POLITECNICO MILANO 1863

Exploring Concept Drift Visualization and Explanation in Image Streams

Discovering Drift Phenomena in Evolving Landscape (DELTA 2024) Workshop at ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2024)

Giacomo Ziffer and Emanuele Della Valle

DEIB – Politecnico di Milano

26.08.2024

Agenda

1. **Beyond i.i.d. assumption in image streams**
2. **Problem Statement**
3. **Proposed Methodologies**
4. **Preliminary Results**
5. **Conclusions and Future Works**

Offline Learning vs Online Learning

Offline Learning

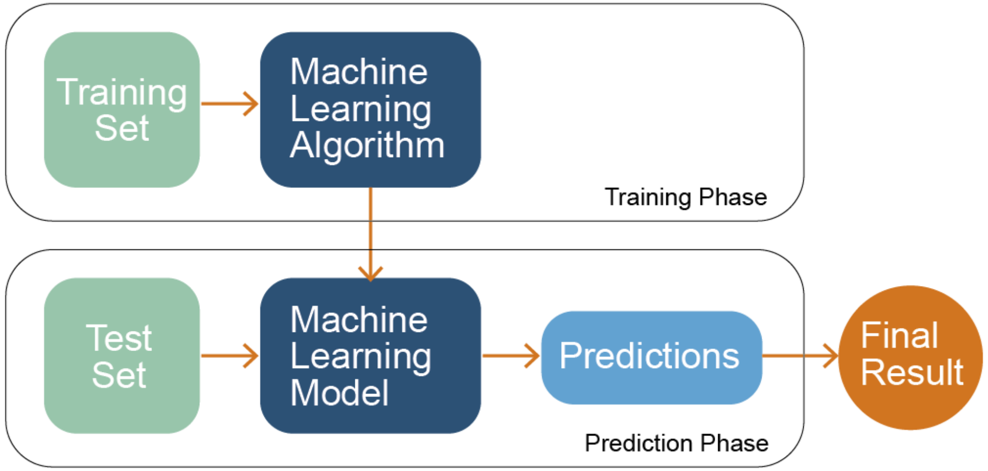

Online Learning

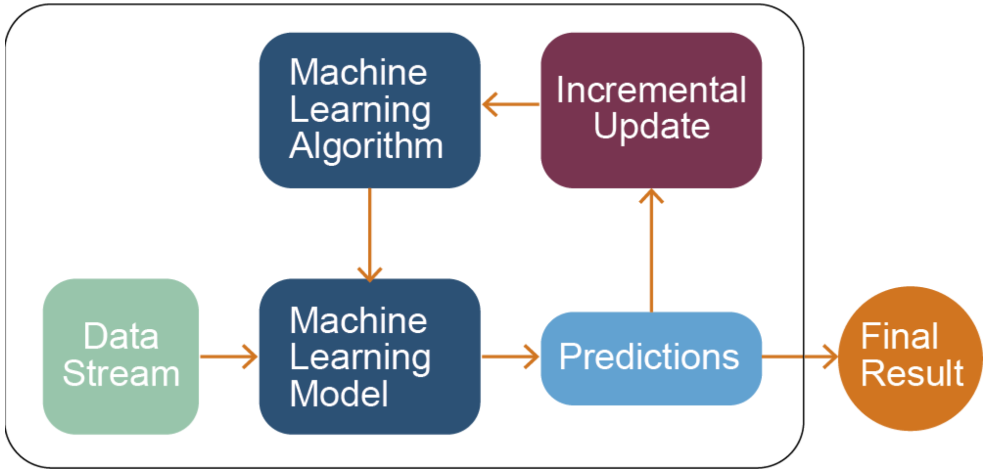

The relevance of image streams

Visual perception systems need to cope with **changing environments**

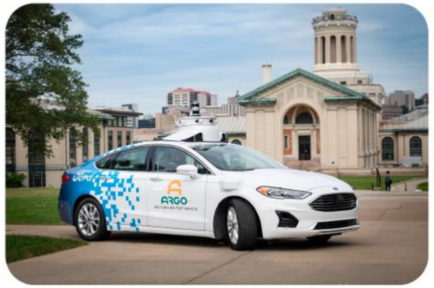

A self-driving car

Pittsburgh

New cities?

