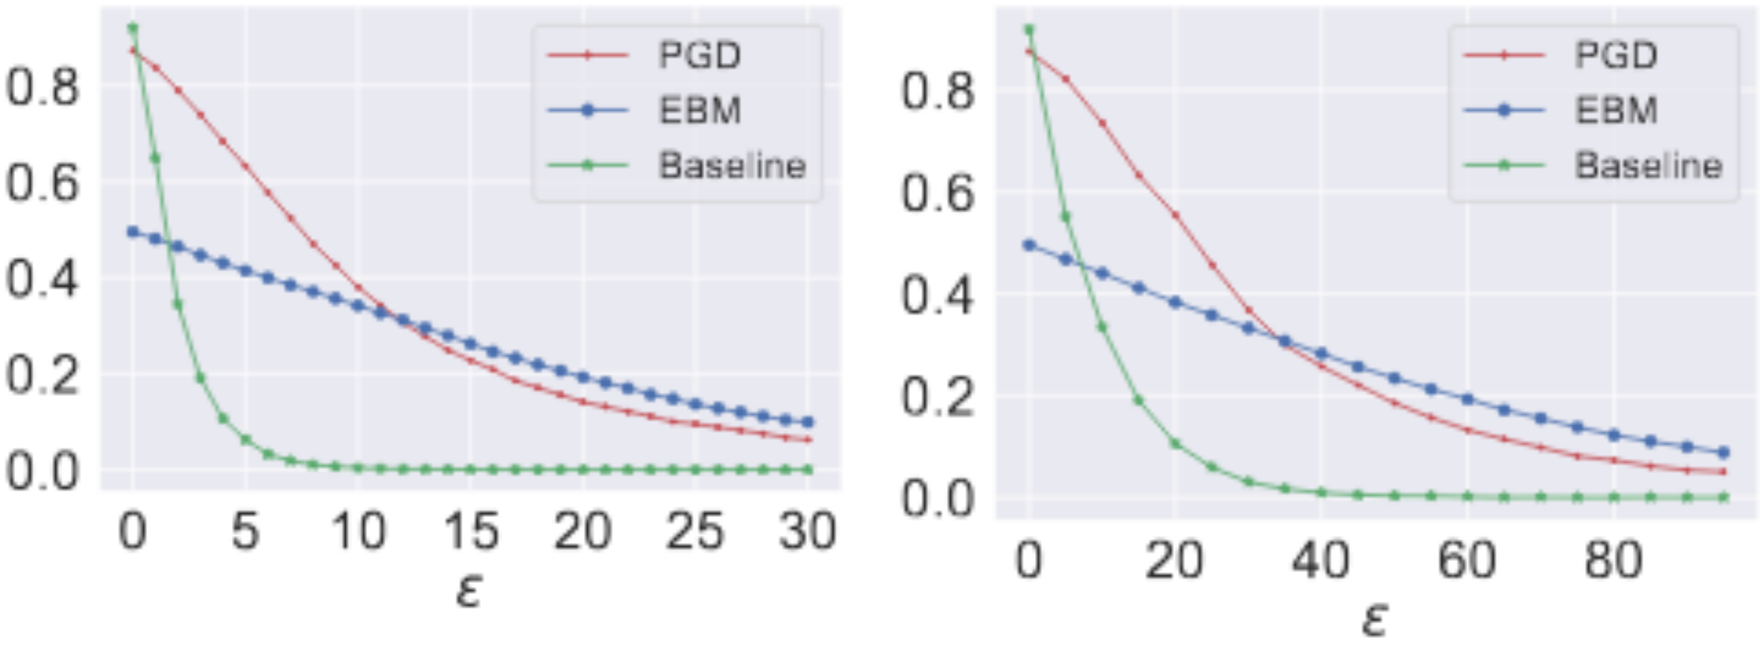

(a) L_∞ robustness (b) L_2 Robustness

Robust Classification

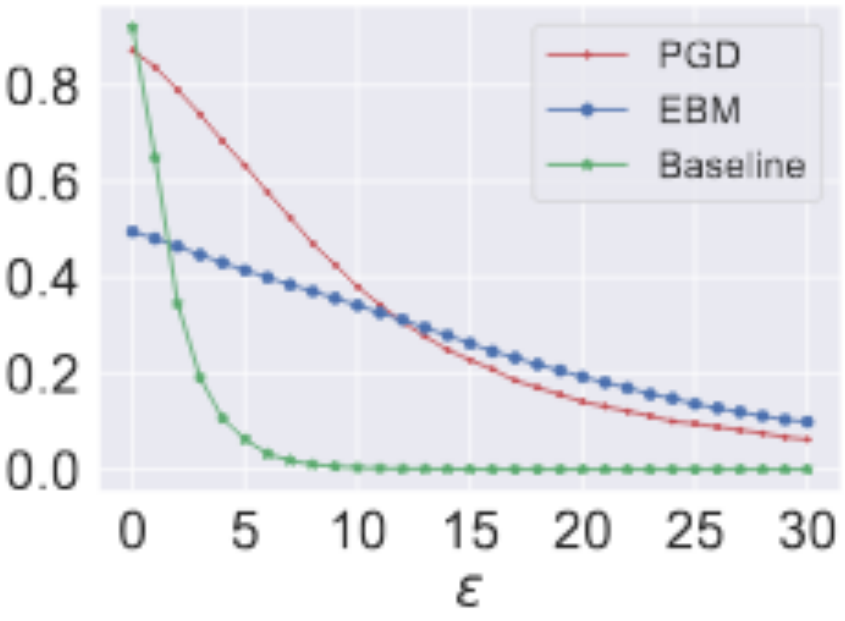

