Contextual Similarity Learning for Image Retrieval and Classification: Applications in Person Re-Identification

Lucas Pascotti Valem[1,2*]
Advisor: Daniel Carlos Guimarães Pedronette[1]

[1]Instituto de Geociências e Ciências Exatas – Universidade Estadual Paulista (IGCE-UNESP)
1515, 24-A Ave., Rio Claro - SP, 13506-900, Brazil

[2]Instituto de Ciências Matemáticas e de Computação – Universidade de São Paulo (ICMC-USP)
400, Trabalhador São-Carlense Ave., São Carlos - SP, 13566-590, Brazil

lucas@icmc.usp.br, daniel.pedronette@unesp.br

Abstract. *The rapid expansion of image collections has driven the need for advanced machine learning and image retrieval applications. However, many existing methods depend on large labeled datasets for training, which are costly. To address this, various techniques have been developed, with a key challenge being the effective definition of image similarity. Most approaches still rely on pairwise similarity measures, overlooking valuable contextual information. This work enhances image retrieval and classification by leveraging contextual similarity, proposing seven novel methods applied to general-purpose and person re-identification (Re-ID) datasets. Extensive experiments show that the proposed methods perform comparably to or better than state-of-the-art approaches.*

1. Introduction

In recent years, there has been an exponential increase in the volume of image data, primarily due to advancements in technologies for generating, storing, and sharing visual information [Tripathi and King 2024]. Additionally, there are numerous applications (e.g., surveillance cameras, medical imaging, remote sensing systems, social media) that generate vast amounts of visual data.

In this scenario, image retrieval and machine learning tasks such as image classification are increasingly being utilized in many applications [Peng et al. 2022]. Remarkable progress has been made in these methods, particularly due to the consistent evolution of deep learning. However, the majority of them are supervised and depend on large volumes of labeled data for training. In contrast, the production of labeled data is challenging since it is often expensive and time-consuming to obtain. It may also require a specialist for labeling, especially according to the specificity of the domain. Aiming at filling this gap, many unsupervised, semi-supervised, and even self-supervised approaches have been proposed to deal with such a challenge. In most of these methods, effectively modeling data is crucial for exploiting the information available in the unlabeled data.

For most approaches, the essence of learning hinges on the ability to model data accurately, which involves different concepts, in particular, representation approaches

[*] Lucas Pascotti Valem defended his PhD on June 28, 2024, at São Paulo State University (UNESP) and is currently a Professor at the University of São Paulo (USP).

and distance or similarity measures [Uelwer et al. 2023]. This is especially important for Content-Based Image Retrieval (CBIR) systems, which retrieve images based on visual content rather than metadata. These systems usually employ feature extraction and representation methods, which have evolved considerably. Such methods have transitioned from traditional hand-crafted features to more advanced deep learning approaches, including Vision Transformers. However, most comparison tasks still rely on pairwise measures [Srivastava et al. 2023], which do not exploit contextual information. In general, the term context can be broadly understood as all the relevant information pertinent to an application and its users. This work considers the idea of *contextual similarity* that consists of exploiting the relationships beyond pairwise analysis, involving other elements, such as the neighborhood or more related additional information. The term *contextual similarity learning* is used for the learning process that employs contextual similarity for more effectively capturing the underlying relationships among elements.

This work exploits contextual similarity learning for improving classification and image retrieval, particularly in cases where labeled data is limited or non-existent. Besides general-purpose image retrieval, we also focus on person re-identification, usually abbreviated as person Re-ID. Person Re-ID is a type of surveillance application that has been gaining a lot of attention and nowadays is of fundamental importance in many camera surveillance systems [Liao and Shao 2022]. The task consists of identifying individuals across multiple cameras that have no overlapping views. A Re-ID system broadly consists of three main steps: person detection, feature extraction, and person retrieval or matching. This work focuses on the final step, which can be viewed as a specific image retrieval application.