Miami

Domino’s car (2013)

New car models?

2013

2025

Domino’s car (2025?)

I.i.d. assumption in data streams

Typical assumption (in data streams):

$(x_t, y_t) \sim P_t$ is iid wrt P_t

i.e., no temporal dependence and same distribution **per concept**.

So, drift is something to detect and 'deal with' , by rebuilding/adapting models.

Learning the concept in its temporal unfolding

Concept drift

Learning the concept in its temporal unfolding

Concept drift + **Temporal dependence**

Problem statement

The field of Streaming Machine Learning faces significant challenges when dealing with **high-dimensional data streams**, particularly in the domain of **image processing**.

- Complex high-dimensional data streams are underinvestigated
- Lack of methods to detect and visualize evolving patterns
- Need to learn the concept in its temporal unfolding
- Importance of interpretable outputs in evolving visual landscapes

Proposed Methodologies

Methodology 1 - Zero Shot Classification for Drift Visualization

Zero-shot classification leverages a **pre-trained ResNet-50** model in **inference-only mode** to detect and visualize **concept drift** in image streams without additional training.

- **No fine-tuning** required (attention to the dataset used in the pretraining!)
- Tracks changes in **predicted class distribution** over time
- Exploits pretrained model's **generalization capability**
- Provides **real-time insights** into evolving data patterns
- Enables the **evaluation on patterns evolution** wrt pre-learned patterns

Proposed Methodologies

Methodology 2 - Feature Extraction and UMAP for Drift Analysis

This approach combines **low-dimensional embeddings** extracted from neural networks with **Uniform Manifold Approximation and Projection (UMAP) dimensionality reduction** to visualize and analyze distributional shifts in image streams.

- Maintains **temporal order** of data streams
- Visualizes **reduced-dimensional embeddings**
- Enables detection of **distributional shifts** in feature space
- Potential for **quantifying variation** using distance metrics
- Suitable for both **real-time monitoring** and **exploratory analysis**

Experimental results – CLEAR10 dataset

- Real-world images with **smooth temporal evolution**
- Captures **technological advancements** and **shifting trends**
- Datasets with **temporal distribution shifts** (real drift)
- Classifier must adapt to the **temporal coherence** of the stream

