Nonmetric Earth Mover's Distance for Efficient Similarity Search

Jakub Lokoč and Tomáš Skopal
SIRET Research Group, Faculty of Mathematics and Physics
Charles University in Prague, Czech Republic
{lokoc,skopal}@ksi.mff.cuni.cz

Christian Beecks and Thomas Seidl
Data Management and Data Exploration Group
RWTH Aachen University, Germany
{beecks,seidl}@cs.rwth-aachen.de

***Abstract*—The Earth Mover's Distance is a well-known distance-based similarity measure employed in various domains of data management, especially in computer vision and content-based multimedia retrieval. However, as the computation of the Earth Mover's Distance is a considerably expensive task, efficient processing of content-based similarity queries in large multimedia databases remains a challenging issue. In this paper, we propose to use nonmetric ground distances within the computation of the Earth Mover's Distance in order to speedup its computation, thus improving the efficiency of the entire retrieval process. Moreover, by investigating the inner workings of the Earth Mover's Distance, we show how to balance the trade-off between effectiveness and efficiency in order to adapt the retrieval process to individual user requirements. By making use of metric access methods in combination with the Rubner filter, we empirically show an improvement in efficiency by two orders of magnitude according to the sequential scan, while keeping the retrieval error below 5%.**

***Index Terms*—Earth Mover's Distance; Similarity Search; Indexing.**

I. Introduction

Distance-based similarity search has been successfully utilized in various domains including computer vision and content-based multimedia retrieval, where a database consists of objects represented by nontrivial feature descriptors extracted from unstructured complex data (such as multimedia data). To further improve applicability and effectiveness of such distance-based similarity models, domain experts shift from traditional feature histograms to feature signatures [4], [15], which can flexibly describe the content of multimedia objects. Feature signatures in combination with adaptive distance-based similarity measures [4] then form a powerful tool for effective content-based multimedia retrieval.

However, as the volume of multimedia data grows exponentially, content-based retrieval systems have to provide users with new and more sophisticated exploration facilities. To this end, distance-based similarity models are expected to provide additional trade-off parameters for tuning the precision/performance of the retrieval systems. The system administrators then use the parameters to adapt the retrieval model to better fit user requirements, e.g., more effective but less efficient retrieval, or vice versa. The parameters often affect both quality of the similarity measure and its behavior within an indexing structure for fast retrieval.

In this paper, we focus on the *Earth Mover's Distance* (EMD) [15] – an adaptive distance function for measuring similarity that utilizes a user-defined *ground distance* to penalize some operations within the similarity assessment. In order to speedup the similarity search process using the EMD, we could benefit from indexing by *metric access methods* [6], [20] or by processing queries in a *filter-and-refinement* scheme [2], [3], [16]. Most methods assume *metric* ground distances and some are limited just to feature *histograms*. However, considering also *nonmetric* ground distances could lead to more robust behavior of the distance measure. The arguments for nonmetric ground distances follow from previous image retrieval studies, showing that nonmetric distances performed better than the metric ones [8], [10].

Moreover, using the *parameterized version* of the Earth Mover's Distance [12], we can investigate and tune properties of the distance space, like the "indexability" (i.e., how fast are we able to search/prune that space), measured via the intrinsic dimensionality [6], or the "metricity" that affects the retrieval error exhibited by the metric access methods [18]. In general, for distance spaces suffering from high intrinsic dimensionality, as the traditional (non-parameterized) Earth Mover's Distance [15] on feature signatures [4], [15] does, it is impossible to create an efficient metric index for exact search, and for this reason it is more promising to employ domain-specific filters. Hence, in this paper we investigate the impact of the parameterized Earth Mover's Distance on the distance space properties and point to the promising trade-offs between indexability and retrieval quality. We also combine the distance-based approach with the domain-specific filters.

