Computational Geometry 51 (2016) 40–54

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

Bottleneck partial-matching Voronoi diagrams and applications

Matthias Henze [1], Rafel Jaume [*,2]

Institut für Informatik, Freie Universität Berlin, Takustraße 9, D-14195 Berlin, Germany

A R T I C L E   I N F O

Article history:
Received 4 December 2014
Accepted 29 September 2015
Available online 9 October 2015

Keywords:
Bottleneck matching
Voronoi-type diagram
Partial point matching under translations
Bipartite graph

A B S T R A C T

We study the minimization of the bottleneck distance between a point set B and an equally-sized subset of a point set A under translations. We relate this problem to a Voronoi-type diagram and derive polynomial bounds for its complexity that are optimal in the size of A. We devise efficient algorithms for the construction of such a diagram and its lexicographic variant, which generalize to higher dimensions. We use the diagram to find an optimal bottleneck matching under translations, to compute a connecting path of minimum bottleneck cost between two positions of B, and to determine the maximum bottleneck cost in a convex polygon.

1. Introduction

Applications often demand algorithms to find an occurrence of a point pattern in a given cloud of points. Using a suitable cost function, it is common to define a similarity measure between the pattern and the point cloud as the minimum cost among the images of the pattern under a set of allowed transformations. One of the most studied similarity measures between finite point sets A and B in $\mathbb{R}^d$ is the *directed Hausdorff distance*, which is the maximum of the (Euclidean) distances from each point in B to its nearest neighbor in A.

For some applications in robotics and pattern recognition, it is required that each point of the smaller set is matched to a distinct point in the bigger one. This is modeled by the *bottleneck distance*, which is defined for point sets A and B with $|B| \le |A|$, as

$$\Delta(B, A) = \min_{\sigma: B \hookrightarrow A} \max_{b \in B} \|b - \sigma(b)\|,$$

where $\|\cdot\|$ denotes the Euclidean norm and the minimum is taken over all injections from B into A. This notion was defined in [2] for the case when A and B have the same size.

In contrast to the directed Hausdorff distance, the bottleneck distance has the advantage of being symmetric for equally-sized sets. On the other hand, it is harder to compute, since the points cannot be regarded independently. Note that there might be several matchings that minimize the bottleneck distance, even when all the distances between points are distinct. This can be avoided by considering the matching that lexicographically minimizes the distances between matched points; cf. [7,11,21].

* Corresponding author.
E-mail addresses: matthias.henze@fu-berlin.de (M. Henze), jaume@inf.fu-berlin.de (R. Jaume).

1 Supported by the ESF EUROCORES program EuroGIGA-VORONOI, (DFG): RO 2338/5-1.
2 Supported by "Obra Social la Caixa" and DAAD.

http://dx.doi.org/10.1016/j.comgeo.2015.10.001
0925-7721/

In this paper, we are interested in a dynamic version of the bottleneck distance. More precisely, we want to efficiently compute, among all translated copies of B with respect to A, one attaining the minimum bottleneck distance; that is, $\min_{t\in\mathbb{R}^d} \Delta(B+t, A)$. This problem will be called *bottleneck partial-matching under translations*. It was introduced for equally-sized point sets in the plane by Alt, Mehlhorn, Wagener & Welzl [2], who gave an algorithm running in $O(n^6 \log n)$ time for point sets of size n. Their bound was later improved to $O(n^5 \log^2 n)$ by Efrat, Itai & Katz [14].

To the best of our knowledge, bottleneck matching under translations has not been studied with the focus on algorithms whose complexity is sensitive to the size of the smaller set. In order to do so, we associate Voronoi-type diagrams to the problem, which we call *bottleneck diagrams* and *lex-bottleneck diagrams*, respectively. This follows an idea of Rote [20] who partitioned the space of translations according to the (partial) matching that minimizes the least-squares distance between translated copies of B and A (cf. [17,4,3] for follow-up studies). Our bottleneck diagrams partition $\mathbb{R}^d$ into polyhedral cells that can be mapped to locally-optimal lexicographic bottleneck matchings.

