Enhancing Continual Learning in Diabetic Retinopathy: Multimodal Zero-shot Clustering and Strategic Experience Replay

Gusseppe Bravo-Rocca, Peini Liu, Jordi Guitart
Ajay Dholakia, David Ellison & Rodrigo Castillo-Larco
Barcelona Supercomputing Center, Universitat Politecnica de Catalunya
Lenovo Infrastructure Solutions Group, Emory University

Introduction

- **Continual learning** enables machine learning models to adapt to new data without **forgetting** previous knowledge and **revisiting** all past samples, which is crucial for applications like **diabetic retinopathy detection**.
- In this paper, we introduce a framework that gets **LLM-generated descriptions** and **zero-shot clustering** to improve Experience Replay on new tasks.
- Our method combines a **strategic experience replay** designed to learn from multiple tasks whilst retaining a good performance and protecting **data privacy**.

Method

Our method uses a **LLM** to generate **descriptions** d_i for each image x_i, using **metadata** m_i and y_i for initial **domain learning** in Task 0 (**supervised phase**). These descriptions underpin **unsupervised zero-shot clustering**, forming $|\{y_i\}|$ clusters. **Key points** from these clusters are buffered for **replay** (e_i). A **multi-head classifier** leverages this buffer in an **Experience Replay** strategy, learning the pertinent head i for **predictions** y, thus preserving **knowledge** across successive tasks (**unsupervised phase**).