2. Related Work

The related work presented in the dissertation covers five key topics: image retrieval, unsupervised person re-identification, query performance prediction (QPP), semi-supervised classification, and contrastive learning. From that, it is evident that effectively defining the similarity between images is a central challenge in retrieval and machine learning applications. This issue is deeply connected to *(i)* how information is represented [Uelwer et al. 2023] and *(ii)* the measures used to compare these representations [Srivastava et al. 2023]. Across these areas, different approaches have been proposed to improve similarity estimations, whether by refining feature representations, leveraging contextual relationships, or employing advanced learning paradigms such as contrastive learning and graph-based models.

Despite recent advancements, appropriately representing and exploiting contextual information in each scenario poses significant difficulties. There is still a wide range of methodologies that can be proposed to exploit contextual information in numerous applications. Similarity learning applied to retrieval is generally explored by re-ranking tasks. Despite the increasing popularity of these methods, more robust structures have not yet been extensively employed in most cases. Structures that represent higher-order similarities, such as relationships among neighbors of neighbors, can be particularly advantageous. Hypergraphs [Antelmi et al. 2023], for example, allow edges to connect multiple vertices, offering a sophisticated technique for capturing these relationships.

In addition to retrieval, it is also imperative to address scenarios with limited labeled data for classification [Song et al. 2023]. Graph convolutional networks (GCNs) offer a promising solution for semi-supervised classification by learning from both labeled and unlabeled data, considering graph structures. Moreover, GCNs can learn node and graph embeddings that capture complex dependencies and structural relationships. However, GCNs are not widely used for image classification since graphs are typically not available in image domains [Rodrigues and Carbonera 2024]. Therefore, effectively modeling these graphs, which can be utilized to exploit contextual information, is a crucial topic for research.

Another approach that has recently demonstrated continuous advances for improving classification results is the use of contrastive learning [Tripathi and King 2024]. Unlike the commonly used cross-entropy loss, which aims to minimize the difference between the predicted and true class probabilities, contrastive loss focuses on learning similarities and dissimilarities between data points rather than merely categorizing them. Despite this, most contrastive losses consider only pairwise measures, with only a few incorporating some type of neighborhood information. Moreover, these approaches often require huge volumes of data for training (i.e., labeled or unlabeled), even in self-supervised scenarios, which is a challenge in circumstances where data is scarce.

3. Objective, Research Questions, and Dissertation Statement

In light of the presented discussion and all challenges, the objective of this work is to exploit the use of contextual similarity information to improve the effectiveness of image retrieval and classification, particularly in cases where labeled data is limited or nonexistent. This study primarily concentrates on unsupervised learning, while also proposing semi-supervised and supervised approaches.

However, effectively representing and leveraging contextual information in each scenario presents substantial challenges. Addressing these challenges requires careful consideration of how contextual similarity can be integrated to enhance retrieval and classification tasks. In this context, the research questions formulated to guide this work are:

- **Q1:** How can contextual similarity information be used for selection and fusion in unsupervised person Re-ID?
- **Q2:** How can data be modeled using contextual similarity information for QPP?
- **Q3:** How can contextual similarity information be used to generate synthetic data?
- **Q4:** How can contextual similarity learning be employed on synthetic data?
- **Q5:** How can more complex structures, which encode contextual information more effectively, be applied to unsupervised similarity learning?
- **Q6:** How can contextual similarity information be encoded to generate embeddings that are useful for tasks beyond retrieval, such as classification?
- **Q7:** How can contextual similarity be incorporated into the input graph utilized by Graph Convolutional Networks (GCNs) and improve classification results?
- **Q8:** Can rank-based information be utilized to measure the correlation between images more effectively?
- **Q9:** Can a correlation measure be proposed and applied to enhance image retrieval with manifold learning?
- **Q10:** How can contextual similarity information be incorporated into metric learning, including its direct integration into losses such as contrastive loss?

Given the challenges of obtaining large amounts of labeled data and the growing need for methods that leverage contextual information, this work explores the application of contextual similarity learning in different scenarios. The main hypothesis of the work is briefly stated as follows:

> *Contextual similarity learning can improve the effectiveness of image retrieval and classification tasks across general-purpose and person re-identification (Re-ID) applications. This concept is applicable to unsupervised, semi-supervised, and supervised approaches, particularly in contexts where labeled data is limited.*

The hypothesis is validated by the proposed approaches and a comprehensive experimental evaluation presented in the dissertation.