(a) L_∞ robustness

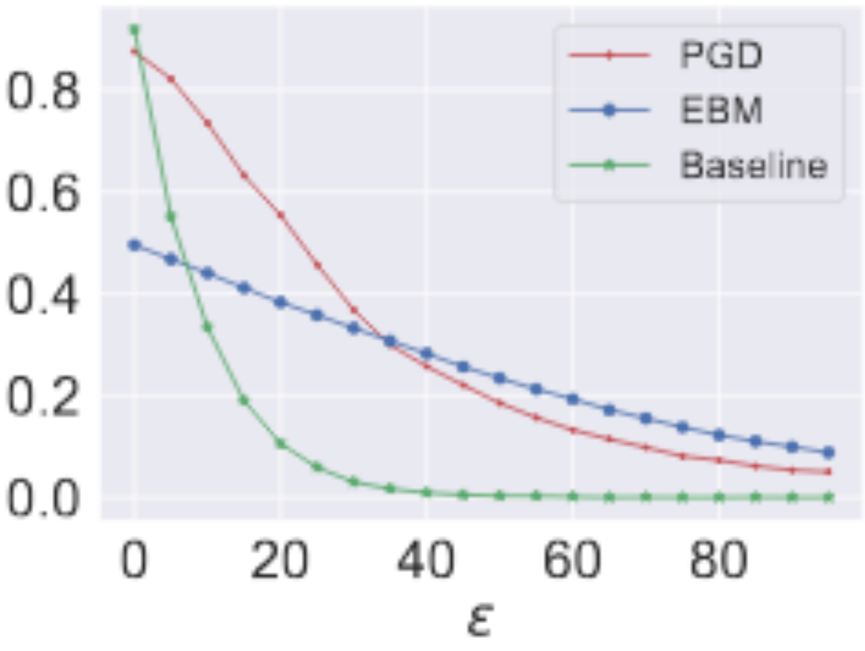

(b) L_2 Robustness

(recent follow-up submission at ICLR 2020 improves baseline EBM performance)