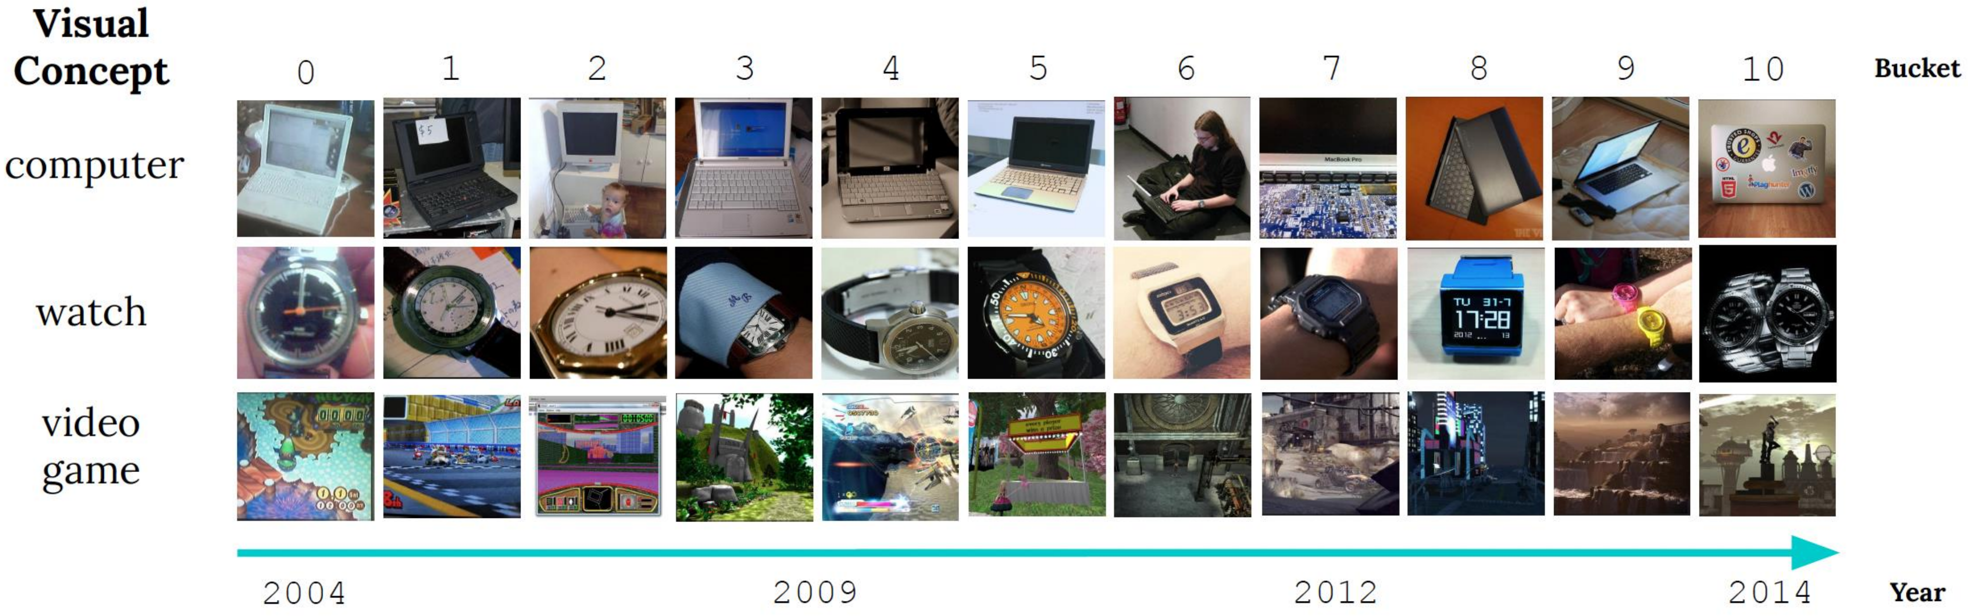

Lin, Z., Shi, J., Pathak, D., & Ramanan, D. (2021). *The CLEAR Benchmark: Continual LEArning on Real-World Imagery*. In NeurIPS Datasets and Benchmarks.

Experimental results – Drift visualization

- Baseball class show **consistent mapping** to ImageNet classes: ~200 images per year are consistently mapped to same ImageNet class
- Indicates **stability** of these classes **over time**

