Semi-supervised classification using generative diffusion models

Paul-Dumitru Orășan
Babeș-Bolyai University

WeADL 2024 Workshop

The workshop is organized by the Machine Learning research group (www.cs.ubbcluj.ro/ml) and the Romanian Meteorological Administration (https://www.meteoromania.ro/)

Machine Learning Research Group
MLyRE

Outline

Semi-supervised Image Classification

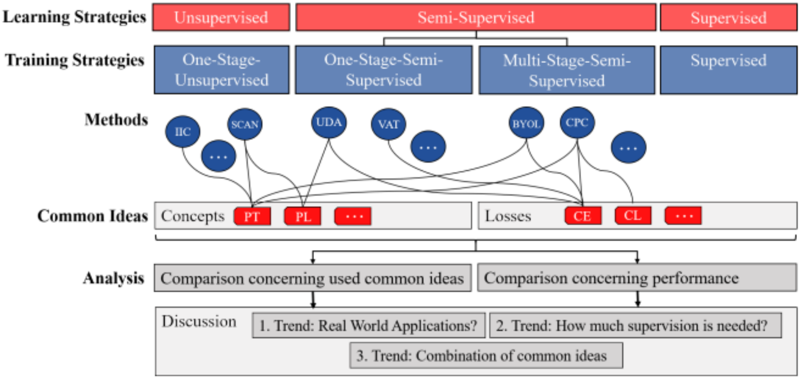

Figure: Highlight of learning and training strategies in correlation with trends [SSSK21]

Semi-supervised Image Classification

- efficient use of data
- cost-effective labelling
- adaptability to new data
- robustness to noisy labels
- real-world applications: medical imaging, satellite imagery, surveillance

"The goal of any semi-supervised model is to have better performance after training on both labeled and unlabeled data than that of a supervised model trained only on labeled data."

Rainfall nowcasting

- modeled as semi-supervised image classification
- remote sensing data can be viewed as grid maps
- example labelling:
  - **no** rain (0 mm/hr of precipitation)
  - light rain (1-4 mm/hr of precipitation)
  - moderate rain (4-10 mm/hr of precipitation)
  - heavy rain (more than 10 mm/hr of precipitation)

Rainfall nowcasting - visualization

Rainfall nowcasting - visualization

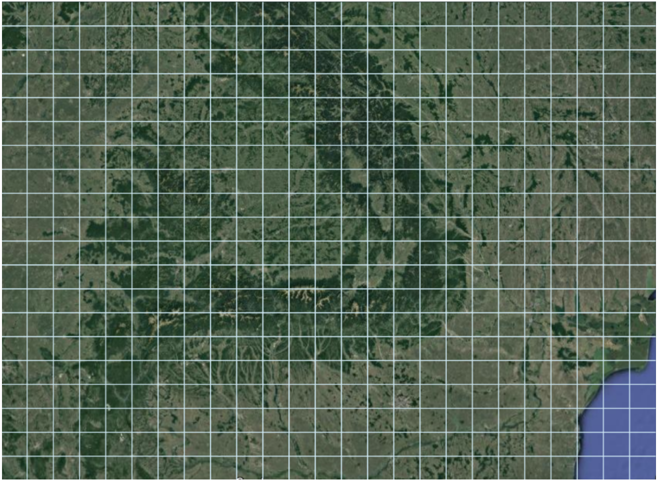

Rainfall nowcasting - visualization

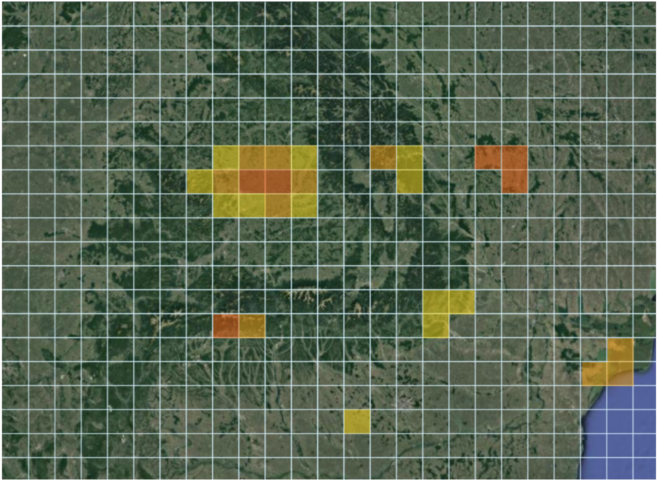

Pseudo-labelling based methods

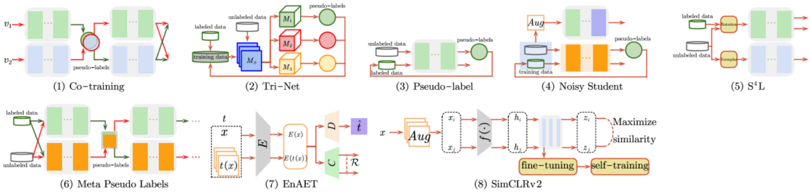

Figure: Overview of pseudo-labeling semi-supervised mehods [YSKX23]