Surprising Benefits of Energy-Based Models

- Robustness
- Continual Learning
- Compositionality
- Trajectory Modeling

Continual Learning: Split MNIST

	Method	Memory	Incremental task learning	Incremental domain learning	Incremental class learning
Baselines	Adam		93.46 ± 2.01	55.16 ± 1.38	19.71 ± 0.08
	SGD		97.98 ± 0.09	63.20 ± 0.35	19.46 ± 0.04
	Adagrad		98.06 ± 0.53	58.08 ± 1.06	19.82 ± 0.09
	L2		98.18 ± 0.96	66.00 ± 3.73	22.52 ± 1.08
	Naive rehearsal	✓	99.40 ± 0.08	95.16 ± 0.49	90.78 ± 0.85
	Naive rehearsal-C	✓	**99.57** ± 0.07	**97.11** ± 0.34	**95.59** ± 0.49
Continual learning methods	EWC		97.70 ± 0.81	58.85 ± 2.59	19.80 ± 0.05
	Online EWC		98.04 ± 1.10	57.33 ± 1.44	19.77 ± 0.04
	SI		98.56 ± 0.49	64.76 ± 3.09	19.67 ± 0.09
	MAS		99.22 ± 0.21	68.57 ± 6.85	19.52 ± 0.29
	LwF		99.60 ± 0.03	71.02 ± 1.26	24.17 ± 0.33
	GEM	✓	98.42 ± 0.10	96.16 ± 0.35	92.20 ± 0.12
	DGR	✓	99.47 ± 0.03	95.74 ± 0.23	91.24 ± 0.33
	RtF	✓	**99.66** ± 0.03	**97.31** ± 0.11	**92.56** ± 0.21
Offline (upper bound)			99.52 ± 0.16	98.59 ± 0.15	97.53 ± 0.30

Evaluation by [Hsu at al, 2019]

Continual Learning: Split MNIST

	Method	Memory	Incremental task learning	Incremental domain learning	Incremental class learning
Baselines	Adam		93.46 ± 2.01	55.16 ± 1.38	19.71 ± 0.08
	SGD		97.98 ± 0.09	63.20 ± 0.35	19.46 ± 0.04
	Adagrad		98.06 ± 0.53	58.08 ± 1.06	19.82 ± 0.09
	L2		98.18 ± 0.96	66.00 ± 3.73	22.52 ± 1.08
	Naive rehearsal	✓	99.40 ± 0.08	95.16 ± 0.49	90.78 ± 0.85
	Naive rehearsal-C	✓	**99.57** ± 0.07	**97.11** ± 0.34	**95.59** ± 0.49
Continual learning methods	EWC		97.70 ± 0.81	58.85 ± 2.59	19.80 ± 0.05
	Online EWC		98.04 ± 1.10	57.33 ± 1.44	19.77 ± 0.04
	SI		98.56 ± 0.49	64.76 ± 3.09	19.67 ± 0.09
	MAS		99.22 ± 0.21	68.57 ± 6.85	19.52 ± 0.29
	LwF		99.60 ± 0.03	71.02 ± 1.26	24.17 ± 0.33
	GEM	✓	98.42 ± 0.10	96.16 ± 0.35	92.20 ± 0.12
	DGR	✓	99.47 ± 0.03	95.74 ± 0.23	91.24 ± 0.33
	RtF	✓	**99.66** ± 0.03	**97.31** ± 0.11	**92.56** ± 0.21
Offline (upper bound)			99.52 ± 0.16	98.59 ± 0.15	97.53 ± 0.30

Evaluation by [Hsu at al, 2019]

