Semi-supervised learning

Samuel Cheng

University of Oklahoma

March 6, 2025

What is Semi-Supervised Learning?

Definition: Semi-supervised learning (SSL) is a paradigm that leverages a small amount of labeled data and a large amount of unlabeled data to train a model.

- Bridges the gap between supervised and unsupervised learning.
- Useful when labeled data is scarce but unlabeled data is abundant.

Goal: Improve learning efficiency by using both types of data.

Annotation Challenges

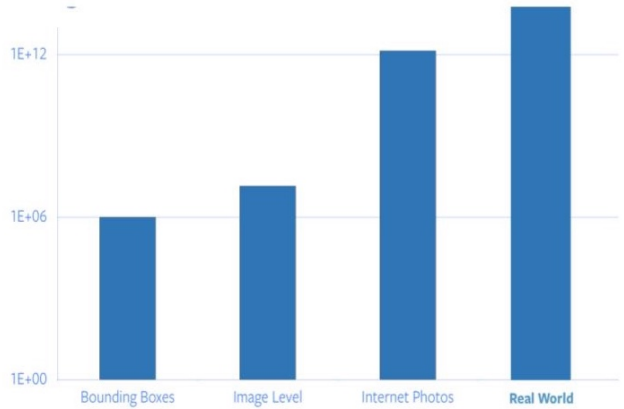

Variation in available data quantum basis annotation complexity

Annotation Challenges

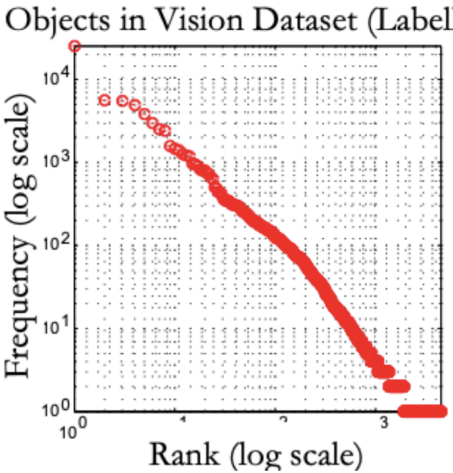

10% of the classes account for 93% of the data

Why Semi-Supervised Learning?

Challenges in Supervised Learning:

- Labeled data is often expensive and time-consuming to acquire.
- Certain domains (e.g., medical imaging, genomics) require expert annotation.

Solution: Semi-supervised learning reduces reliance on labeled data while leveraging the abundance of unlabeled data.

Key Approaches in Semi-supervised Learning

- Model-Generated Labeling (Pseudo-Labeling)
- Consistency Regularization

Common Assumptions:

- Clustering assumption: Similar data points belong to the same class.
- Smoothness assumption: Decision boundaries should lie in low-density regions.

Model-Generated Labeling (Pseudo-Labeling)

Process:

1. Train an initial model on the small labeled dataset.
2. Use the model to predict labels for unlabeled examples.
3. Select confident predictions and treat them as ground truth.
4. Retrain the model with the expanded dataset.

Benefit: Expands training data without human intervention.

Reference: Lee, D. H. (2013). "Pseudo-Label: The simple and efficient semi-supervised learning method for deep neural networks." Workshop on Challenges in Representation Learning, ICML 2013.

Model-Generated Labeling (Pseudo-Labeling)

Process:

1. Train an initial model on the small labeled dataset.
2. Use the model to predict labels for unlabeled examples.
3. Select confident predictions and treat them as ground truth.
4. Retrain the model with the expanded dataset.

Benefit: Expands training data without human intervention.

Reference: Lee, D. H. (2013). "Pseudo-Label: The simple and efficient semi-supervised learning method for deep neural networks." Workshop on Challenges in Representation Learning, ICML 2013.

Model-Generated Labeling (Pseudo-Labeling)

Process:

1. Train an initial model on the small labeled dataset.
2. Use the model to predict labels for unlabeled examples.
3. Select confident predictions and treat them as ground truth.
4. Retrain the model with the expanded dataset.

Benefit: Expands training data without human intervention.

Reference: Lee, D. H. (2013). "Pseudo-Label: The simple and efficient semi-supervised learning method for deep neural networks." Workshop on Challenges in Representation Learning, ICML 2013.

