Class-incremental learning with generative classifiers

Gido van de Ven, Zhe Li & Andreas Tolias

Open-source code: https://github.com/GMvandeVen/class-incremental-learning

Class-incremental learning

- Main challenge:
  → Learn to distinguish between classes that are not observed together

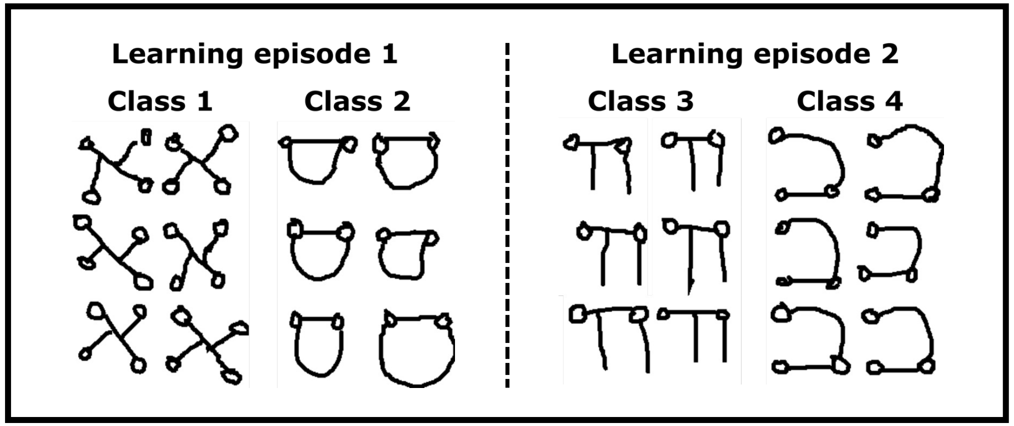

- Existing strategies:
  - ~~Storing previously seen examples~~
  - Generative replay (e.g., **DGR**; Shin et al., 2017 *NeurIPS*; **BI-R**; van de Ven et al., 2020 *Nature Communications*)
  - Parameter regularization (e.g., **EWC**; Kirckpatrick et al., 2017 *PNAS*; **SI**; Zenke et al., 2017 *ICML*)
  - Bias-correction (e.g., **CWR**; Lomonaco & Maltoni, 2017 *CoRL*; **AR1** Maltoni & Lomonaco, 2019 *Neural Networks*; "**Labels trick**" Zeno et al., 2019 *arXiv*)

Proposed strategy: generative classification

Discriminative classifiers

- Directly learn $p(y|\boldsymbol{x})$, or $\underset{y}{\operatorname{argmax}}\, p(y|\boldsymbol{x})$.

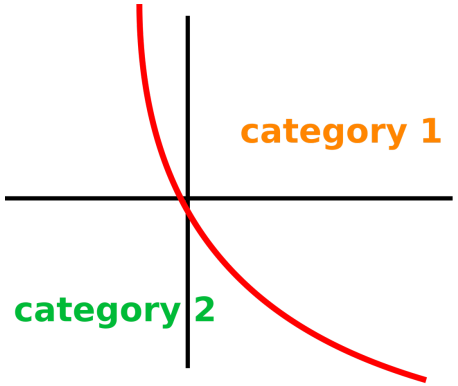

- Learn rules / shortcuts / features to distinguish between the classes to be learned
- Comparison between classes is during *training*

Generative classifiers

- Learn $p(\boldsymbol{x}, y)$, factorized as $p(\boldsymbol{x}|y)p(y)$, and classify using Bayes' rule

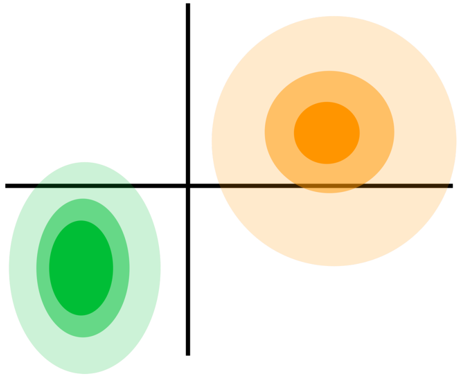

- Learn a model / template / representation for each class to be learned
- Comparison between classes is during *inference*

Generative classification ***rephrases a class-incremental problem as a task-incremental problem***, whereby each 'task' is to learn a class-conditional generative model.