A. Paper Contribution

The main contribution of this paper is an evaluation of the similarity search under the parameterized Earth Mover's Distance utilizing various L_p ground distances (metric and nonmetric). The findings can be summarized as:

- The parameterized Earth Mover's Distance can be used to tune the retrieval quality and can be efficiently processed by metric access methods or domain-specific filtering techniques. Moreover, we can use the parameter of the distance to guarantee exact searching via more effective fractional L_p distances.

 ISBN: 978-1-61208-195-3

- Although the nonmetric ground distances turn both the filtering and indexing techniques into just approximate methods, the retrieval error caused by the ”ground nonmetricity” is not significant.
- The parameterized Earth Mover’s Distance can be tuned to better fit metric access methods, however, the retrieval error obtained by metric access methods is higher than the retrieval error obtained by the filtering approach.

In the next two sections, we review preliminaries and related work that we utilize and combine in this paper. In section 4, we revisit the filter-and-refinement schema used for efficient EMD processing and we discuss the impact of the nonmetric ground distances. After that, we report and discuss experimental results in section 5 and finally, we conclude contributions of our approach in section 6.

II. PRELIMINARIES

Before we proceed to the contribution and the related work, we briefly summarize the motivation for feature signatures and the Earth Mover’s Distance.

A. Feature Signatures

Unlike conventional feature histograms, feature signatures [4], [15] are frequently obtained by clustering the objects’ properties [7], [14] within a feature space and storing the cluster representatives and weights. Thus, given a feature space $\mathbb{F}$, the *feature signature* S^o of a multimedia object o is defined as a set of tuples from $\mathbb{F} \times \mathbb{R}^+$ consisting of representatives $r^o \in \mathbb{F}$ and weights $w^o \in \mathbb{R}^+$.

We depict an example of image feature signatures according to a feature space comprising position and color information, i.e., $\mathbb{F} \subseteq \mathbb{R}^5$, in Figure 1. The feature representatives are depicted as circles of the corresponding color, while the weights are reflected by the diameter of the circles. As can be seen in this example, feature signatures adjust to individual image contents by aggregating the features according to their appearance in the underlying feature space.

Although feature signatures are more general than feature histograms, in fact feature histograms are a special case of feature signatures, for similarity search purposes they could be used together with some distances originally designed for histograms, such as the Earth Mover’s Distance or the Quadratic Form Distance (generalized to Signature Quadratic Form Distance [5]). In this paper, we put our attention to the Earth Mover’s Distance, which we will explain in the next section.

B. Earth Mover’s Distance

The Earth Mover’s Distance is a distance-based similarity measure originated in the computer vision domain [15]. Its successful utilization, however, gave raise to numerous applications also in different domains. This distance describes the cost for transforming one feature signature (or histogram) into another one. The distance is considered to be a transportation problem and thus is the solution to a linear optimization problem which can be solved via a specialized simplex algorithm.

Fig. 1. Three example images with their corresponding feature signature visualizations.

Given a *ground distance* d that measures the dissimilarity of two features within a feature space $\mathbb{F}$, the *Earth Mover’s Distance* (EMD) is defined between two feature signatures $S^q = \{c_i^q, w_i^q\}_{i=1}^n$ and $S^o = \{c_j^o, w_j^o\}_{j=1}^m$ as a minimum cost flow over all flows $f_{ij} \in \mathcal{R}$ as:

$$EMD_d(S^q, S^o) = \min_{f_{ij}} \left\{ \frac{\sum_i \sum_j f_{ij} \cdot d(c_i^q, c_j^o)}{\min\{\sum_i w_i^q, \sum_j w_j^o\}} \right\},$$

subject to the constraints: $\forall i : \sum_j f_{ij} \leq w_i^q$, $\forall j : \sum_i f_{ij} \leq w_j^o$, $\forall i, j : f_{ij} \geq 0$, and $\sum_i \sum_j f_{ij} = \min\{\sum_i w_i^q, \sum_j w_j^o\}$.

