Discovering Global False Negatives On the Fly for Self-supervised Contrastive Learning

Vicente Balmaseda[1], Bokun Wang[1], Ching-Long Lin[2], Tianbao Yang[1] [1]Texas A&M University, [2]University of Iowa

Abstract. We propose GloFND, a scalable method that discovers false negatives globally in contrastive learning by learning per-sample thresholds on the fly, with minimal overhead.

Contrastive Learning (CL)

CL learns generalizable representations for data by pushing apart *negative* pairs (augmentations from different images) and pulling together *positive* samples (augmentations from same image)

$$-\tau \log \frac{\exp\left(\text{sim}(\mathbf{z}_i, \mathbf{z}_i')/\tau\right)}{\sum_{\mathbf{x}\in\mathcal{S}_i^-} \exp(\text{sim}(\mathbf{z}_i, E_{\mathbf{w}}(\mathbf{x}))/\tau)}$$

False Negative Problem

Some *negatives* are semantically similar and pushing them apart hurts learning

Local/batch-based methods → unreliable and require large batch size
Global clustering → high computational cost

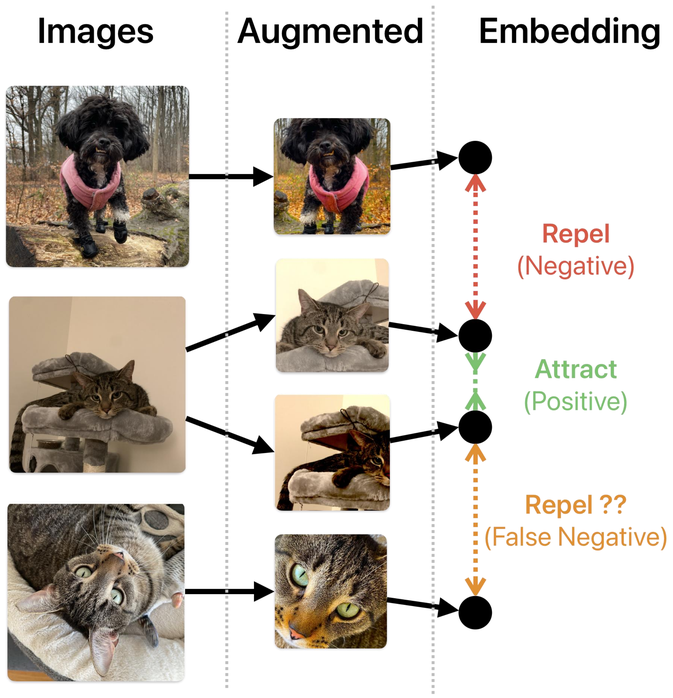

Global False Negative Detection (GloFND)

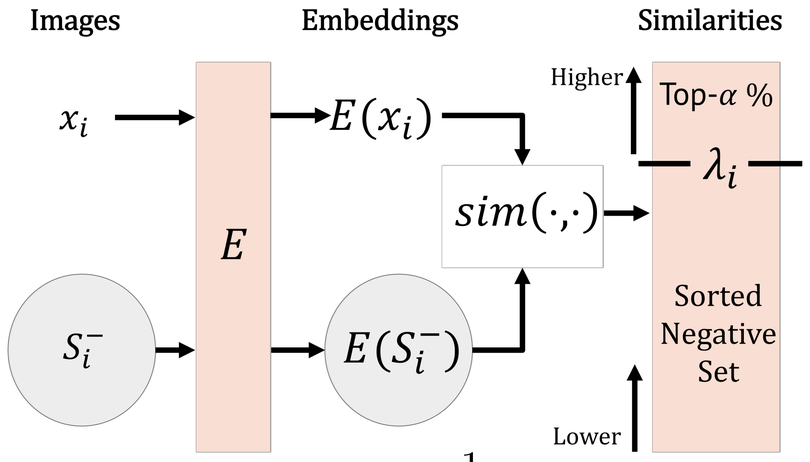

$$\lambda_i = \arg\min_{\nu\in[-1,1]} \nu\alpha + \frac{1}{|R|}\sum_{r\in R_i}(r-\nu)_+$$

For each anchor within the mini-batch

1. **Similarity** computation within mini-batch
2. **Update λ_i** through stochastic gradient descent
3. **Mark as false negatives** negative samples with similarity higher than λ_i
4. **Filter** false negatives from contrastive loss
5. Update model