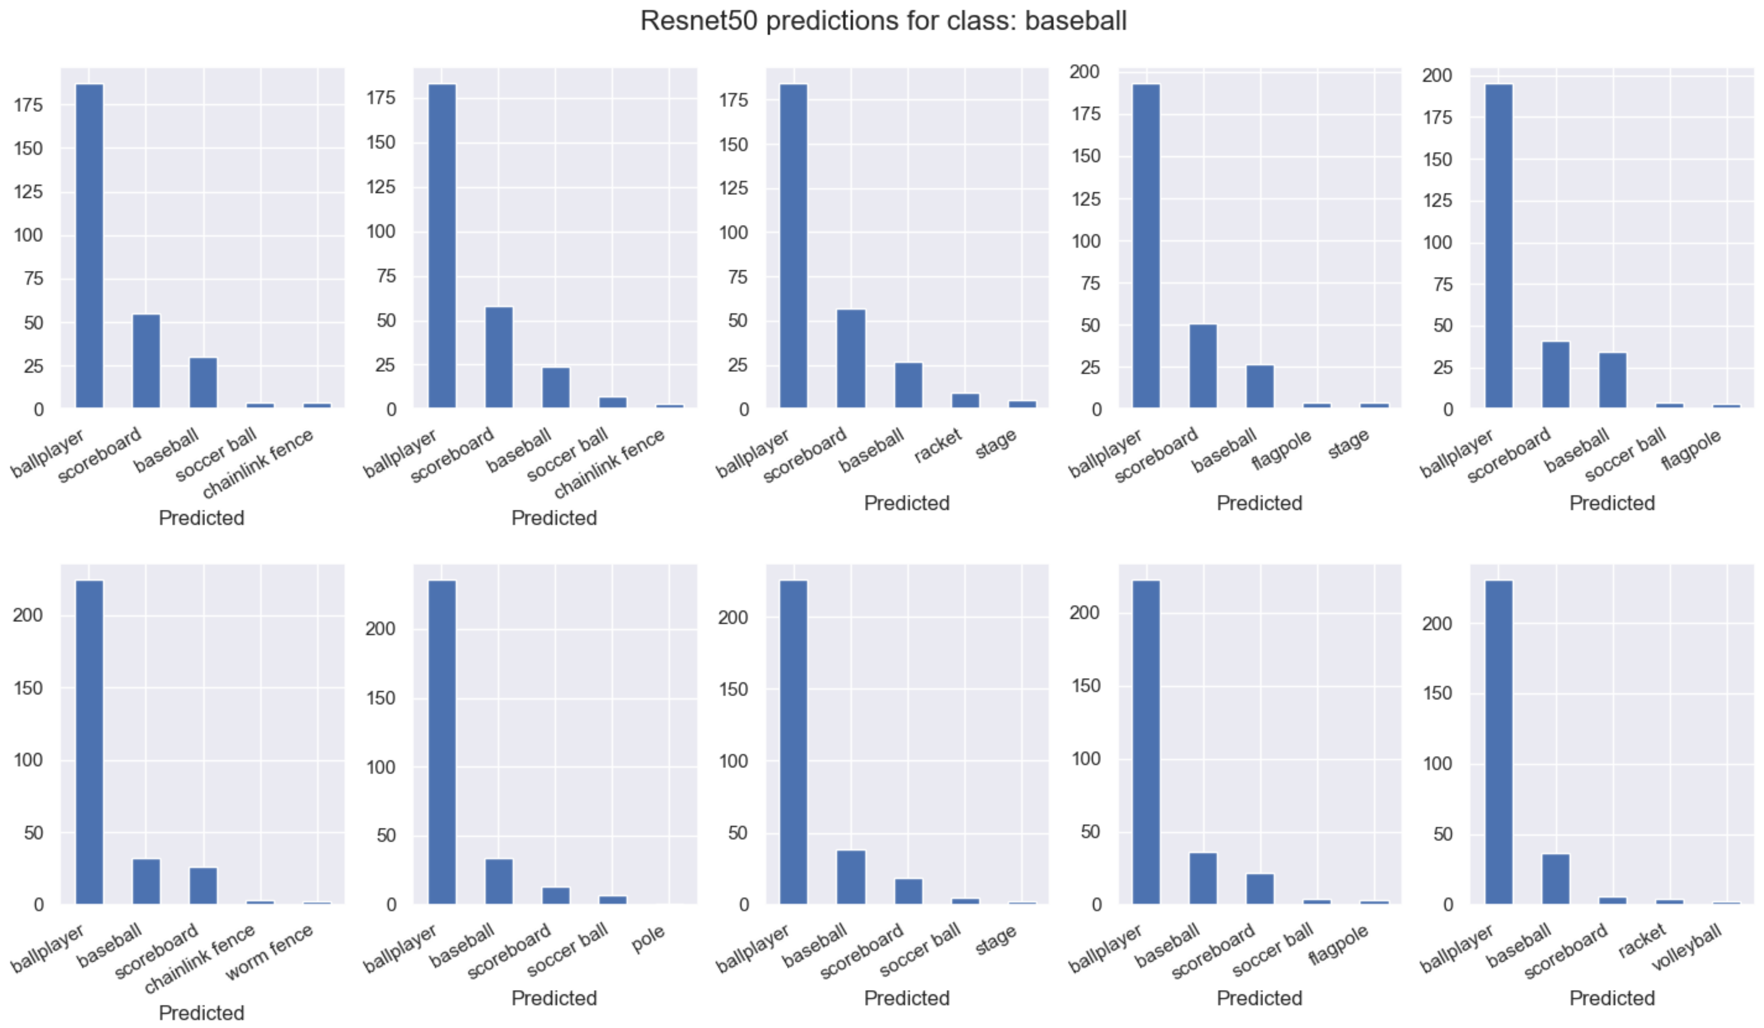

Experimental results – Drift visualization

- Cosplay class absent from ImageNet training set
- Predictions **spread across multiple** ImageNet **classes**, primarily various clothing categories
- Often confused with CLEAR10 dress class (indicates similarity in visual features)
- Fluctuating predictions over time (suggests **potential gradual** and **recurrent drift**)