These constraints guarantee a feasible solution, i.e., all costs are positive and do not exceed the limitations given by the weights in both feature signatures. In this paper, we assume the ground distance d is the L_p distance over a D-dimensional feature space $\mathbb{F} = \mathbb{R}^D$, defined as

$$L_p(u, v) = \left(\sum_{i=1}^{D} |u_i - v_i|^p \right)^{1/p}$$

While for the parameter $p \geq 1$ the L_p distance is a metric (so-called *Minkowski metric*), for the parameter $0 < p < 1$ it becomes nonmetric (so-called *fractional* L_p *distance [1]*) as it violates the triangle inequality.

Based on the notion of the EMD and its inherent (non)metric ground distance, we continue with summarizing related work in the next section.

 ISBN: 978-1-61208-195-3

III. Related Work

As there is a minimization problem to solve within the EMD evaluation, the computation time complexity is considerably high (between $O(n^3)$ and $O(n^4)$). Hence, techniques providing efficient similarity search in large multimedia databases using the Earth Mover's Distance are necessary. In the following paragraphs, we shortly summarize the state-of-the-art orthogonal approaches used to efficiently process similarity queries in EMD-based distance spaces.

A. Domain Specific Filters

If we assume the ground distance d as a Minkowski metric (L_p distance, $p \geq 1$), a very simple yet efficient lower-bound for EMD is the *Rubner filter* [15], defined as

$$EMD(S^q, S^o) \geq \sqrt[1/p]{\sum_{k=1}^{D} \left| \sum_{i=1}^{n} w_{q_i} q_{ik} - \sum_{j=1}^{m} w_{o_j} o_{jk} \right|^p},$$

where m, n are the sizes of S^o, S^q, respectively. The Rubner filter holds only if $\sum_{i=1}^{n} w_{q_i} = \sum_{j=1}^{m} w_{o_j}$, otherwise, it becomes an approximate method.

A novel dimensionality reduction techniques for the EMD in a two-step filter-and-refine architecture for efficient exact search can be found in [2], [3]. The authors assume feature histograms and metric ground distances. In [19], Wichterich et al. utilized dimensionality reduction to improve the time of EMD evaluation. They proved that EMD evaluated in a low-dimensional subspace lower-bounds the EMD in the original space. Again, only histograms are applicable to that technique.

B. Transformation to Wavelet Domain

In [17], Shirdhonkar and Jacobs presented a linear-time algorithm for approximating the EMD for low-dimensional histograms using the sum of absolute values of the weighted wavelet coefficients of the difference histogram. Since the EMD computation is a special case of the Kantorovich-Rubinstein transshipment problem, the method can exploit the Hölder continuity constraint in its dual form to convert it into a simple optimization problem with an explicit solution in the wavelet domain. The authors proved the resulting wavelet EMD metric is equivalent to EMD, i.e., the ratio of the two is bounded. The bound estimates were also provided.

C. Metric access methods

Another approach for efficient indexing of the Earth Mover's Distance could be the distance-based indexing, especially the *metric access methods* [6], [20]. These methods utilize precomputed distances stored in a metric index to estimate the lower-bound of the original distance between a query object q and a database object o. In Figure 2, we depict an example of $\delta(q, o)$ lower-bound estimation using one reference point p, where $\delta(p, o)$ is the precomputed distance stored in the metric index and $\delta(q, p)$ is evaluated at the beginning of query processing. In the case $LB(\delta(q, o))$ is greater than the actual query radius r (considering range or kNN query), the original expensive distance $\delta(q, o)$ does not have to be evaluated.

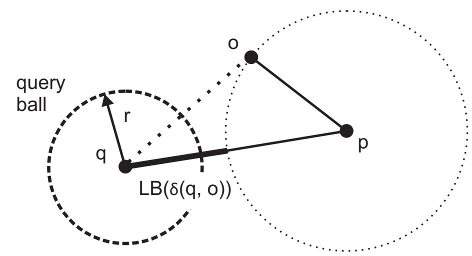

Fig. 2. Lower-bound estimation using a reference point p.