Methods applied for rainfall nowcasting

- Supervised deep-learning approaches
  - Convolutional Neural Networks [ACM+22]
  - Vision Transformers [CSZ+23]
- Generative approaches
  - Generative Adversarial Networks [NZW+23]
  - Denoising Diffusion Probabilistic Models
- Semi-supervised deep-learning approaches
  - Graph Neural Networks [MWH+20]

Diffusion models

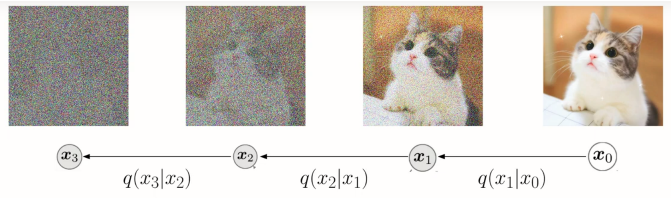

Figure: Forward diffusion process - adding noise each step

$$q(x_t|x_{t-1}) = \mathcal{N}(\sqrt{\alpha_t}x_{t-1}, (1 - \alpha_{t-1})I) \tag{1}$$

Diffusion models

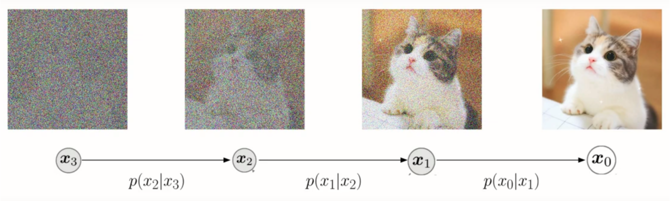

Figure: Backward diffusion process - removing noise each step

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (2)$$

LRA-Diffusion

- CARD [HZZ22] extends diffusion models to classification and regression tasks
- LRA-Diffusion [CZY+23] robust to noisy labels

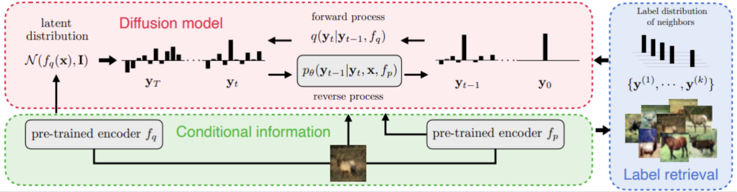

Figure: Overview of LRA-Diffusion [CZY+23]

Proposal

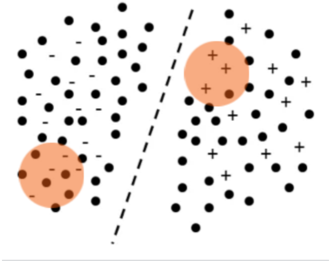

Figure: Learning neighboring distribution

$$\boxed{\begin{aligned} p(\bar{y}|x) &= \mathcal{N}(f_\phi(x), I) \\ \bar{y} &= \{y_1, y_2, \ldots, y_k\} \end{aligned}} \tag{3}$$

Semi-supervised pipeline

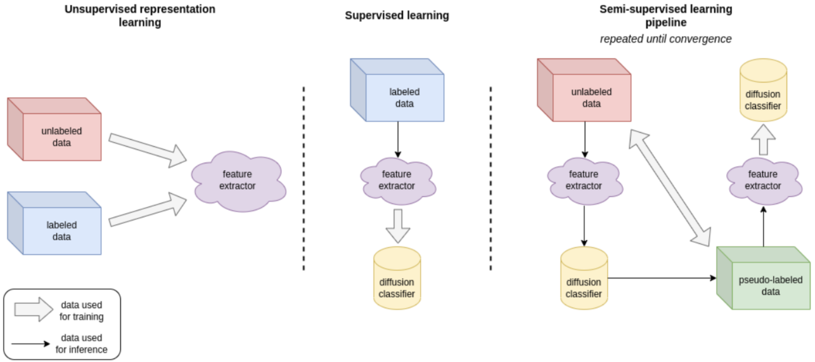

Figure: Full overview of the semi-supervised pipeline

Roadmap

- Phase 1 - implementation and validation (completed)
  - Building the codebase for the training framework
  - Running experiments to validate the semi-supervised hypothesis
- Phase 2 - rainfall nowcasting application (in progress)
  - Training feature extractors for remote sensing data using contrastive learning
  - Running experiments to validate the semi-supervised hypothesis

Datasets

- CIFAR-10 (60000 labels, 10 classes)
  - 40 labels
  - 250 labels
  - 4000 labels
- CIFAR-100 (60000 labels, 20 classes divided into 600 subclasses)
  - 400 labels
  - 2500 labels
  - 10000 labels
- Remote sensing data

Experiments

Method \ Dataset	CIFAR-10			CIFAR-100		
	40	250	4000	400	2500	10000
Supervised	62.73	75.19	82.64	33.92	38.12	53.65
Semi-supervised	87.15	92.37	**96.79**	47.83	69.42	75.84
SoTA	**94.40**	**95.16**	96.04	**62.19**	**74.93**	**79.42**

Table: Quantitative evaluation of our method. Top row describes the results achieved by training our method using only labeled data. The middle row shows the results achieved by training our method using both labeled and unlabeled data. The last row describes the current state-of-the-art benchmark results.