Continual Learning: Split MNIST

	Method	Memory	Incremental task learning	Incremental domain learning	Incremental class learning
Baselines	Adam		93.46 ± 2.01	55.16 ± 1.38	19.71 ± 0.08
	SGD		97.98 ± 0.09	63.20 ± 0.35	19.46 ± 0.04
	Adagrad		98.06 ± 0.53	58.08 ± 1.06	19.82 ± 0.09
	L2		98.18 ± 0.96	66.00 ± 3.73	22.52 ± 1.08
	Naive rehearsal	✓	99.40 ± 0.08	95.16 ± 0.49	90.78 ± 0.85
	Naive rehearsal-C	✓	**99.57** ± 0.07	**97.11** ± 0.34	**95.59** ± 0.49
Continual learning methods	EWC		97.70 ± 0.81	58.85 ± 2.59	19.80 ± 0.05
	Online EWC		98.04 ± 1.10	57.33 ± 1.44	19.77 ± 0.04
	SI		98.56 ± 0.49	64.76 ± 3.09	19.67 ± 0.09
	MAS		99.22 ± 0.21	68.57 ± 6.85	19.52 ± 0.29
	LwF		99.60 ± 0.03	71.02 ± 1.26	24.17 ± 0.33
	GEM	✓	98.42 ± 0.10	96.16 ± 0.35	92.20 ± 0.12
	DGR	✓	99.47 ± 0.03	95.74 ± 0.23	91.24 ± 0.33
	RtF	✓	**99.66** ± 0.03	**97.31** ± 0.11	**92.56** ± 0.21
Offline (upper bound)			99.52 ± 0.16	98.59 ± 0.15	97.53 ± 0.30

EBM: 64.99 ± 4.27

(10 seeds)

Evaluation by [Hsu at al, 2019]

Continual Learning: Split MNIST

	Method	Memory	Incremental task learning	Incremental domain learning	Incremental class learning
Baselines	Adam		93.46 ± 2.01	55.16 ± 1.38	19.71 ± 0.08
	SGD		97.98 ± 0.09	63.20 ± 0.35	19.46 ± 0.04
	Adagrad		98.06 ± 0.53	58.08 ± 1.06	19.82 ± 0.09
	L2		98.18 ± 0.96	66.00 ± 3.73	22.52 ± 1.08
	Naive rehearsal	✓	99.40 ± 0.08	95.16 ± 0.49	90.78 ± 0.85
	Naive rehearsal-C	✓	**99.57** ± 0.07	**97.11** ± 0.34	**95.59** ± 0.49
Continual learning methods	EWC		97.70 ± 0.81	58.85 ± 2.59	19.80 ± 0.05
	Online EWC		98.04 ± 1.10	57.33 ± 1.44	19.77 ± 0.04
	SI		98.56 ± 0.49	64.76 ± 3.09	19.67 ± 0.09
	MAS		99.22 ± 0.21	68.57 ± 6.85	19.52 ± 0.29
	LwF		99.60 ± 0.03	71.02 ± 1.26	24.17 ± 0.33
	GEM	✓	98.42 ± 0.10	96.16 ± 0.35	92.20 ± 0.12
	DGR	✓	99.47 ± 0.03	95.74 ± 0.23	91.24 ± 0.33
	RtF	✓	**99.66** ± 0.03	**97.31** ± 0.11	**92.56** ± 0.21
Offline (upper bound)			99.52 ± 0.16	98.59 ± 0.15	97.53 ± 0.30

EBM: 64.99 ± 4.27

Would any generative model work instead? Doesn’t look like it:

VAE: 40.04 ± 1.31

Evaluation by [Hsu at al, 2019]

Surprising Benefits of Energy-Based Models

- Robustness
- Continual Learning
- Compositionality
- Trajectory Modeling

Compositionality via Sum of EBMs

[Hinton, 1999]