However, the distance spaces based on the Earth Mover's Distance usually suffer from high *intrinsic dimensionality* [6] (low indexability, i.e., $LB(\delta(q, o))$ is mostly lower than actual query radius r), so that only approximate techniques can be used for efficient search. To overcome this limitation, we proposed a parameterized version of the Earth Mover's Distance (pEMD) [12], defined as:

$$pEMD_d(S^q, S^o, w) = \min_{f_{ij}} \left\{ \frac{\sum_i \sum_j f_{ij} \cdot \mathrm{FP}(d(c_i^q, c_j^o), w)}{\min\{\sum_i w_i^q, \sum_j w_j^o\}} \right\},$$

where FP is the *fractional-power modifier* [18] defined as:

$$\mathrm{FP}(x, w) = \begin{cases} x^{\frac{1}{1+w}} & \text{for } w > 0 \\ x^{1-w} & \text{for } w \leq 0 \end{cases}$$

The intuition behind the FP-modifier is rather simple. Depending on the weight parameter w, we can either suppress ($w > 0$) or strengthen ($w < 0$) the transportation costs to outlier features when comparing feature signatures. Hence, the robustness of the measure can be tuned by setting the impact of outliers (noise bins or clusters) on the overall distance. However, the parameterized Earth Mover's Distance is no longer a metric when employing the FP-modifier ($w < 0$), or fractional L_p distance ($p < 1$), or not-normed weights.

All the techniques mentioned above are based on domain-specific solutions for (often low-dimensional) feature histograms, metric L_p distances ($p \geq 1$), or they utilize approximate similarity search to speed up query processing.

IV. Indexing the Earth Mover's Distance

In this paper, we investigate the general problem of EMD-based distance spaces employing feature signatures and non-metric ground distances. Since we consider feature signatures, the original Rubner filter and the distance-based indexing (metric access methods) can only be used to speed up query processing in large multimedia databases.

A. Revisiting the filter-and-refinement scheme

First of all, we would like to reconsider the filter-and-refinement scheme proposed in [2], where Assent et al. used a chain of filters. In general, when trying to exclude irrelevant database objects from the search process, we apply lower-bounding filters that progressively reduce the candidate set of the results. To optimize this process, we first apply the

 ISBN: 978-1-61208-195-3

cheapest filters (e.g., taking less than $O(n)$ time to compute the lower bound, where n is the signature size) and then apply the more expensive ones (which are usually also more effective in filtering). This corresponds to the optimal multi-step search process described in [16].

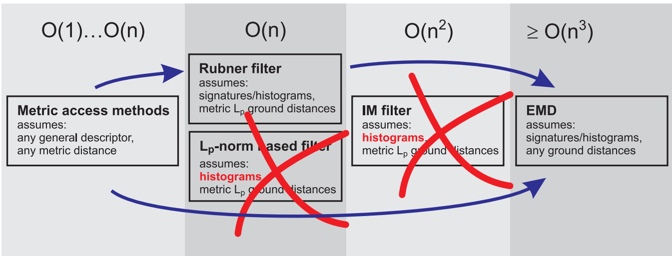

Fig. 3. Filter chain used in the filter-and-refinement scheme for EMD search.

In the aforementioned paper [2], Assent et al. proposed similarity query processing using the cheap Rubner filter ($O(n)$) and subsequently the more expensive independent minimization (IM) filter ($O(n^2)$). The two filters significantly reduce the final set of candidate objects, so that the expensive EMD evaluation is needed only for a fraction of the database. One of the contributions of this paper is the extension of the class of cheap filters by the metric access methods (MAM), which reduce the complexity of lower bound computation from $O(1)$ to $O(n)$. We depict an overview of the proposed filter-and-refinement scheme for efficient EMD-based similarity search in Figure 3. Note, that we do not consider the IM filter and L_p-norm filter anymore, as they only support histograms.

