Energy-based Latent Aligner for Incremental Learning

K J Joseph[‡△], Salman Khan[△], Fahad Shahbaz Khan[△], Rao Anwer[△], Vineeth N Balasubramanian[‡]
[‡]Indian Institute of Technology Hyderabad, [△]MBZ University of AI, Abu Dhabi.
{cs17m18p100001,vineethnb}@iith.ac.in,{salman.khan,fahad.khan,rao.anwer}@mbzuai.ac.ae

CVPR JUNE 19-24 2022 NEW ORLEANS • LOUISIANA

TL;DR:

We propose an energy-based aligner which counteracts the representational shift that happens while learning incrementally.

Motivation and Overview

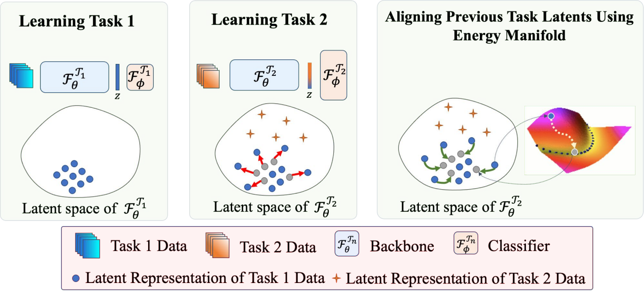

ELI: **E**nergy-based **L**atent Aligner for **I**ncremental Learning

Formulation

EBM: $p_\psi(\mathbf{z}) = \dfrac{\exp\left(-E_\psi(\mathbf{z})\right)}{\int_{\mathbf{z}} \exp\left(-E_\psi(\mathbf{z})\right) d\mathbf{z}}$ Likelihood: $L(\psi) = \mathbb{E}_{\mathbf{z}\sim p_{true}}[\log p_\psi(\mathbf{z})]$.

Gradient of Likelihood:

$$\partial_\psi L(\psi) = \mathbb{E}_{\mathbf{z}\sim p_{true}}[-\partial_\psi E_\psi(\mathbf{z})] + \mathbb{E}_{\mathbf{z}\sim p_\psi}[\partial_\psi E_\psi(\mathbf{z})].$$

Sampling z:

$$\mathbf{z}_{i+1} = \mathbf{z}_i - \frac{\lambda}{2}\partial_{\mathbf{z}} E_\psi(\mathbf{z}) + \sqrt{\lambda}\omega_i,\ \omega_i \sim \mathcal{N}(0, \mathbf{I})$$

Learning the EBM

Algorithm Algorithm LearnEBM

Input: Feature extractor trained till current task: $\mathcal{F}_\theta^{\mathcal{T}_t}$; Feature extractor trained till previous task: $\mathcal{F}_\theta^{\mathcal{T}_{t-1}}$; Current task data: $p_{data}^{\mathcal{T}_t}$

$E_\psi \leftarrow$ Initialize the Energy function.
while until required iterations **do**
  $\mathbf{x} \sim p_{data}^{\mathcal{T}_t}$
  $\mathbf{z}^{\mathcal{T}_{t-1}} \leftarrow \mathcal{F}_\theta^{\mathcal{T}_{t-1}}(\mathbf{x})$
  $\mathbf{z}^{\mathcal{T}_t} \leftarrow \mathcal{F}_\theta^{\mathcal{T}_t}(\mathbf{x})$
  $\mathbf{z}_{sampled}^{\mathcal{T}_t} \leftarrow$ Sample from EBM with $\mathbf{z}^{\mathcal{T}_t}$ as starting points.
  in_dist_energy $\leftarrow E_\psi(\mathbf{z}^{\mathcal{T}_{t-1}})$
  out_of_dist_energy $\leftarrow E_\psi(\mathbf{z}_{sampled}^{\mathcal{T}_t})$
  Likelihood $\leftarrow$ (−in_dist_energy + out_of_dist_energy)
  Optimize E_ψ to maximize *Likelihood*.
return E_ψ