Proportion of data

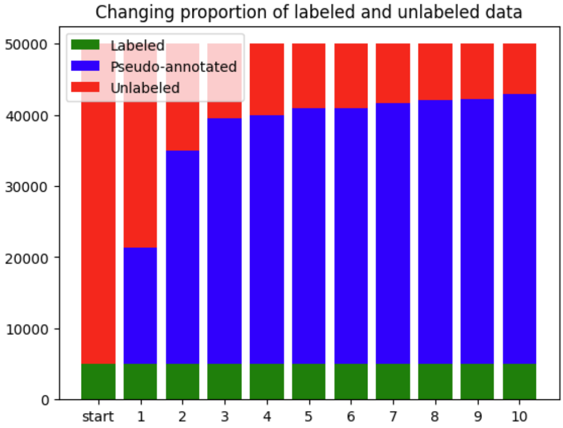

Figure: Proportion of training dataset per model iteration

Correctness of pseudo-labels

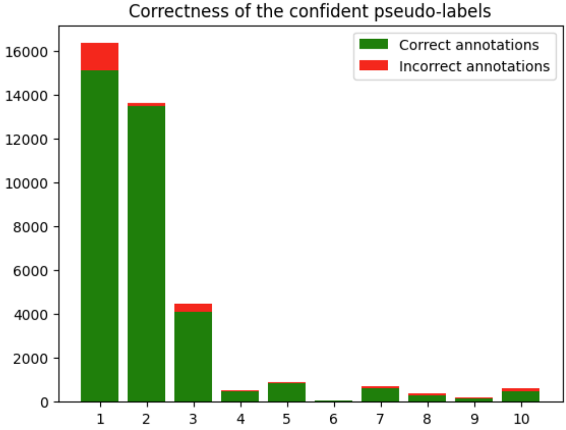

Figure: Robustness of confidence evaluation methods

Limitations and challenges

- high computational demands;
- lack of resources;
- scarcity of public, open-source feature extractors for remote sensing data
- amplified performance hit under heavy stress (less than 1% labeled data)

References I

Alexandra-Ioana Albu, Gabriela Czibula, Andrei Mihai, Istvan Gergely Czibula, Sorin Burcea, and Abdelkader Mezghani, *NeXtNow: A Convolutional Deep Learning Model for the Prediction of Weather Radar Data for Nowcasting Purposes*, Remote Sensing **14** (2022), no. 16, ID: 3890.

Shengchao Chen, Ting Shu, Huan Zhao, Guo Zhong, and Xunlai Chen, *TempEE: Temporal–Spatial Parallel Transformer for Radar Echo Extrapolation Beyond Autoregression*, IEEE Transactions on Geoscience and Remote Sensing **61** (2023), 1–14.

References II

Jian Chen, Ruiyi Zhang, Tong Yu, Rohan Sharma, Zhiqiang Xu, Tong Sun, and Changyou Chen, *Label-retrieval-augmented diffusion models for learning from noisy labels*, Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023 (Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, eds.), 2023.

References III

Xizewen Han, Huangjie Zheng, and Mingyuan Zhou, *CARD: classification and regression diffusion models*, Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022 (Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, eds.), 2022.

Kaichao Miao, Wei Wang, Rui Hu, Lei Zhang, Yali Zhang, Xiang Wang, and Fudong Nian, *Multimodal semisupervised deep graph learning for automatic precipitation nowcasting*, Mathematical Problems in Engineering **2020** (2020), 1–9.

References IV

- Xianghua Niu, Lixia Zhang, Chunlin Wang, Kailing Shen, Wei Tian, and Bin Liao, *A Generative Adversarial and Spatiotemporal Differential Fusion Method in Radar Echo Extrapolation*, Remote Sensing **15** (2023), ID: 5329, 1–17.
- Lars Schmarje, Monty Santarossa, Simon-Martin Schröder, and Reinhard Koch, *A survey on semi-, self- and unsupervised learning for image classification*, IEEE Access **9** (2021), 82146–82168.
- Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu, *A survey on deep semi-supervised learning*, IEEE Trans. Knowl. Data Eng. **35** (2023), no. 9, 8934–8954.