In particular, we combine the Rubner filter and an MAM-based filter (e.g., the pivot tables as detailed in the Section V-C), both cheap and supporting feature signatures. Having a Rubner filter lower-bound LB_{Rub} and the corresponding MAM lower-bound LB_{MAM}, we select the larger estimate $max(LB_{Rub}, LB_{MAM})$. To further improve the performance of the whole filter-and-refinement scheme, we also plan to generalize the IM filter for feature signatures (subject of our future work).

B. Nonmetric ground distances

As another contribution of this paper, we investigate the impact of the nonmetric L_p ($p < 1$) ground distances that have been frequently used in the image retrieval for robust image matching [8], [10]. For instance, consider the three feature signatures depicted in Figure 4, where S^q stands for a query signature and S^x, S^y for database signatures. The first two signatures S^q, S^x consist of very similar distributions of feature clusters, while the third one S^y is slightly different. To provide robust image ranking by the EMD, we have to employ such a ground distance, that will neglect the impact of the outlying noisy cluster c_4 in S^x. The fractional L_p ground distances are a good choice for such purpose, as they decrease the impact of outlying distances to the overall aggregation provided by the EMD. A similar robust behavior can be achieved by using the FP-modifier within the parameterized EMD. Hence, we have tools for fine-tuning the quality of the EMD-based similarity retrieval in a particular database.

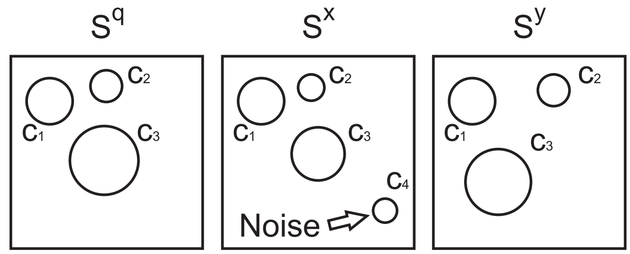

Fig. 4. Comparing a query signature with a database signature.

The main disadvantage of the fractional L_p distances is that they are not metrics. In particular, they lose the triangle inequality, assumed by both the Rubner filter and metric access methods to guarantee exact search. In turn, the nonmetric EMD (incorrectly) employed as a metric brings the possibility of a retrieval error – false dismissals and/or false positives in the query result. Nevertheless, we can control the retrieval error by tuning the parameter w of the parameterized EMD, while as showing in the following experimental section, the retrieval error caused by the ”nonmetricity” of the EMD is not significant.

V. Experimental Results

We conducted an experimental evaluation on the MIR Flickr [11] and ALOI [9] databases comprising 25,000 and 72,000 images, respectively. We have extracted feature signatures based on color, position, and texture information, similar to [4], where each image was represented by several 7-dimensional feature centroids and each centroid was assigned a weight. The feature signatures of the ALOI database consist of 12-140 centroids, 54 centroids on average, and those for the MIR Flickr database consist of 8-150 centroids, 57 centroids on average. Although the selected databases were not very large, they provided the ground truth for evaluating the search effectiveness. We have examined 6 variants of the Earth Mover's Distance, each using different L_p ground distances. Three variants were nonmetrics $\{L_{0.25}, L_{0.5}, L_{0.75}\}$ and the other three were metrics $\{L_1, L_2, L_5\}$. As a metric access method, we used simple Pivot tables (using 50 pivots) [6], [13]. In each test, we performed 100 kNN queries ($k = 10$) and averaged the results. As the main observables, we measured both the retrieval efficiency and the retrieval effectiveness. The efficiency was measured in terms of real time as well as in the number of EMD computations. The effectiveness was measured using the mean average precision (i.e., employing the ground truth) and also using the retrieval error defined as the deviation from the referential result obtained by the sequential scan. The tests ran on a workstation 2x Intel Xeon X5660 2.8 Ghz, 24GB RAM, Windows Server 2008 R2 64bit (non-virtualized).

A. Basic tests

The graphs in Figure 5 depict the intrinsic dimensionality and the mean average precision values for the aforementioned databases by changing the parameter w of pEMD (denoted FP weight). The intrinsic dimensionality decreases for both databases with decreasing w, while mean average precision stays at a considerably high level of greater than 0.6 for the ALOI database and 0.26 for the MIR Flickr database. Since pEMD changes the ground distance matrix and thus changes the number of iterations needed to find the optimal solution of pEMD, we have also observed query processing times for the sequential scan.