GloFND's Characteristics

- **Per-sample**: each sample has its own λ_i threshold
- **Global**: threshold for top-α% across entire dataset
- **Dynamic**: learns with the network and adapts to it
- **Scalable computation**: independent of dataset size
- **Efficient**: simple and cheap updates and filtering
  (2% overhead on Imagenet100 with 128 batch size)

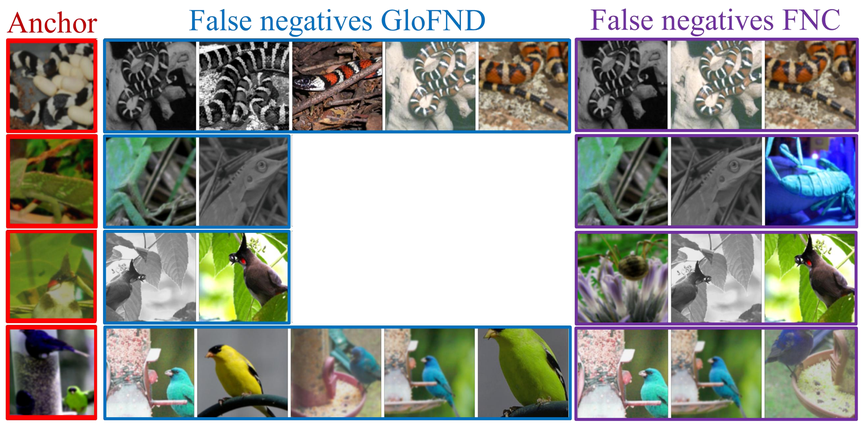

Contributions

1. **Global** false negative discovery as top-α% problem
   - Detects false negatives of each anchor across the entire dataset
2. **Efficient algorithm** with minimal overhead
   - Learns a **dynamic global per-anchor** threshold while keeping computation within the mini-batch
3. **Out-of-the-box better** false negative discovery
   - Improves representations and downstream performance without major changes to contrastive algorithm

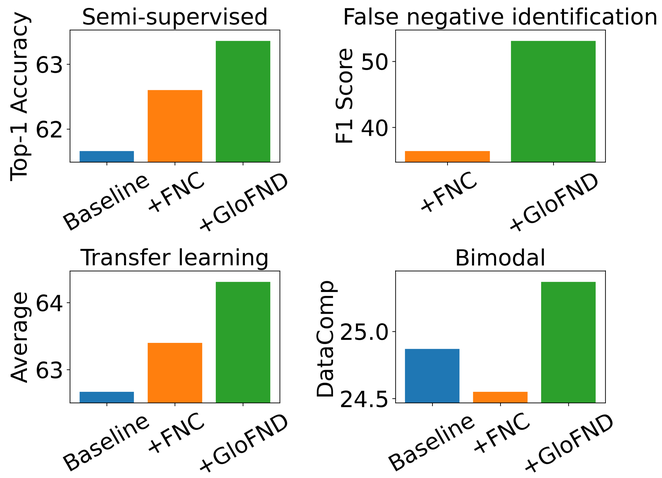

(Left) **GloFND improves detection of false negative pairs**
(Right) **GloFND learns higher quality threshold distribution compared to mini-batch approach**

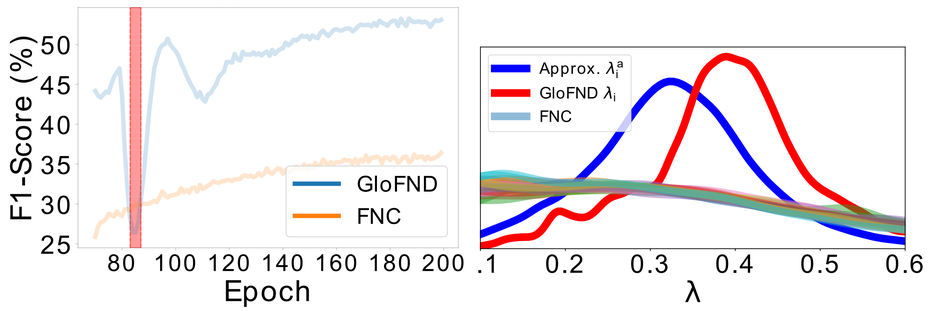

Key Takeaways

1.**GloFND** is a lightweight, optimization-based method that dynamically learns per-sample similarity thresholds—**detecting false negatives globally with only mini-batch computations**

2.**More accurate false negative detection = better representations** in unimodal, bimodal, and semi-supervised settings **without major changes to contrastive algorithms**