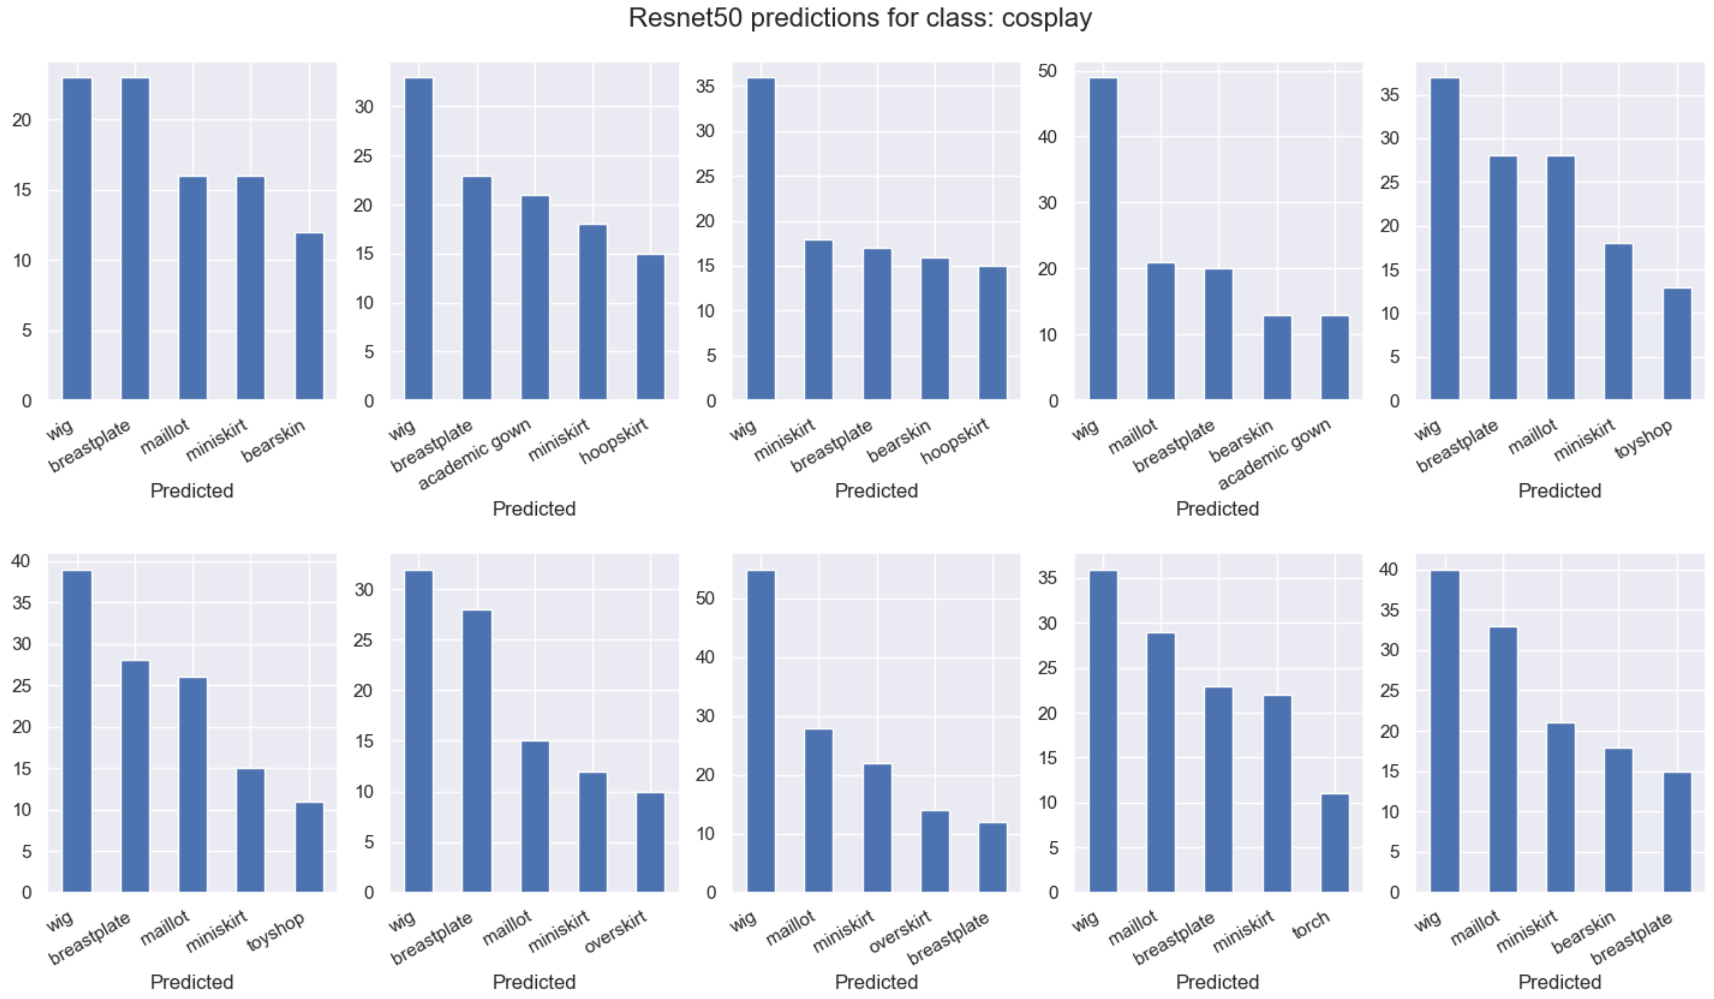

Experimental results – Visualization of all of CLEAR10 classes

- X-axis: CLEAR10 classes by year
- Y-axis: Relevant ImageNet classes
- Circle size/color: Prediction frequency

- Identifies **classes** with **consistent** or **evolving** mappings to ImageNet
- Reveals **stability** or **change** in **class representations over time**

- **Potential for unsupervised scenarios**: Leverages pre-trained model's classes without needing labeled data