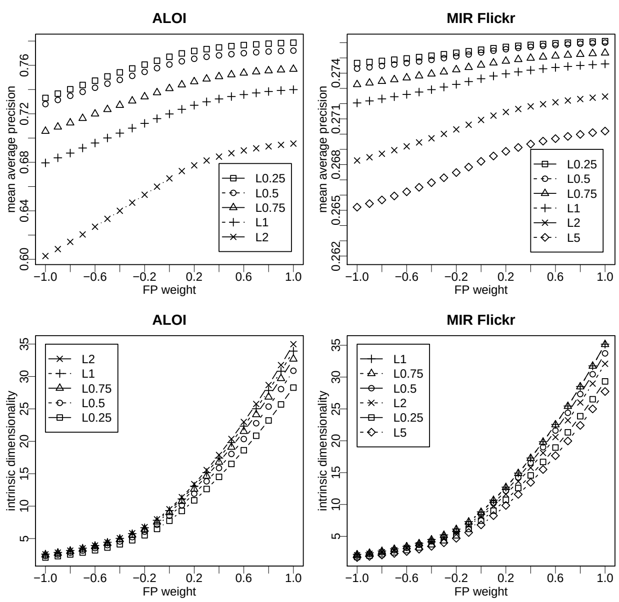

Fig. 5. The intrinsic dimensionality and mean average precision for MIR Flickr and ALOI databases.

As can be seen from the graphs in Figure 6, increasing the parameter w leads to a lower number of iterations, which results in decreased realtime cost. Also note that the L_1 ground distances (i.e., $L_p, p = 1$) led to fastest responses due to the absence of powering/rooting by p in the L_p formula.

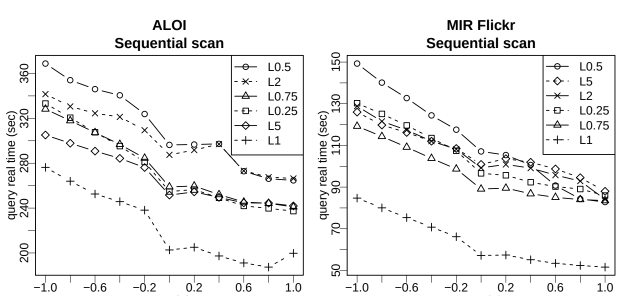

Fig. 6. The realtime of sequential scan for MIR Flickr and ALOI databases.

B. Rubner filter

In the second set of experiments, we evaluated the performance of the Rubner filter for various values of pEMD's parameter w. As we can observe in Figure 7, both realtime and the number of EMD evaluations decreases for higher w. This is caused by the fact, that the Rubner estimation of EMD is more similar to the real EMD value. However, this leads also to a small retrieval error, because the estimation is no longer guaranteed a lower bound and so some relevant objects may be filtered, as shown in Figure 8. We can also observe a good performance of $L_{0.25}$, however this ground distance results also in higher retrieval error.

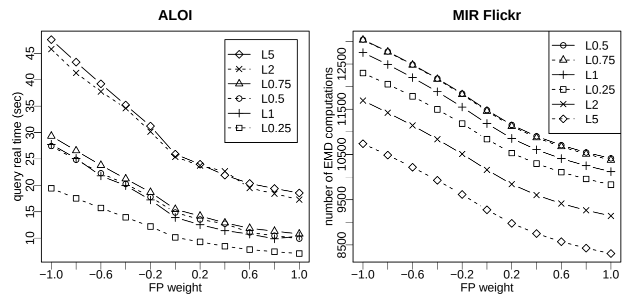

Fig. 7. Effect of the Rubner filter – query real time.