Our motivation to investigate these diagrams does not restrict to solving the bottleneck partial-matching problem under translations only. We moreover aim to provide a structure that may be either used for preprocessing or may be adjusted towards other problems that are based on the computation of the bottleneck distance in various translated positions of the point sets. The applications at the end of the paper exemplify this utility of the bottleneck diagrams.

A non-archival abstract containing parts of our studies appeared in [16].

Our results In Section 3, we formally introduce the Voronoi-type diagrams before investigating their basic properties and combinatorial complexity. It turns out that there exists a lex-bottleneck diagram (and, hence, a bottleneck diagram) of complexity $O(n^2k^6)$ for any given planar point sets $A, B \subset \mathbb{R}^2$ with $k = |B| \le |A| = n$ (see Theorem 3.12), and that this bound cannot be improved with respect to the size of A. For any pair of point sets $A, B \subset \mathbb{R}^d$ of higher dimensions we obtain in Corollary 3.6 that there is a lex-bottleneck diagram of complexity $O(n^{2d}k^{2d})$. Based on this complexity result, we devise algorithms in Section 4 that construct these polyhedral subdivisions of $\mathbb{R}^d$ and at the same time compute a lexicographic bottleneck matching for each of the cells of the subdivision, which we will use to solve the bottleneck matching problem under translations. This is achieved with an overhead of $O(k^2)$ for the bottleneck diagrams, and $O(k^4)$ for the lexicographic variant (see Theorems 4.8 and 4.9). The matching problem under translations can then be solved for the bottleneck case in time $O(n^2k^8)$, and for the lexicographic variant in time $O(n^2k^{10})$, if the point sets are planar (see Theorem 5.1). In higher dimensions the time bounds are $O(n^{2d}k^{2d+2})$ and $O(n^{2d}k^{2d+4})$, respectively. Finally, Theorems 5.5 and 5.7 show how we can use the bottleneck diagrams to solve two related problems. The first one is to efficiently compute a path of minimum bottleneck cost between two positions of a pattern in the plane. The second application is to determine what we call the cover radius of a polygon: the maximum of the bottleneck costs attained by a pair of point sets when one of them is allowed to be translated only inside a given polygon.

Comparison to previous work Although neither Alt et al. [2] nor Efrat et al. [14] consider the bottleneck matching problem for different-sized point sets, their methods can be adapted to this situation without major difficulties. In Appendix A, we elaborate on an analysis of such adaptations and derive the time bounds $O(n^3k^3 \log n)$ and $O(n^2k^3 \log^2 n)$, respectively, where k is the size of the smaller set and n the size of the bigger one. Hence, the bound we derive for our algorithm is better when $k = o(\log^{2/5} n)$. In addition to improving the bound for small k, the use of bottleneck diagrams is conceptually different from previous methods, and has the advantage of being applicable to solve the matching under translations problem in any dimension and, moreover, with respect to the lexicographic bottleneck cost. Only approximation algorithms for the bottleneck matching problem in higher dimensions were known previously (see [14]).

Organization of the paper In the next section, we introduce (lexicographic) bottleneck matchings between two finite point sets and investigate corresponding minimization diagrams and their properties. After these basics, we define our main objects of study, bottleneck partial-matching Voronoi diagrams, and analyze their combinatorial complexity in Section 3, before addressing construction algorithms for these structures in Section 4. Finally, in Section 5, we apply the bottleneck diagrams to solve the bottleneck partial-matching problem and related questions.

2. Bottleneck and lexicographic bottleneck matchings

In this section, we introduce bottleneck matchings and discuss the minimization diagram corresponding to the bottleneck partial-matching problem under translation. The issues we encounter explain our approach to the definition of the bottleneck diagrams in Section 3.

Throughout the paper, we assume that we are given two point sets $A, B \subset \mathbb{R}^d$ with $k = |B| \le |A| = n$ and that B is allowed to be translated. We use the term *edge* for a pair of points $(a, b) \in A \times B$ and denote it by ab for short. The *length* of the edge ab is defined as the Euclidean distance $\|b - a\|$. In this context, we identify every injection of B into A with the *matching*, i.e., the set of edges, it induces. The cost of such a matching varies according to a parameter representing the position of the point set B.