Aligning the Latents

Algorithm Algorithm AlignLatents

Input: Latent vector to be adapted: $\mathbf{z}$; EBM: E_ψ; Number of Langevin steps: L_{steps}; Learning rate: λ

while until L_{steps} iterations **do**
  $grad \leftarrow \nabla_{\mathbf{z}} E_\psi(\mathbf{z})$
  $\mathbf{z} \leftarrow \mathbf{z} - \lambda * grad$
return z

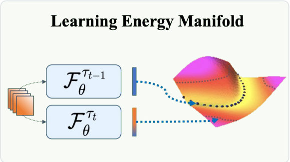

Please find more detailed analysis, results and explanations in our paper!

Toy Experiment

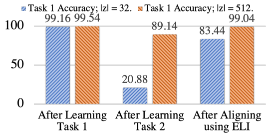

Task 1 contains the first five classes of MNIST, while Task 2 contains the rest.

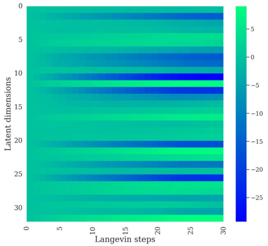

While plotting each latent dimension across Langevin iterations, we note that EBM implicitly identify which latent dimension is important to be preserved or modified.

Results on Incremental Classification and Object Detection

Settings →		Half of all the classes is used to learn the first task						Same number of classes for each task					
Datasets →		CIFAR-100			ImageNet subset			CIFAR-100			ImageNet subset		
Methods	Venue	5 Tasks	10 Tasks	25 Tasks	5 Tasks	10 Tasks	25 Tasks	5 Tasks	10 Tasks	20 Tasks	5 Tasks	10 Tasks	20 Tasks
iCaRL [41]	CVPR 17	56.97	53.28	50.98	58.24	51.6	49.02	61.59	60.05	57.81	71.46	65.25	60.21
iCaRL + ELI		63.68 + 6.71	58.92 + 5.64	54.00 + 3.02	68.94 + 10.73	61.48 + 9.88	56.11 + 7.08	**70.13** + 8.54	**67.81** + 7.75	**63.06** + 5.25	**78.51** + 7.04	**71.66** + 6.41	**66.77** + 6.56
LUCIR [20]	CVPR 19	64.37	62.57	59.91	71.38	68.99	64.65	62.01	58.95	54.2	74.22	67.97	62.2
LUCIR + ELI		66.06 + 1.69	63.50 + 0.93	60.30 + 0.39	**74.58** + 3.21	71.62 + 2.61	66.35 + 1.71	64.55 + 2.49	59.51 + 0.56	54.98 + 0.78	75.38 + 1.16	70.28 + 2.31	65.51 + 3.31
AANet [31]	CVPR 21	67.53	66.25	64.28	70.84	70.3	69.07	63.89	60.94	56.88	65.86	54.13	44.96
AANet + ELI		**68.78** + 1.25	**66.62** + 0.37	**64.72** + 0.44	73.54 + 2.73	**71.82** + 1.52	**70.32** + 1.25	66.36 + 2.47	61.72 + 0.78	57.65 + 0.77	67.43 + 1.57	55.47 + 1.34	46.93 + 1.97

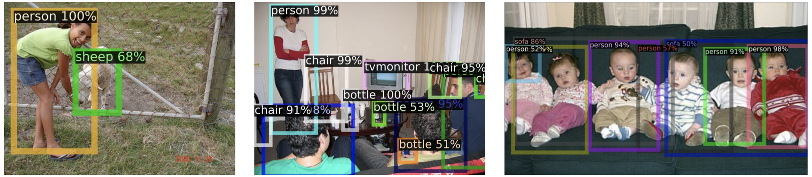

In these qualitative results of Incremental Object Detection, instances of *plant*, *sheep*, *sofa*, *train* and *tvmonitor* were introduced to a detector trained on the rest. We detect instances of old and new classes alike.