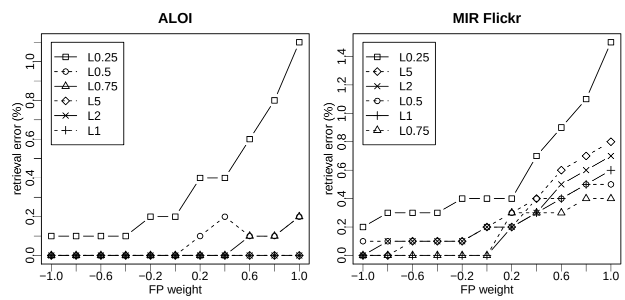

Fig. 8. Effect of the Rubner filter – retrieval error.

C. Rubner Filter combined with the Metric Access Methods

In the last experiments, we combined the Rubner filter with a simple metric access method – the Pivot Table (the original LAESA) [13]. From the Figure 9, we may observe the impact of low intrinsic dimensionality on pivot table filtering. However, the efficiency is coupled with high retrieval error (a lot of false dismissals), see Figure 10. Nevertheless, if the user accepts 5% error then the query is more than twice as fast as using just the Rubner filter. If we require the highest retrieval precision, it is better to utilize the Rubner filter and nonmetric ground distances rather than metric access methods and metric ground distances.

D. Discussion

In the experiments, we have focused on the effective similarity search using nonmetric ground distances employed in the Earth Mover's Distance. To the best of our knowledge, this

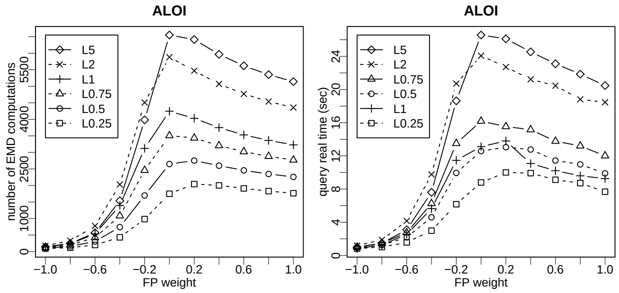

Fig. 9. Effect of the Rubner filter combined with Pivot tables – query cost.

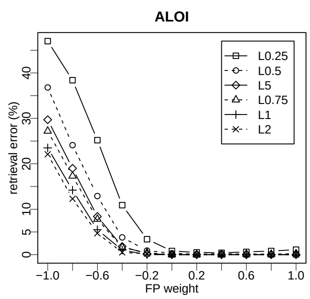

Fig. 10. Effect of the Rubner filter combined with Pivot tables – retrieval error.

paper is the first one, that combines and investigates the effect of two filtering approaches used for efficient query processing – the general metric spaces approach and the domain specific Rubner filter. We also more deeply investigate the effect of the recently introduced Earth Mover's Distance parameter w, that in connection with the Rubner filter can result in more efficient and more effective similarity search. In the case a fast query processing is required and the precision is not preferred, we can employ the metric access methods as an approximate search technique for lower values of w.

Our experimental evaluation reveals that the combination of filter-and-refinement schemes and nonmetric ground distances within the EMD provides the best retrieval results. In particular, the trade-off between efficiency and effectiveness is given by comparatively low query response times and small retrieval errors. Thus, we conclude that our proposed approach is able to outperform state-of-the-art metric indexing solutions for the EMD.

VI. Conclusion and Future Work

We have investigated a parameterized version of the Earth Mover's Distance in combination with nonmetric ground distances. In the experimental evaluation, we showed that nonmetric ground distances can be utilized for effective and efficient similarity search in multimedia databases by the parameterized Earth Mover's Distance. More specifically, using a revisited filter-and-refinement schema, the system administrators can provide user with more sophisticated exploration facilities, e.g., more effective but less efficient retrieval, or vice versa. In the future, we would like to formally describe the observed behavior and investigate other modifying functions that can provide more desirable properties (e.g., preserve metric axioms). We also want to generalize other EMD filters for feature signatures and nonmetrics.