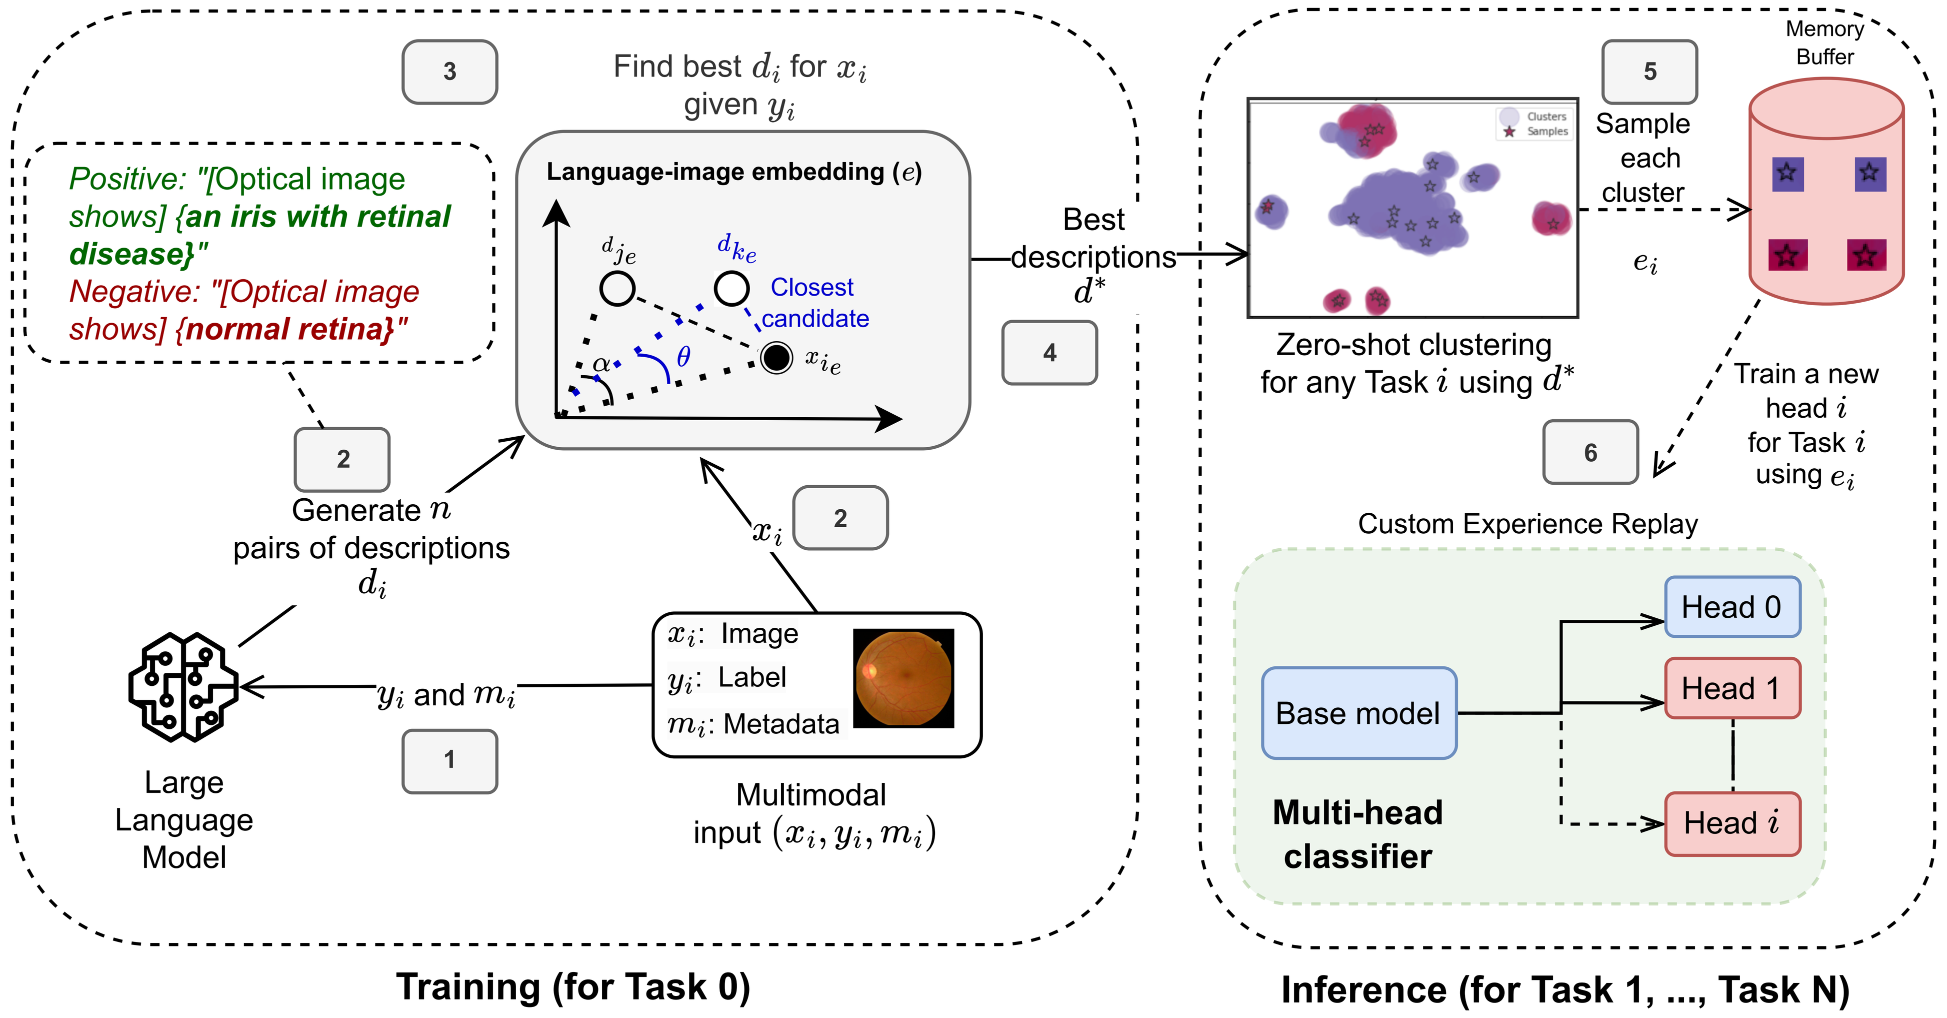

To improve continual learning in **diabetic retinopathy** detection, we combine LLM-generated descriptions and **zero-shot clustering** to obtain **better points** to **replay**.