(see van de Ven & Tolias, 2018 *NeurIPS workshop*)

Naïve implementation for proof-of-principle

- Separate VAE model for each class
- If a pretrained network is available, the VAE models are trained on the latent features
- Class-conditional likelihoods are estimated using importance sampling
- The *total* number of parameters is similar to that of generative replay

Schematic

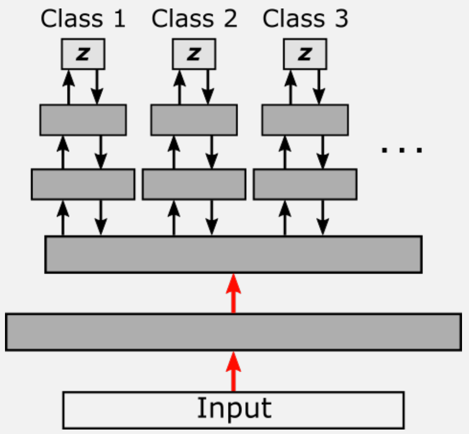

Making a classification decision

(1) Estimate class-conditional likelihoods

$$p(\boldsymbol{x}_{\text{test}}|y=1) \quad p(\boldsymbol{x}_{\text{test}}|y=2) \quad p(\boldsymbol{x}_{\text{test}}|y=3) \quad \dots$$

(2) Classify using Bayes' rule

$$\hat{y}^{(\boldsymbol{x}_{\text{test}})} = \operatorname*{argmax}_i p(\boldsymbol{x}_{\text{test}}|y=i)$$

Results

Strategy	**Method**	**MNIST**	**CIFAR-10**	**CIFAR-100**	**CORe50**
Baselines	*None*	*19.92 (± 0.02)*	*18.74 (± 0.29)*	*7.96 (± 0.11)*	*18.65 (± 0.26)*
	Joint	*98.23 (± 0.04)*	*82.07 (± 0.15)*	*54.08 (± 0.27)*	*71.85 (± 0.30)*
Generative Replay	DGR	91.30 (± 0.60)	17.21 (± 1.88)	9.22 (± 0.24)	-
	BI-R	-	-	21.51 (± 0.25)	60.40 (± 1.04)
	BI-R + SI	-	-	34.38 (± 0.21)	62.68 (± 0.72)
Regularization	EWC	19.95 (± 0.05)	18.63 (± 0.29)	8.47 (± 0.09)	18.56 (± 0.31)
	SI	19.95 (± 0.11)	18.14 (± 0.36)	8.43 (± 0.08)	18.69 (± 0.26)
Bias-correction	CWR	32.48 (± 2.64)	18.37 (± 1.61)	21.90 (± 0.68)	40.28 (± 1.13)
	CWR+	37.20 (± 3.11)	22.32 (± 1.08)	9.34 (± 0.25)	40.12 (± 1.06)
	AR1	48.84 (± 2.55)	24.44 (± 1.08)	20.62 (± 0.45)	45.27 (± 1.02)
	Labels Trick	32.46 (± 1.95)	18.43 (± 1.31)	23.68 (± 0.26)	42.59 (± 1.03)
Other	SLDA	87.30 (± 0.02)	38.35 (± 0.03)	44.49 (± 0.00)	**70.80 (± 0.00)**
Generative Classifier		**93.79 (± 0.08)**	**56.03 (± 0.04)**	**49.55 (± 0.06)**	**70.81 (± 0.11)**

All experiments were performed 10 times with different random seeds, reported is the mean (± SEM) accuracy over these runs. Code to replicate: https://github.com/GMvandeVen/class-incremental-learning.

Results: generative replay *vs.* generative classification

Strategy	Method	MNIST	CIFAR-10	CIFAR-100	CORe50
Baselines	*None*	*19.92 (± 0.02)*	*18.74 (± 0.29)*	*7.96 (± 0.11)*	*18.65 (± 0.26)*
	Joint	*98.23 (± 0.04)*	*82.07 (± 0.15)*	*54.08 (± 0.27)*	*71.85 (± 0.30)*
Generative Replay	DGR	91.30 (± 0.60)	17.21 (± 1.88)	9.22 (± 0.24)	-
	BI-R	-	-	21.51 (± 0.25)	60.40 (± 1.04)
	BI-R + SI	-	-	34.38 (± 0.21)	62.68 (± 0.72)
	EWC	19.95 (± 0.05)	18.63 (± 0.29)	8.47 (± 0.09)	18.56 (± 0.31)
Other	SLDA	87.30 (± 0.02)	38.35 (± 0.03)	44.49 (± 0.00)	**70.80 (± 0.00)**
Generative Classifier		**93.79 (± 0.08)**	**56.03 (± 0.04)**	**49.55 (± 0.06)**	**70.81 (± 0.11)**

How to best use a learned generative model?

Generate samples to train a discriminative classifier?

Generative classification?

All experiments were performed 10 times with different random seeds, reported is the mean (± SEM) accuracy over these runs. Code to replicate: https://github.com/GMvandeVen/class-incremental-learning.

Take-home message

- Generative classification is a promising strategy for class-incremental learning and a fruitful direction for future research

Funding acknowledgements

This research project has been supported by the Lifelong Learning Machines (L2M) program of the Defence Advanced Research Projects Agency (DARPA) via contract number HR0011-18-2-0025 and by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBC) contract number D16PC00003. Disclaimer: The views and conclusions contained in this presentation are those of the presenter and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, IARPA, DoI/IBC, or the U.S. Government.

DARPA

IARPA