Semi-Supervised Class Incremental Learning

Alexis Lechat[1,2], Stéphane Herbin[1] and Frédéric Jurie[2]
[1]ONERA, [2]Normandie Université

MOTIVATIONS

- Models trained with Continual Learning (CL) suffer from *Catastrophic Forgetting*.
- Learning good representation is crucial for competitive performance. It requires a large visual diversity. However, incremental learning involves a restricted amount of data available to train on at each instant.
- Current fully-supervised CL methods still requires large amount of labels, making the data collection an expensive process.
- Access to a large amount of unlabeled data is inexpensive

OUR APPROACH

We consider a *Class Incremental* (CI) scenario[1]: the supervised data stream is sorted in a sequence of subsets with each one containing a new batch of classes.

Our *Semi-Supervised Incremental Learning* (SSIL) method proposes to make available an additional stream of unlabeled data that can be leveraged by the model through *self-supervision* in order to learn better features and regularize the continual training process.

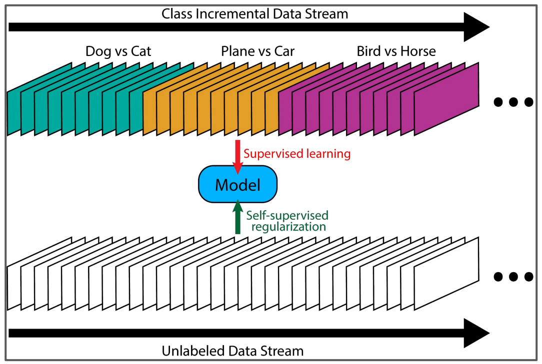

Aside from the self-supervised regularization, we also use *rehearsal learning* to alleviate the forgetting: an episodic memory stores a few samples from past classes and replays them to learn the model.

METHOD

The core of our method relies on an Adversarial Autoencoder [2] (AAE) in a semi-supervised setting. For each batch, the training is a 3-step process: reconstruction, adversarial regularization and supervised classification.

The latent representation is divided in two parts: a one-hot vector **y** representing the class/cluster and a continuous style encoding variable **z**. The adversarial losses impose that **y** follows a categorical distribution and **z** follows a gaussian distribution. This constraint on **y** forces the model to cluster the data in the representation space, labeled and unlabeled alike, given **C** clusters (the dimension of **y**) with **C** initialized at value higher than the potential maximal number of classes.

In order to learn the main incremental classification task, clusters are assigned to actual labels, i.e. when a new class is encountered, the unassigned cluster with the highest average response is attributed to this given label for the supervised learning.

The training process is illustrated in the following figure:

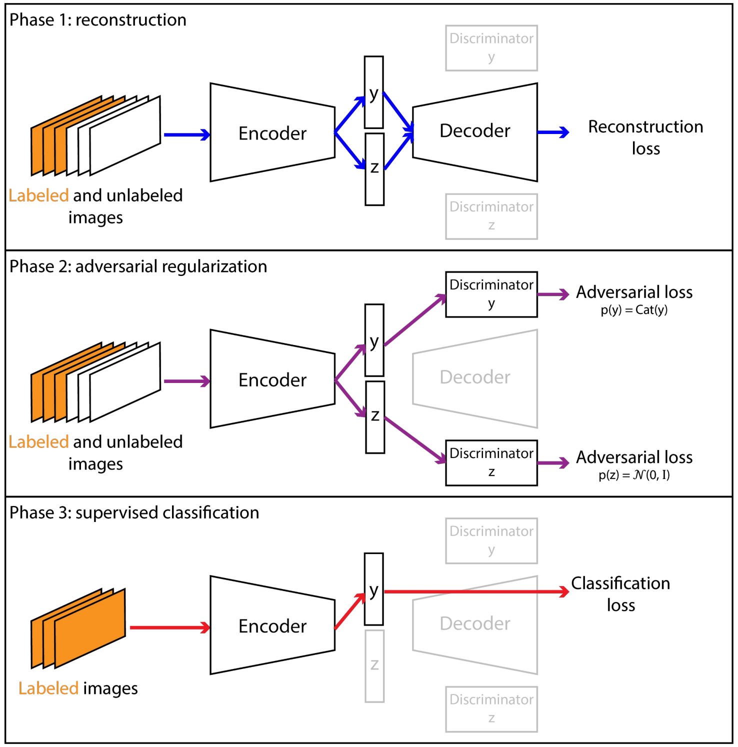

RESULTS

We test our methods on MNIST and STL-10 with a 2-class incremental. Our SSIL only uses 2,000 annotated data for MNIST while compared methods learn on the whole 60,000. We use EMNIST as a source of unlabeled data for MNIST and the provided 100,000 unlabeled images for STL-10.

Comparison of Latest and Average Accuracy of Different Class-Incremental Learning Methods on MNIST and STL-10

Method	MNIST		STL-10	
	Latest (%)	Average (%)	Latest (%)	Average (%)
Oracle	99.4	99.7	67.2	73.5
Fine-Tuning	19.8	44.9	16.2	38.3
LwF	71.3	85.2	17.9	42.5
DMC	81.1	87.4		
Naive Rehearsal	93.7	97.6	43.8	62.0
iCaRL	95.3	97.9	42.6	63.0
WA	96.0	98.3	47.3	63.5
Ours[a]	96.9	98.5	**57.3**	**72.0**
Ours[b] (EMNIST-digits)	**98.1**	**99.0**		
Ours[b] (EMNIST-letters)	95.9	98.5		

[a] Our standard baseline on MNIST uses EMNIST-full as unlabeled data stream.
[b] Additional results on MNIST benchmark when using EMNIST-digits and EMNIST-letters as unlabeled data stream instead of the whole EMNIST.

We also pre-train the models with RotNet[3] on the whole unlabeled dataset so that all methods see both the unlabeled and unlabeled data stream:

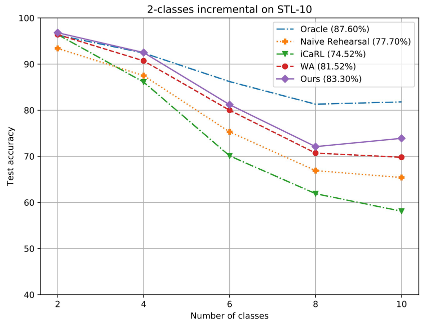

REFERENCES

[1] Rebuffi et al., *iCaRL: Incremental classifier and representation learning*, CVPR, 2017

[2] Makhzani et al., *Adversarial autoencoders*, CoRR, vol. abs/1511.05644, 2015.

[3] Gidaris et al., *Unsupervised representation learning by predicting image rotations*, ICLR, 2018