Acknowledgment

This research has been partially supported by Czech Science Foundation (GAČR) projects 202/11/0968 and P202/12/P297, and by the Deutsche Forschungsgemeinschaft within the Collaborative Research Center SFB 686.

References

[1] C. C. Aggarwal, A. Hinneburg, and D. A. Keim. On the surprising behavior of distance metrics in high dimensional spaces. In *Proc. 8th International Conference on Database Theory*, ICDT '01, pages 420–434, London, UK, 2001. Springer-Verlag.

[2] I. Assent, A. Wenning, and T. Seidl. Approximation techniques for indexing the earth mover's distance in multimedia databases. In *IEEE ICDE*, pages 11–22, 2006.

[3] I. Assent, M. Wichterich, T. Meisen, and T. Seidl. Efficient similarity search using the earth mover's distance for large multimedia databases. In *IEEE ICDE*, pages 307–316, 2008.

[4] C. Beecks, M. S. Uysal, and T. Seidl. A comparative study of similarity measures for content-based multimedia retrieval. In *IEEE ICME*, pages 1552–1557, 2010.

[5] C. Beecks, M. S. Uysal, and T. Seidl. Signature quadratic form distance. In *Proc. ACM CIVR*, pages 438–445, 2010.

[6] E. Chávez, G. Navarro, R. Baeza-Yates, and J. L. Marroquín. Searching in Metric Spaces. *ACM Computing Surveys*, 33(3):273–321, 2001.

[7] T. Deselaers, D. Keysers, and H. Ney. Features for image retrieval: an experimental comparison. *Information Retrieval*, 11(2):77–107, 2008.

[8] M. Donahue, D. Geiger, T. Liu, and R. Hummel. Sparse representations for image decomposition with occlusions. In *CVPR*, pages 7–12, 1996.

[9] J.-M. Geusebroek, G. J. Burghouts, and A. W. M. Smeulders. The Amsterdam Library of Object Images. *IJCV*, 61(1):103–112, 2005.

[10] P. Howarth and S. Ruger. Fractional distance measures for content-based image retrieval. In *ECIR 2005*, pages 447–456. LNCS 3408, Springer-Verlag, 2005.

[11] M. J. Huiskes and M. S. Lew. The mir flickr retrieval evaluation. In *Proc. ACM MIR*, pages 39–43, 2008.

[12] J. Lokoč, C. Beecks, T. Seidl, and T. Skopal. Parameterized earth mover's distance for efficient metric space indexing. In *Proceedings of the Fourth International Conference on SImilarity Search and APplications*, SISAP '11, pages 121–122, New York, NY, USA, 2011. ACM.

[13] M. L. Mico, J. Oncina, and E. Vidal. A new version of the nearest-neighbour approximating and eliminating search algorithm (aesa) with linear preprocessing time and memory requirements. *Pattern Recogn. Lett.*, 15(1):9–17, 1994.

[14] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(10):1615–1630, 2005.

[15] Y. Rubner, C. Tomasi, and L. J. Guibas. The Earth Mover's Distance as a Metric for Image Retrieval. *IJCV*, 40(2):99–121, 2000.

[16] T. Seidl and H.-P. Kriegel. Optimal multi-step k-nearest neighbor search. In *Proc. ACM SIGMOD*, pages 154–165, 1998.

[17] S. Shirdhonkar and D. Jacobs. Approximate earth movers distance in linear time. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1 –8, june 2008.

[18] T. Skopal. Unified framework for fast exact and approximate search in dissimilarity spaces. *ACM Transactions on Database Systems*, 32(4):1–46, 2007.

[19] M. Wichterich, I. Assent, P. Kranen, and T. Seidl. Efficient emd-based similarity search in multimedia databases via flexible dimensionality reduction. In *Proc. ACM SIGMOD*, pages 199–212, 2008.

[20] P. Zezula, G. Amato, V. Dohnal, and M. Batko. *Similarity Search: The Metric Space Approach (Advances in Database Systems)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.

 ISBN: 978-1-61208-195-3