Week 15 Report

Thomas Deatherage
NFHM/BioCosmos
29 November 2024

Time Slot

1) What progress did you make in the last week?

- Catchup with Roman
- Explore HiPerGator
- Short week: Thanksgiving holiday so not too much coding or analysis

2) What are you planning on working on next?

- Make the unicom code we have distributed again

3) Is anything blocking you from getting work done?

- Nope.

Abstracts & Summaries

SimpleShot: Revisiting Nearest-Neighbor Classification for Few-Shot Learning
Abstract: Learning feature embedding directly from images without any human supervision is a very challenging and essential task in the field of computer vision and machine learning. Following the paradigm in supervised manner, most existing unsupervised metric learning approaches mainly focus on binary similarity in Euclidean space. However, these methods cannot achieve promising performance in many practical applications, where the manual information is lacking and data exhibits non-Euclidean latent anatomy. To address this limitation, we propose an Unsupervised Hyperbolic Metric Learning method with Hierarchical Similarity. It considers the natural hierarchies of data by taking advantage of Hyperbolic metric learning and hierarchical clustering, which can effectively excavate richer similarity information beyond binary in modeling. More importantly, we design a new loss function to capture the hierarchical similarity among samples to enhance the stability of the proposed method. Extensive experimental results on benchmark datasets demonstrate that our method achieves state-of-the-art performance compared with current unsupervised deep metric learning approaches.

Summary (Claude AI): The paper introduces a novel approach called "Unsupervised Hyperbolic Metric Learning with Hierarchical Similarity" that addresses a fundamental challenge in computer vision: learning how to measure similarity between images without having labeled training data.
The key innovation lies in how it handles the natural hierarchical relationships between images. For example, if you have pictures of different birds, some species are more closely related than others - a Green Jay is more similar to a Florida Jay than to an Ovenbird because they belong to the same family. Previous methods typically treated similarity as binary (similar/not similar), missing these nuanced relationships.
The authors solve this problem through two main technical innovations:
First, they use hyperbolic geometry instead of traditional Euclidean space. Hyperbolic space naturally represents hierarchical relationships better - imagine it like a curved surface that expands exponentially as you move outward from the center, allowing more room to represent complex hierarchical structures.
Second, they develop a new loss function that captures different degrees of similarity between images. Rather than just saying two images are similar or different, their method can express that some images are more different than others, creating a more nuanced understanding of relationships between images.
The method works in two main steps:

It maps images into hyperbolic space using a deep neural network
It performs hierarchical clustering in this space to discover natural groupings and relationships between images

The researchers tested their approach on several standard computer vision datasets (CARS196, CUB-200-2011, and Stanford Online Products) and achieved state-of-the-art results. For example, on the CARS196 dataset, their method improved the Recall@1 metric by 4.7% compared to the previous best approach.
What makes this work particularly significant is that it achieves these results without requiring any labeled training data, which is usually expensive and time-consuming to obtain. The method can automatically discover meaningful relationships between images just by looking at the images themselves.
This approach could be valuable in many real-world applications where labeled data is scarce, such as organizing large collections of images, visual search systems, or identifying similar products in e-commerce platforms.

Relevance: This paper was referenced by the UNICOM paper. But what was particularly appealing is the hierarchical clustering approach. UNICOM uses k-means clustering which is very “flat”. We’re working with ecological data which is taxonomically hierarchical. I’m curious to know if a hierarchical approach to clustering might be a reasonable avenue for experimentation in the future.

Scripts and Code Blocks

No real code contributions to show this week. Explored the HiPerGator interface.

Flow Charts/Diagrams

N/A

Documentation

No documentation to add really just yet

Results Visualization + Proof of Work

Next Week’s Proposal

- Distributed training with Unicom