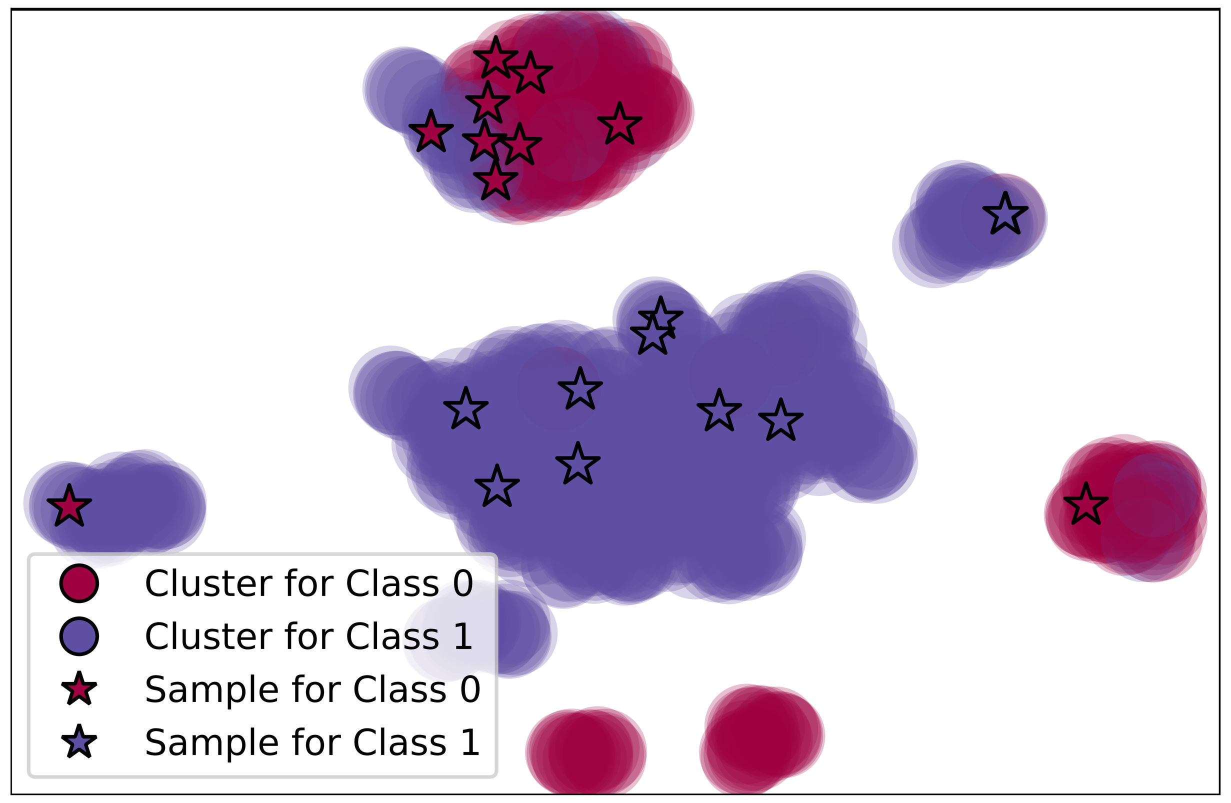

Dataset description: three tasks for continual learning.

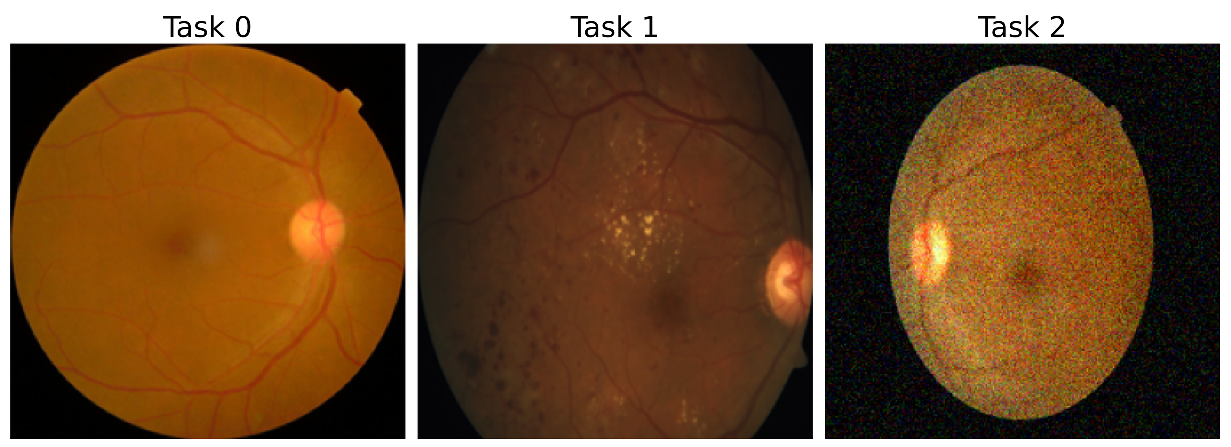

Fundus images representing different tasks with varying image quality and conditions. From left to right: Task 0 shows a fundus photograph with uniform image quality; Task 1 is an image with some variation in lighting; Task 2 displays an image with added Gaussian noise to simulate a challenging imaging condition.