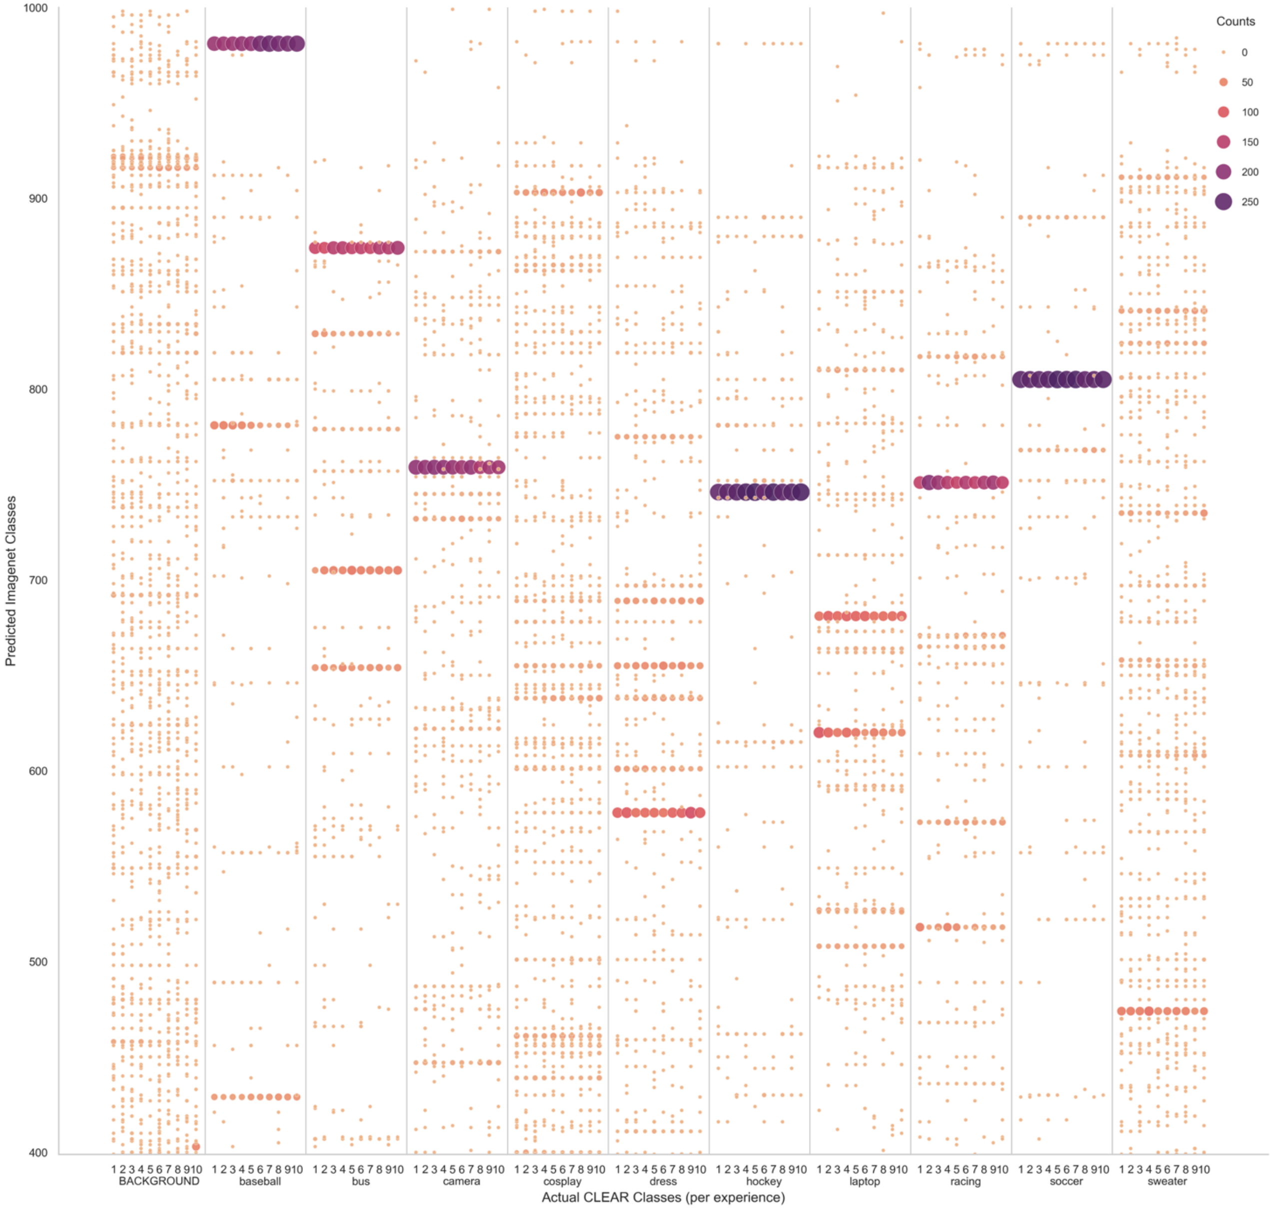

Experimental results - UMAP visualization of Cosplay class

- Cosplay class maintains fairly stable position in the feature space
- Slight shifts suggest the potential presence of gradual drift over time and the absence of abrupt changes