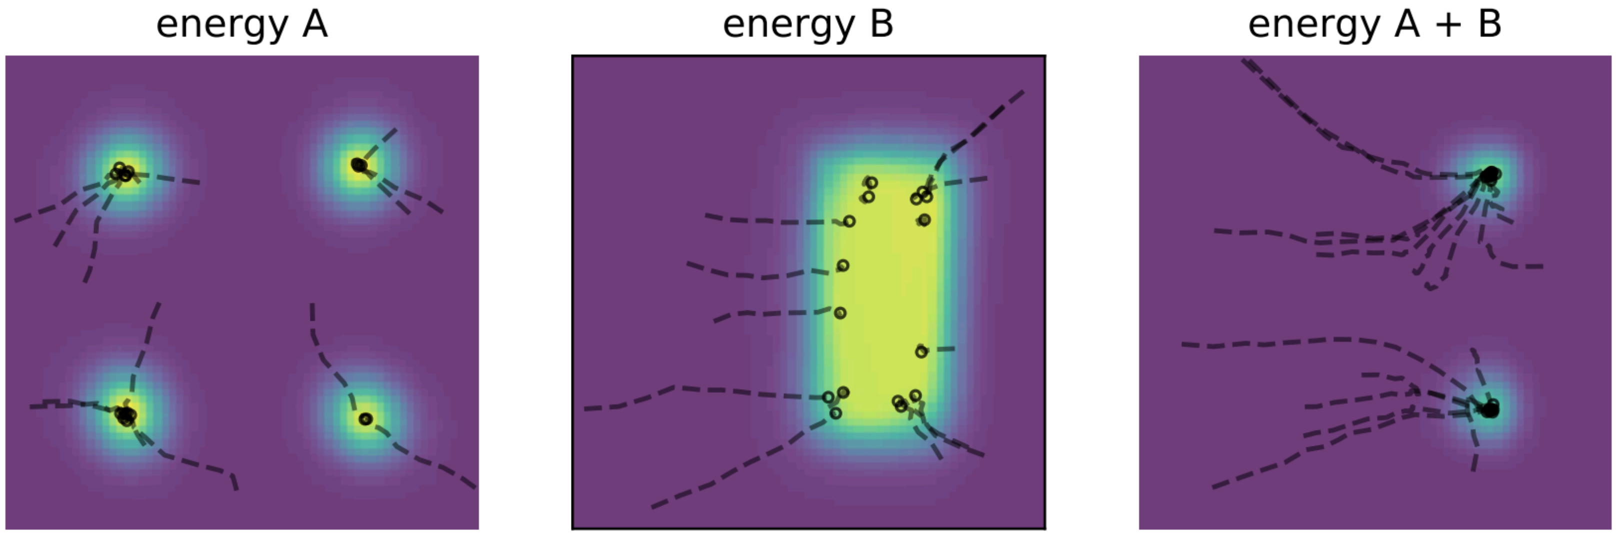

Specify a concept by successively adding constraints

Compositionality via Sum of Energies

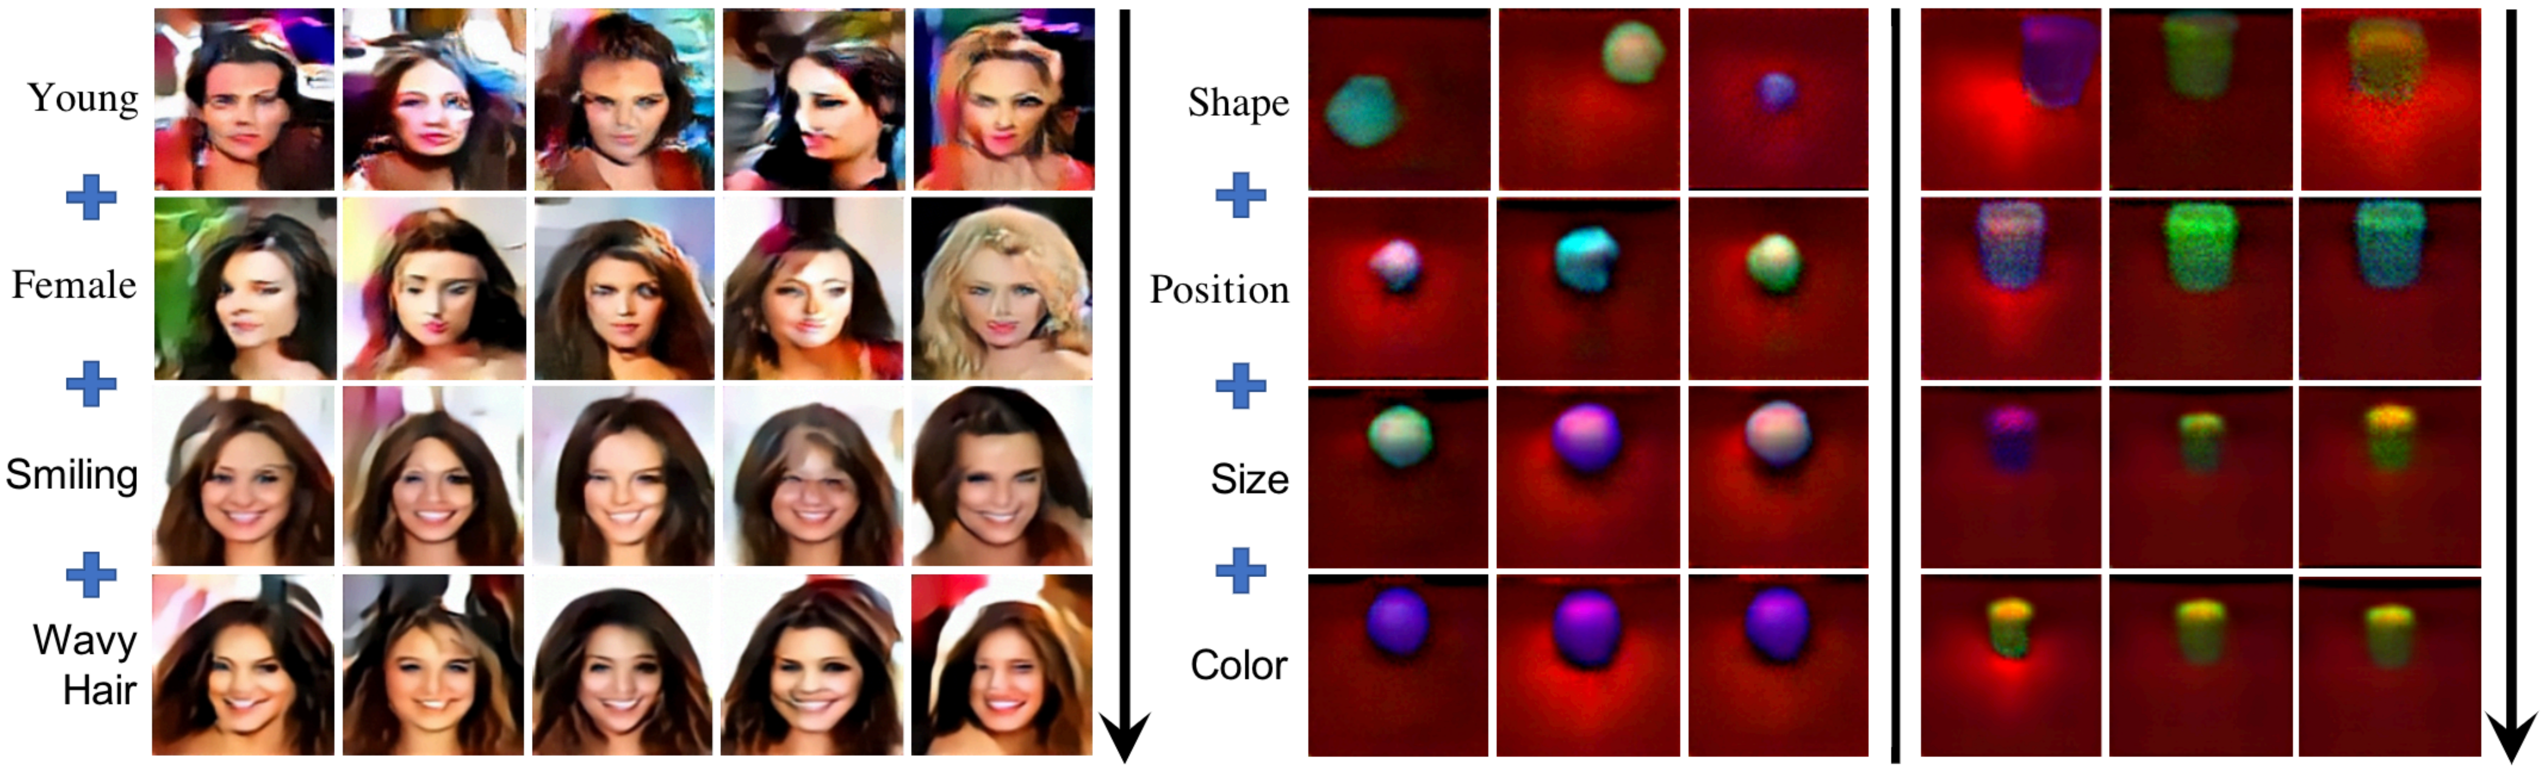

Specify a concept by successively adding constraints

Compositional Visual Generation with EBMs [Du, Li, Mordatch, 2019]

Compositionality via Sum of Energies

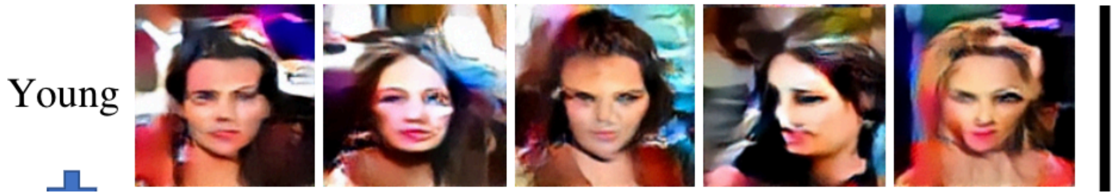

Specify a concept by successively adding constraints

Compositional Visual Generation with EBMs [Du, Li, Mordatch, 2019]

Compositionality via Sum of Energies

Specify a concept by successively adding constraints

Compositional Visual Generation with EBMs [Du, Li, Mordatch, 2019]

Compositionality via Sum of Energies

Specify a concept by successively adding constraints

Compositional Visual Generation with EBMs [Du, Li, Mordatch, 2019]

Compositionality via Sum of Energies

Specify a concept by successively adding constraints

Compositional Visual Generation with EBMs [Du, Li, Mordatch, 2019]

Surprising Benefits of Energy-Based Models

- Robustness
- Continual Learning
- Compositionality
- Trajectory Modeling

EBMs for Trajectory Modeling and Control

[Du, Lin, Mordatch, 2019]

- Train energy to model pairwise state transitions s_t, s_{t+1}
- Trajectory probability:

$$p_\theta(\tau) = p_\theta(s_1, s_2, \ldots, s_T) = \prod_{t=1}^{T-1} p_\theta(s_t, s_{t+1})$$