Model-Generated Labeling (Pseudo-Labeling)

Process:

1. Train an initial model on the small labeled dataset.
2. Use the model to predict labels for unlabeled examples.
3. Select confident predictions and treat them as ground truth.
4. Retrain the model with the expanded dataset.

Benefit: Expands training data without human intervention.

Reference: Lee, D. H. (2013). "Pseudo-Label: The simple and efficient semi-supervised learning method for deep neural networks." Workshop on Challenges in Representation Learning, ICML 2013.

Consistency Regularization

Core Idea: The model should produce the same output for an input and its perturbed version.

- Enforces robustness to small changes in the input.
- Encourages smooth decision boundaries.

Technique:

- Training consistency constraints (e.g., Mean Teacher, VAT, Pi-Model).

Reference: Sajjadi, M., Javanmardi, M., and Darrell, T. (2016). "Regularization with stochastic transformations and perturbations for deep semi-supervised learning." Advances in Neural Information Processing Systems (NeurIPS), 2016.

Consistency Regularization vs. Data Augmentation

Key Differences:

- Data Augmentation: Introduces variations in input data to improve generalization but does not explicitly enforce consistency in model predictions.
- Consistency Regularization: Actively enforces stable model predictions under perturbations, ensuring smooth decision boundaries.

Example:

- Data Augmentation: Random cropping, flipping, color jittering.
- Consistency Regularization: Enforcing prediction similarity under transformations via loss terms.

Examples of Consistency Loss

Common Loss Functions Used in Consistency Regularization:

- Mean Squared Error (MSE): Penalizes differences between the original and perturbed outputs.

$$L_{\text{consistency}} = \|f(x) - f(\tilde{x})\|^2$$

- KL Divergence (Kullback-Leibler): Measures the difference between two probability distributions.

$$L_{\text{consistency}} = D_{KL}(p(y|x) \parallel p(y|\tilde{x}))$$

Intuition: The model should produce similar outputs for an input and its augmented version.

Mean Teacher Model

Concept: Uses two networks, a student and a teacher, to enforce prediction consistency.

- The teacher model is an exponential moving average (EMA) of the student.
- The student is trained to minimize the difference between its predictions and the teacher's.
- Helps improve stability and prevents overfitting.

Reference: Tarvainen, A. and Valpola, H. (2017). "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results." NeurIPS 2017.

Virtual Adversarial Training (VAT)

Core Idea: Finds adversarial perturbations that cause maximum change in the model's output and penalizes them.

- First, find the adversarial perturbation that maximizes the divergence in model output:

$$r^* = \arg\max_r D_{KL}(p(y|x) \parallel p(y|x+r))$$

- Solve using gradient ascent:

$$g = \nabla_r D_{KL}(p(y|x) \parallel p(y|x+r))$$

- Normalize and scale the perturbation:

$$r^* = \epsilon \frac{g}{\|g\|}$$

- Then, train the model to minimize the divergence under this perturbation:

$$L_{VAT} = D_{KL}(p(y|x) \parallel p(y|x+r^*))$$

Reference: Miyato, T., Maeda, S.-I., Koyama, M., and Ishii, S. (2018). "Virtual adversarial training: A regularization method for supervised and semi-supervised learning." *IEEE TPAMI 2018.*

Pi-Model

Concept: Ensures prediction consistency by training on multiple perturbed versions of each input.

- Passes the same input through the model twice with different augmentations or dropout.
- Computes a consistency loss to align the predictions.
- Encourages the model to learn stable representations.

Why is it called Pi-Model?

- The name "Pi-Model" is derived from the use of two different perturbed versions of the same input x, enforcing prediction consistency between them.
- The two views form a symmetry, akin to the mathematical notion of π, which is often associated with duality and balance.

Reference: Laine, S., and Aila, T. (2017). "Temporal ensembling for semi-supervised learning." *ICLR 2017.*

Summary

- Semi-supervised learning helps models learn efficiently with minimal labeled data.
- Pseudo-labeling: Model generates labels for unlabeled data.
- Consistency Regularization: Model is trained to be invariant to small input perturbations.
- Techniques such as Mean Teacher, VAT, and Pi-Model help enforce consistency and improve learning.

Impact: Semi-supervised learning improves generalization and reduces reliance on costly labeled datasets.