4. Main Contributions

Effectively defining the similarity between images is a key challenge in machine learning and retrieval applications. This issue is closely tied to how information is represented and the measures used to compare these representations. This work contributes to both aspects by introducing 7 novel approaches. **During the doctoral period, a total of 27 publications were produced in renowned national and international conferences and journals, with impact factors as high as 23.8, h5-index values up to 165, and 242 citations according to Google Scholar, involving a total of 33 co-authors.** Among them, 7 are first-authored publications, each corresponding to one of the proposed methods described in the dissertation, while the remaining 20 resulted from collaborations. **Further details and a complete list of publications can be found in the by-products document.**

Figure 1 outlines the goals, contributions, and their relationships with each approach. There are different addressed tasks: *(i)* content-based image retrieval, which involves obtaining a ranked list of images based on their visual similarity to a query image; *(ii)* query performance prediction (QPP), which aims at evaluating the effectiveness of ranked lists in scenarios where no labels are available; and *(iii)* image classification, which assigns a label to a given image.

Notice that the proposed approaches were investigated to address the research challenges previously discussed. Each method exploits contextual similarity differently. In the following, an overview of each of the seven proposed methods is presented:

1. **Deep Rank Noise Estimator (DRNE) [Valem and Pedronette 2021]:** Proposes a new model architecture for image denoising to perform query performance prediction. The idea is to interpret the incorrectness of a ranked list as noise in a *contextual image*.
2. **Regression for Query Performance Prediction Framework (RQPPF) [Valem et al. 2023c]:** Supports diverse features and regression models for query performance prediction. It computes *meta-features*, that encode reciprocal neighborhood information, based on unsupervised measures.
3. **JaccardMax [Valem et al. 2022]:** Effectively measuring similarity among data samples represented as points in high-dimensional spaces is a major challenge. A novel rank correlation measure, robust to such variations, is a contribution of this study. The proposed measure is suitable for diverse scenarios and is validated on an unsupervised manifold learning algorithm based on the Correlation Graph (CG) approach.