$$\propto \exp(-\sum_{t=1}^{T} E(s_t, s_{t+1}))$$

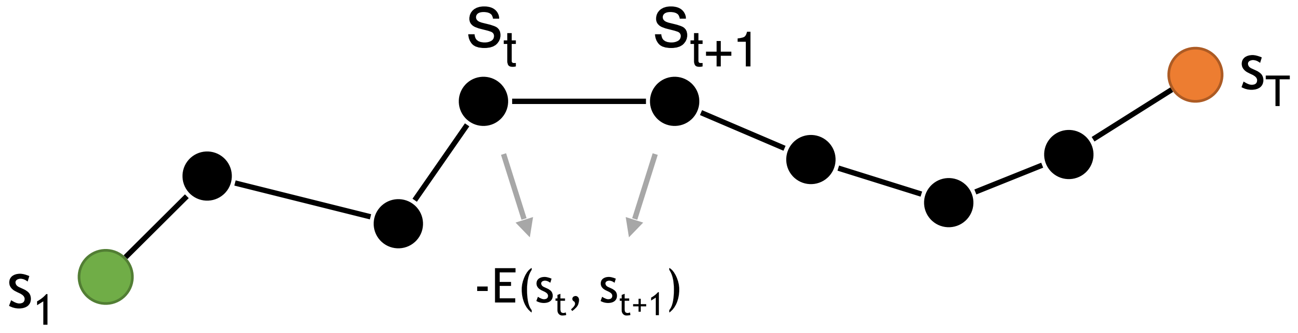

EBMs for Trajectory Modeling and Control

[Du, Lin, Mordatch, 2019]

- Train energy to model pairwise state transitions s_t, s_{t+1}
- Generate trajectories that achieve specific tasks:

$$p_\theta(s_2, \ldots, s_T | s_1, R) \propto \exp(\underbrace{-\sum_{t=1}^{T-1} E(s_t, s_{t+1})}_{\text{EBM}} \underbrace{- \sum_{t=1}^{T} R(s_t)}_{\text{Task}})$$

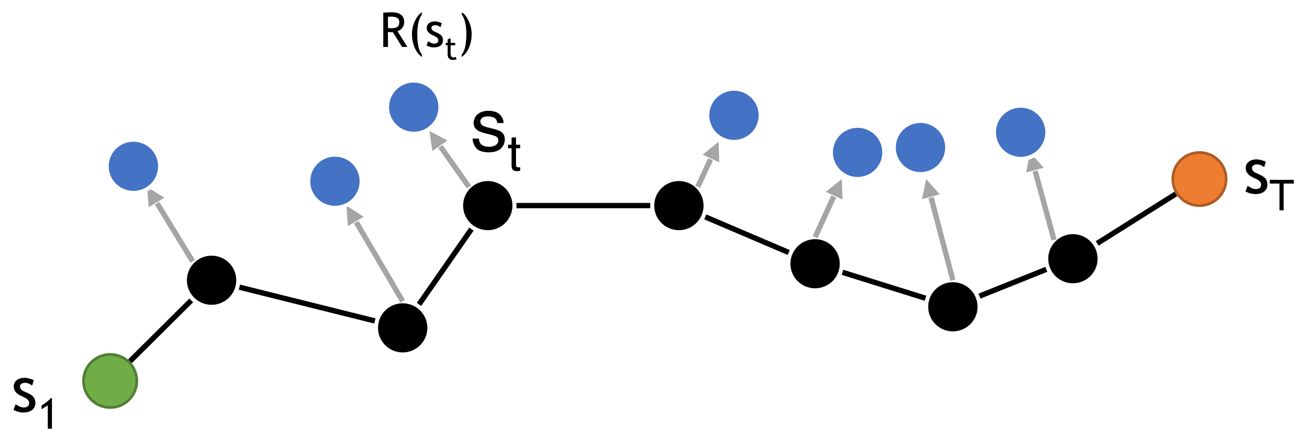

(similar to direct trajectory optimization)

EBMs for Control

Particle Environment | Maze Environment | Reacher Environment

Data	Model	Particle	Maze	Reacher
Pretrained	EBM	**-5.14**	-72.07	**-19.38**
	Action FF	-6.11	**-65.06**	-25.54
Online	EBM	**-20.38**	**-162.97**	**-29.87**
	Action FF	-850.67	-949.99	-42.37

Source Code

- Images
  - https://github.com/openai/ebm_code_release
- Trajectories
  - https://github.com/yilundu/model_based_planning_ebm
- Compositionality
  - https://drive.google.com/file/d/138w7Oj8rQl_e40_RfZJq2WKWb41NgKn3
- Interactive Notebook
  - https://drive.google.com/file/d/1fCFRw_YtqQPSNoqznIh2b1L2baFgLz4W/view