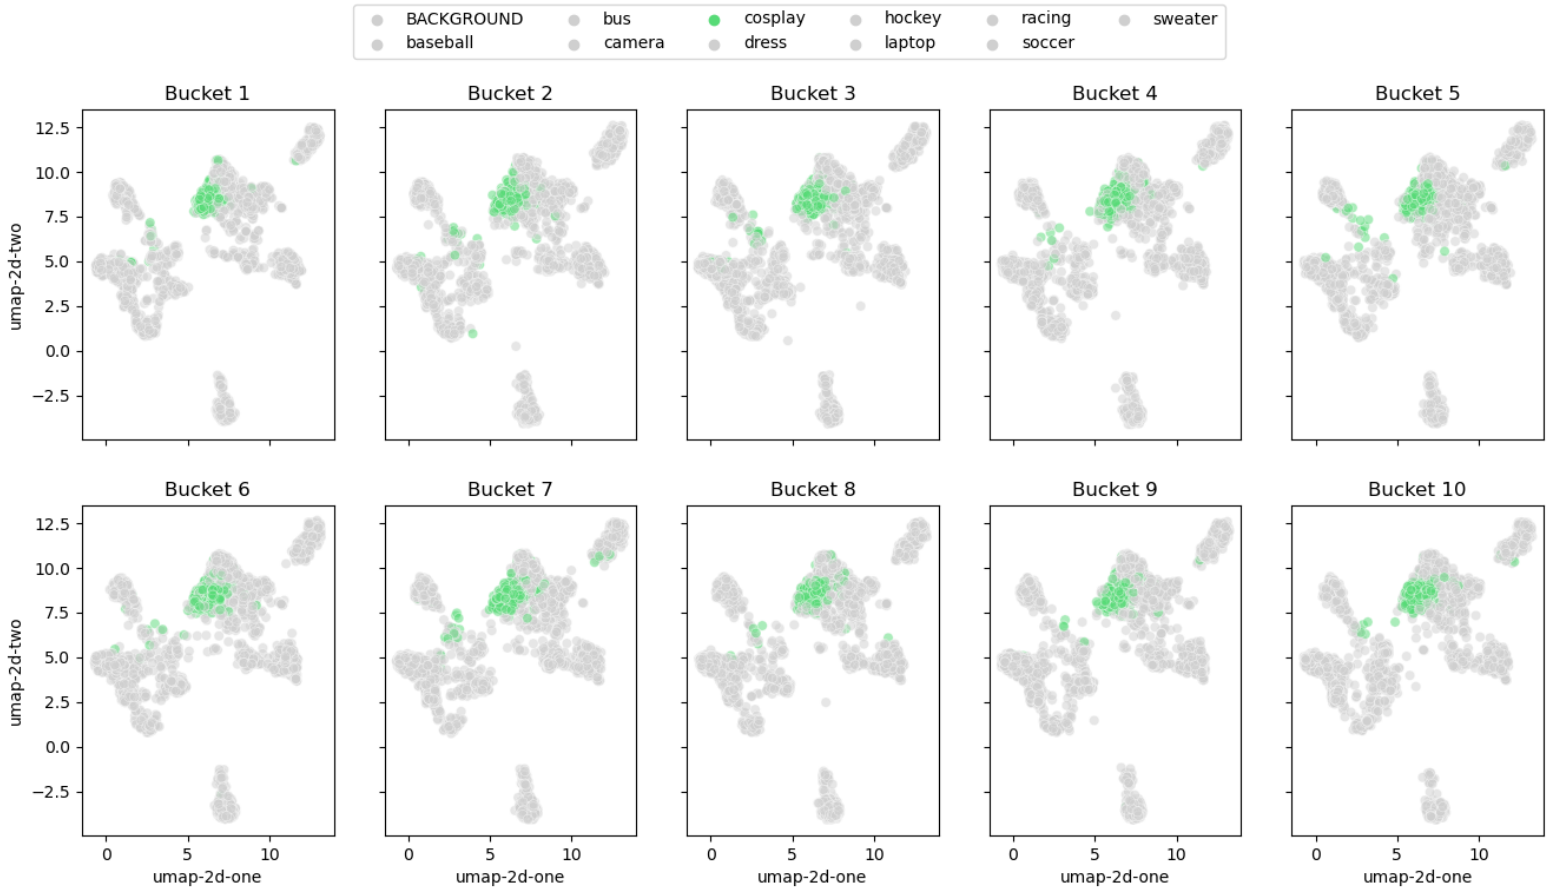

Conclusions and Future works

Conclusions

- Novel methodologies for qualitative drift analysis in image streams:
    - Zero-shot classification with pre-trained DNNs
    - Feature extraction + UMAP visualization
- Preliminary experiments shows potential efficacy in revealing stable or gradually changing concepts

Future Work

- Enlarge experimental campaing using more dataset and pretrained models
- Develop UMAP-based drift detection algorithm
- Integration with adaptive learning systems
- Enhance explainability of detected drifts

POLITECNICO MILANO 1863

Thanks for the attention!

Questions?

Contacts

Giacomo Ziffer

gziffer

giacomo.ziffer@polimi.it