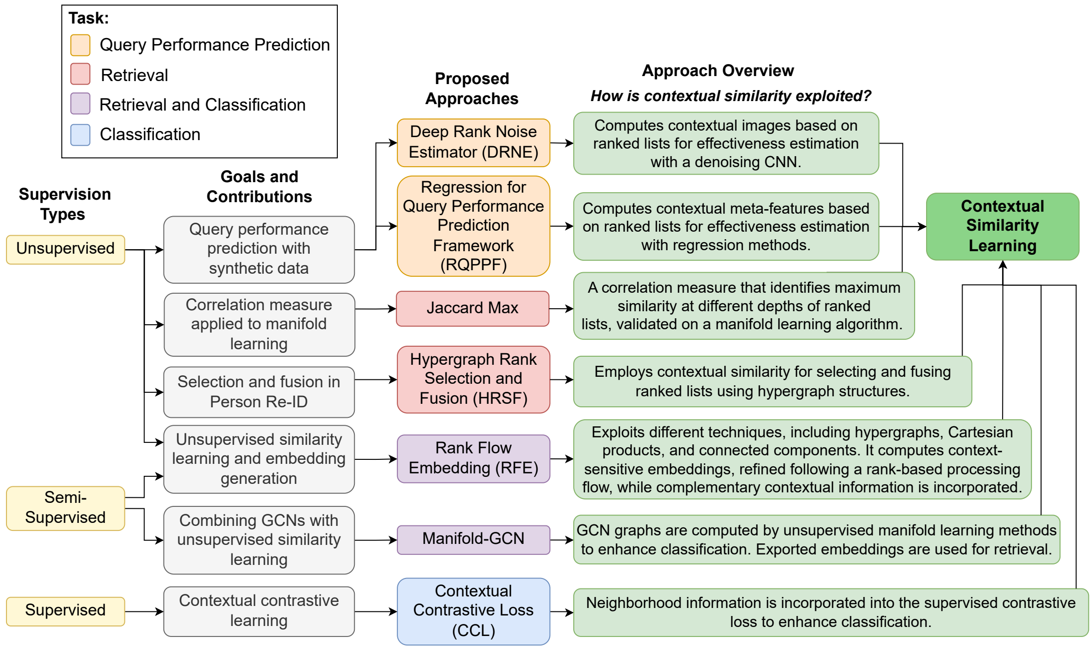

Figure 1. Overview of the contributions and how contextual similarity is exploited for each proposed approach.

4. **Hypergraph Rank Selection and Fusion (HRSF) [Valem and Pedronette 2022]:** This contribution addresses the challenging task of unsupervised selection and fusion of different features for more effective person re-identification. A novel framework is proposed, which combines an unsupervised rank-based formulation for feature selection with a robust hypergraph model for query performance prediction and rank aggregation based on manifold learning.
5. **Rank Flow Embedding (RFE) [Valem et al. 2023b]:** A novel manifold learning algorithm for unsupervised and semi-supervised scenarios. The proposed method is based on ideas recently exploited by manifold learning approaches, which include hypergraphs, Cartesian products, and connected components. The algorithm computes context-sensitive embeddings, which are refined following a rank-based processing flow, while complementary contextual information is incorporated. The generated embeddings can be exploited for more effective unsupervised retrieval or semi-supervised classification based on Graph Convolutional Networks.
6. **Manifold-GCN [Valem et al. 2023a]:** Based on GCNs for semi-supervised image classification. The main objective is to use manifold learning to model the graph structure to further improve the GCN classification. All manifold learning algorithms employed are completely unsupervised, which is especially useful for scenarios where the availability of labeled data is a concern. This method is also evaluated for person Re-ID by utilizing the embeddings exported by the GCNs.
7. **Contextual Contrastive Loss (CCL) [Valem et al. 2024]:** Besides the promising results obtained by contrastive learning approaches, they often consider pairwise measures. The proposed Contextual Contrastive Loss (CCL) replaces pairwise image comparison by introducing a new contextual similarity measure using neighboring elements. The CCL yields a more semantically meaningful image embedding, ensur-

ing better separability of classes in the latent space. Although supervised, the results show that the CCL provides higher gains in cases with fewer labeled data.

5. Experimental Results

This section summarizes the main results from a wide range of experiments presented in the dissertation. Table 1 provides an overview of the evaluation by indicating, with bullets, which methods were applied to each dataset. These methods are grouped into three categories, with colors distinguishing the different types of supervision. Notice that RFE and Manifold-GCN appear in multiple categories, as they were applied to both classification and retrieval. The following subsections detail the results for each task: query performance prediction (QPP), retrieval, and classification. Due to space restrictions, the references for the datasets and baselines were omitted, but are available in the dissertation.

Table 1. Image datasets used in the evaluation of each of the proposed methods, categorized by task and type of supervision.

Supervision Type				Unsupervised					Semi-supervised			Supervised
Task Type				QPP		Retrieval				Classification		
Category	Dataset	Size	Classes	DRNE	RQPPF	HRSF	JaccardMax	RFE	Manifold-GCN	RFE	Manifold-GCN	CCL
General-Purpose	ORL Faces	400	40					•				
	Flowers	1,360	17	•				•		•	•	
	MPEG-7	1,400	70	•	•			•				
	Holidays	1,491	500				•	•				
	Brodatz	1,776	16	•	•			•				
	Corel5k	5,000	50				•	•		•	•	
	UKBench	10,200	2,550				•	•				
	CUB200	11,788	200								•	
	Dogs	20,580	120				•					
	CIFAR-100	60,000	100									•
	MiniImageNet	60,000	100									•
	ALOI	72,000	1,000					•				
	Food101	101,000	101									•
Person Re-ID	CUHK03	14,097	1,467			•		•	•			
	Market1501	32,217	1,501	•	•	•	•	•	•			
	DukeMTMC	36,411	1,812	•	•	•	•	•	•			
	Airport	39,902	9,651			•						