Results and Conclusion

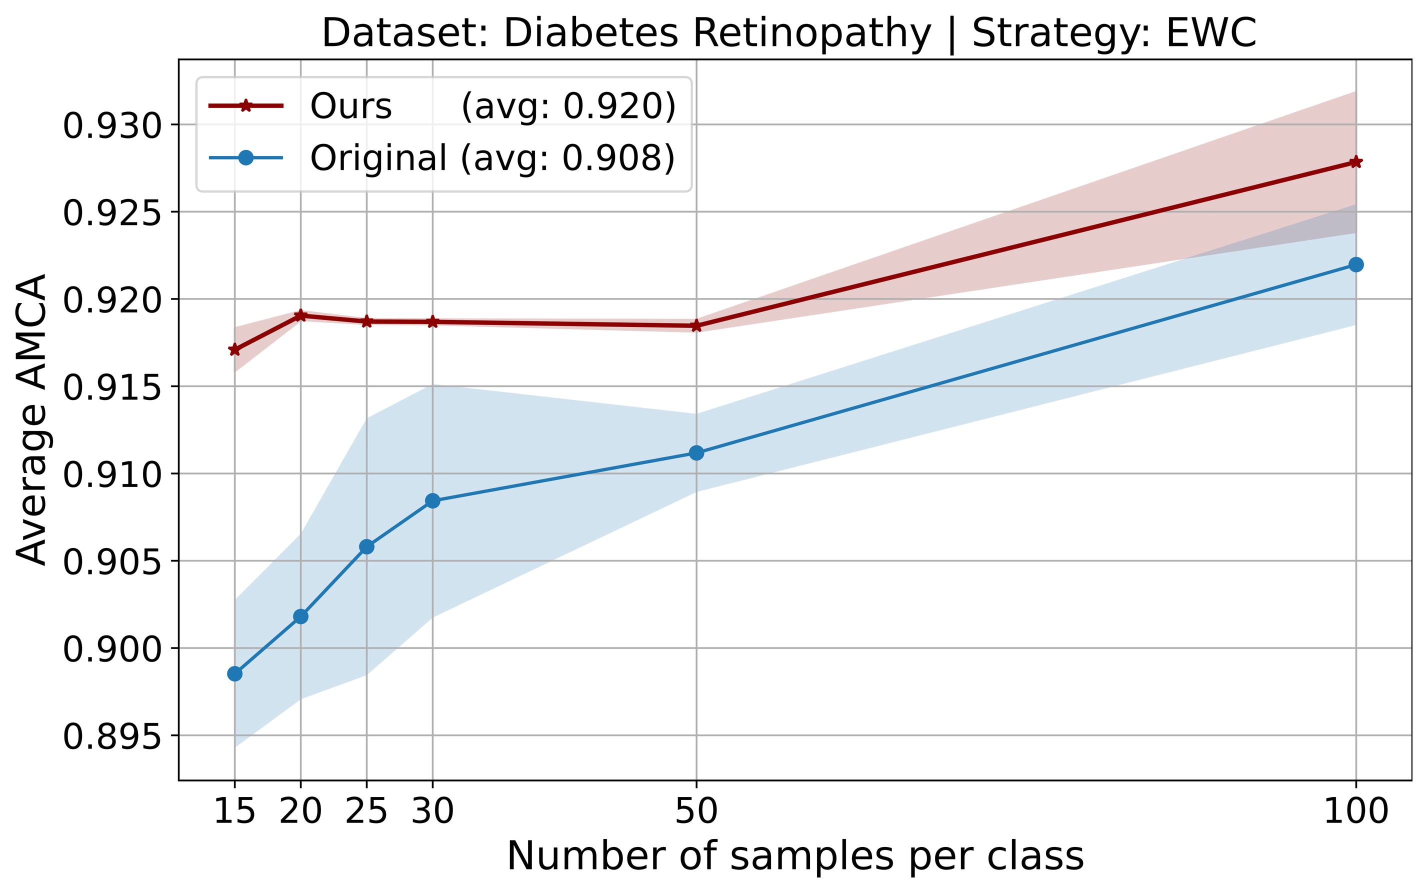

Average AMCA improvement using EWC strategy for diabetic retinopathy detection.

- Our approach significantly improves **Average Mean Class Accuracy (AMCA)** across various continual learning strategies.
- The results show consistent enhancement in **Naive, EWC, LwF,** and **GEM** strategies, even with fewer samples per class.
- This framework is particularly effective for **diabetic retinopathy detection**, showcasing its potential in practical applications.
- Future work should focus on scalability, **privacy considerations**, and further evaluations in clinical settings to ensure robust performance and ethical compliance.