5.1. Query Performance Prediction (QPP)

Experiments were conducted to evaluate the approaches proposed for QPP, DRNE, and RQPPF, showing their capacity to effectively perform QPP in various datasets. Table 2 presents a summary of the relative gains for both approaches in comparison to Authority and Reciprocal Density, which are used as baselines. Notice that RQPPF provided gains in all the evaluated scenarios, while DRNE showed inferior performance in some cases, especially when compared to Reciprocal in the MPEG-7 dataset. However, for the most part, DRNE provided higher and more consistent gains when compared to Reciprocal. Specifically for the AIR descriptor, DRNE revealed superior results in all cases.

5.2. Image Retrieval

A comparison on a general-purpose Holidays dataset is presented in Table 3, where the proposed RFE and Correlation Graph (CG) + JaccardMax methods are evaluated against baseline approaches. Notably, RFE outperformed all of them, while CG + JaccardMax achieved 91.12%, surpassing most other methods.

Table 2. Relative gains of DRNE and RQPPF compared to Authority and Reciprocal Density.

Descriptor	Original MAP	Compared to Auth. DRNE	Compared to Auth. RQPPF	Compared to Rec. DRNE	Compared to Rec. RQPPF
MPEG-7					
AIR	89.39%	+14.81%	+3.50%	+16.84%	+12.99%
ASC	85.28%	-2.50%	+3.76%	-8.29%	+1.84%
IDSC	81.70%	-3.93%	+3.69%	-7.58%	+1.88%
CFD	80.71%	+3.47%	+3.99%	-1.24%	+2.71%
BAS	71.52%	+0.85%	+3.69%	-5.13%	+1.09%
SS	37.67%	+7.08%	+6.01%	+3.32%	+3.52%
Average Gain		+3.30%	+4.11%	-0.35%	+4.01%
Market					
OSNET	43.30%	-2.63%	+1.19%	+5.40%	+5.45%
ResNet	22.82%	-0.89%	+0.13%	+7.95%	+5.29%
Average Gain		-1.76%	+0.66%	+6.68%	+5.37%
DukeMTMC					
OSNET	52.69%	+0.71%	+2.35%	+1.00%	+3.37%
ResNet	32.00%	-2.14%	+0.52%	-0.12%	+2.46%
Average Gain		-0.72%	+1.44%	+0.44%	+2.92%

Table 3. Comparison of approaches on the Holidays dataset (MAP).

Method	MAP
Paulin *et al.* (IJVC 2017)	82.90%
Sun *et al.* (Info. Sciences 2017)	85.50%
Pedronette *et al.* (PR 2017)	86.16%
Arandjelovic *et al.* (CVPR 2016)	87.50%
Li *et al.* (CVPR 2015)	89.20%
Pedronette *et al.* (PR 2021)	90.02%
Gordo *et al.* (IJCV 2017)	90.30%
Valem *et al.* (PRL 2020)	90.51%
Valem *et al.* (Neurocomputing 2020)	90.51%
Liu *et al.* (TIP 2017)	90.89%
Pedronette *et al.* (TIP 2019)	90.94%
Pedronette *et al.* (Journal of Imaging 2021)	91.25%
Yu *et al.* (Neurocomputing 2017)	91.40%
Berman *et al.* (arXiv 2019)	91.80%
CG + JacMax (Ours)	**91.12%**
RFE (Ours)	**91.97%**

For person Re-ID, Figure 2 presents a plot where each dot represents an unsupervised method and its MAP over the years for the DukeMTMC dataset, the largest Re-ID dataset used in the evaluation. Notably, the proposed approaches are among the best.

From the observed results, we can notice that the proposed approaches are comparable or better than state-of-the-art approaches in most cases. The best method in each scenario varies, since each dataset and descriptor presents different aspects. An important attribute of the proposed approaches is their capacity to improve the input data by employing contextual similarity learning. Figure 3 summarizes the relative gains of RFE and JaccardMax for different datasets and descriptors. This highlights the capacity of contextual similarity learning to improve the results across multiple scenarios.

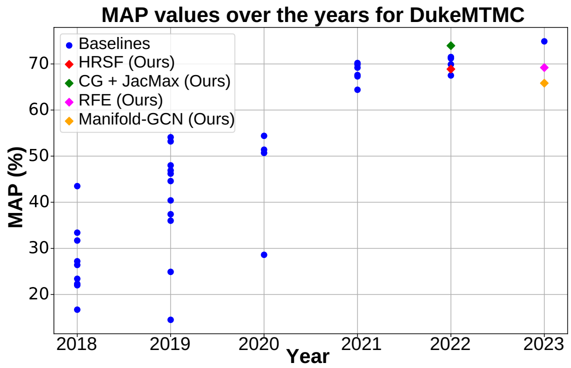

Figure 2. Evolution of MAP for DukeMTMC dataset over the years for unsupervised approaches.

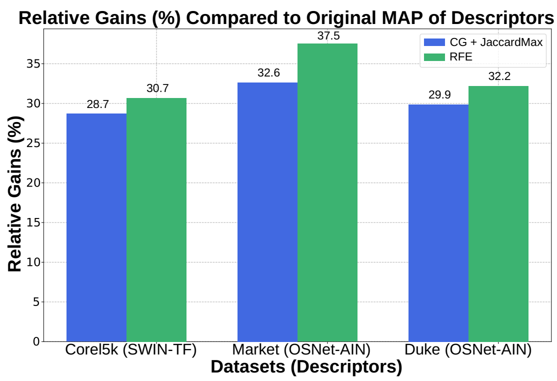

Figure 3. RFE and JaccardMax relative gains (%) over MAP of descriptors.

Experiments were also conducted to visualize the performance of RFE in retrieval tasks. Figure 4 presents an example of a ranked list before and after the execution of RFE for the OSNET-AIN model and the DukeMTMC dataset. The query images are highlighted with green borders, while the incorrect ones are marked with red borders.

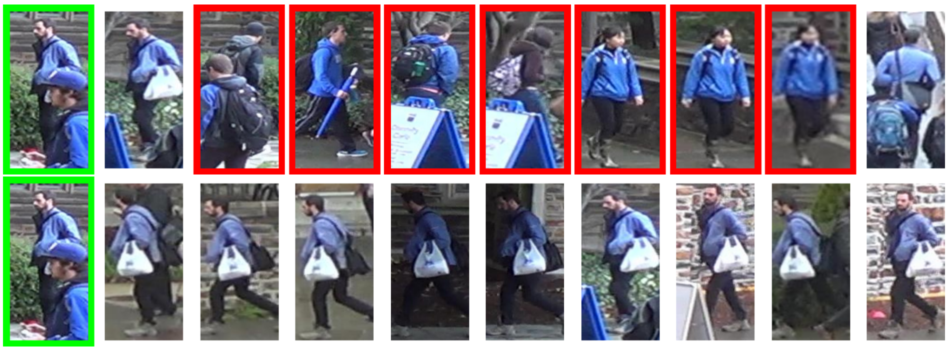

Figure 4. Ranked list before and after RFE for DukeMTMC dataset. Query images are highlighted with green borders, and wrong results with red borders.

5.3. Image Classification

Since both Manifold-GCN and RFE were employed for semi-supervised classification, Table 4 compares them to baselines, both traditional and recent, on Flowers and Corel5k datasets. The values achieved by Manifold-GCN are the highest in all the cases, and they are closely followed by RFE. These results reveal the high effectiveness of the proposed approaches that, besides the significant gains, are also comparable or superior to various methods in the literature.

The CCL was also proposed and evaluated for classification, but in supervised scenarios. Since it is the only method proposed in this category and uses different datasets in the protocol, a direct comparison of it with other proposed methods is not feasible. Therefore, a discussion about its gains is presented. The experimental evaluation in the dissertation showed that the results are consistently better than those of SupCon, which CCL is based on, and SimCLR, another method commonly used as a baseline in this task. Figure 5 presents a plot that evinces the capacity of CCL to provide gains when compared to SupCon for three datasets and with higher values as the training set size decreases. The integration of contextual information within the contrastive loss significantly improved the results, as initially hypothesized, with gains up to 10.759%.

6. Conclusions and Future Work

The dissertation explored contextual similarity learning techniques across various tasks, including query performance prediction, image retrieval, and classification. The proposed approaches — DRNE, RQPPF, HRSF, JaccardMax, RFE, Manifold-GCN, and CCL — demonstrated significant improvements over traditional methods. The findings highlight that incorporating contextual information enhances similarity measures, improves ranked list effectiveness, and enables the learning of robust feature embeddings. Furthermore, these methods achieved competitive results in both unsupervised and semi-supervised settings, often outperforming state-of-the-art baselines. This research also introduced novel techniques for metric learning, manifold analysis, and graph-based representations. A summary of the connections between the research questions and key contributions is provided in Table 5. **Further details and a complete list of publications can be found in the by-products document, including a diagram with all the works that originated from collaborations and their respective publications.**

There are several promising directions for future work, including extending these approaches to multi-query re-identification, exploring applications in other types of multimedia data, and refining strategies for synthetic data generation and embedding.

Table 4. Accuracy (%) comparison for semi-supervised classification.

Method	Flowers	Corel5k
CoMatch	82.55	*85.70*
kNN	63.67	76.80
SVM	80.54	88.73
OPF	71.77	83.56
SL-Perceptron	75.44	83.56
ML-Perceptron	78.88	87.10
PseudoLabel+SGD	82.69	89.76
LS+kNN	73.49	83.98
LS+SVM	73.53	83.26
LS+OPF	72.66	82.32
LS+SL-Perceptron	72.34	82.38
LS+ML-Perceptron	73.03	82.53
GNN-LDS	54.98	62.69
GNN-KNN-LDS	79.32	88.94
WSEF	85.12	91.68
RFE (ours)	84.95	91.54
Manifold-GCN (ours)	**85.88**	**93.08**

Figure 5. Relative gains (%) obtained by CCL in comparison to SupCon for different train/test splits.

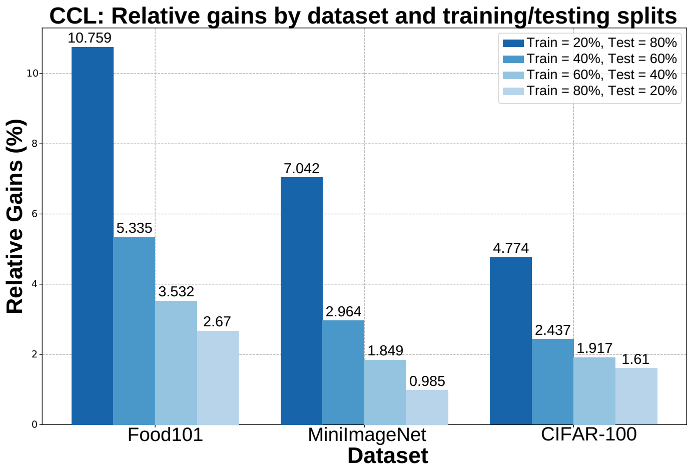

Table 5. Research questions addressed by each of the proposed approaches.

Summarized Research Question	DRNE	RQPPF	JaccardMax	HRSF	RFE	Manifold-GCN	CCL
	Proposed Methods						
Q1. Unsupervised selection and fusion in Re-ID?				•			
Q2. Model data with contextual information for QPP?	•	•					
Q3. Create synthetic data with contextual similarity?	•	•					
Q4. Contextual similarity learning using synthetic data?	•	•					
Q5. More complex structures in unsupervised similarity learning?				•	•		
Q6. How can embeddings be generated for different tasks?					•	•	
Q7. Improve GCN input graph for better classification?						•	
Q8. More effective rank-based correlation measure?			•				
Q9. Correlation measure with manifold learning?			•				
Q10. Similarity information into contrastive loss?							•

Acknowledgments

The authors acknowledge the **financial support from Fulbright, which funded a nine-month research internship at Temple University (Philadelphia, PA, USA)**, and thank Prof. Longin Jan Latecki for his guidance during the internship. We also thank FAPESP – São Paulo Research Foundation (grants #17/25908-6, #18/15597-6, #20/11366-0) and Petrobras for sponsoring projects in which the author participated during his PhD, thereby indirectly contributing to the work presented in the dissertation. We would also like to thank the University of São Paulo (PRPI Ordinance No. 1032) and CNPq – Brazilian National Council for Scientific and Technological Development (grant #313193/2023-1).

References

Antelmi, A., Cordasco, G., Polato, M., Scarano, V., Spagnuolo, C., and Yang, D. (2023). A survey on hypergraph representation learning. *ACM Computing Surveys*, 56(1).

Liao, S. and Shao, L. (2022). Graph sampling based deep metric learning for generalizable person re-identification. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7349–7358.

Peng, L., Qiang, B., and Wu, J. (2022). A survey: Image classification models based on convolutional neural networks. In *2022 14th International Conference on Computer Research and Development (ICCRD)*, pages 291–298.

Rodrigues, J. and Carbonera, J. (2024). Graph convolutional networks for image classification: Comparing approaches for building graphs from images. In *26th Int. Conference on Enterprise Information Systems - Vol. 1: ICEIS*, pages 437–446. SciTePress.

Song, Z., Yang, X., Xu, Z., and King, I. (2023). Graph-based semi-supervised learning: A comprehensive review. *IEEE Transactions on Neural Networks and Learning Systems*, 34(11):8174–8194.

Srivastava, D., Singh, S. S., Rajitha, B., Verma, M., Kaur, M., and Lee, H.-N. (2023). Content-based image retrieval: A survey on local and global features selection, extraction, representation, and evaluation parameters. *IEEE Access*, 11:95410–95431.

Tripathi, S. and King, C. R. (2024). Contrastive learning: Big data foundations and applications. In *Proceedings of the CODS-COMAD '24*, page 493–497, New York, NY, USA. Association for Computing Machinery.

Uelwer, T., Robine, J., Wagner, S. S., Höftmann, M., Upschulte, E., Konietzny, S., Behrendt, M., and Harmeling, S. (2023). A survey on self-supervised representation learning. *arXiv preprint arXiv:2308.11455*.

Valem, L. P., Kawai, V. A. S., Pereira-Ferrero, V. H., and Pedronette, D. C. G. (2022). A novel rank correlation measure for manifold learning on image retrieval and person re-id. In *IEEE International Conference on Image Processing (ICIP)*, pages 1371–1375.

Valem, L. P., Pedronette, D., and Allili, M. S. (2024). Contrastive loss based on contextual similarity for image classification. In *19th International Symposium on Visual Computing (ISVC)*.

Valem, L. P. and Pedronette, D. C. G. (2021). A denoising convolutional neural network for self-supervised rank effectiveness estimation on image retrieval. In *Proc. of the International Conference on Multimedia Retrieval*, ICMR '21, page 294–302. ACM.

Valem, L. P. and Pedronette, D. C. G. (2022). Person re-id through unsupervised hypergraph rank selection and fusion. *Image Vision Computing*, 123(C).

Valem, L. P., Pedronette, D. C. G., and Latecki, L. J. (2023a). Graph convolutional networks based on manifold learning for semi-supervised image classification. *Computer Vision and Image Understanding*, 227:103618.

Valem, L. P., Pedronette, D. C. G., and Latecki, L. J. (2023b). Rank flow embedding for unsupervised and semi-supervised manifold learning. *IEEE Transactions on Image Processing*, 32:2811–2826.

Valem, L. P., Pereira-Ferrero, V., and Pedronette, D. (2023c). Self-supervised regression for query performance prediction on image retrieval. In *2023 IEEE Sixth International Conference on Artificial Intelligence and Knowledge Engineering (AIKE)*, pages 95–98, Los Alamitos, CA, USA. IEEE Computer Society.