Discrete Applied Mathematics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

On the geometry of graph spaces

Brijnesh J. Jain

Technische Universität Berlin, Germany

ARTICLE INFO

Article history:
Received 29 May 2015
Received in revised form 22 June 2016
Accepted 24 June 2016
Available online xxxx

Keywords:
Graphs
Graph metric
Orbit space
Pattern recognition

ABSTRACT

Optimal alignment kernels are graph similarity functions defined as pointwise maximizers of a set of positive-definite kernels. Due to the max-operation, optimal alignment kernels are indefinite graph kernels. This contribution studies how the max-operation transforms the geometry of the associated feature space and how standard pattern recognition methods such as linear classifiers can be extended to those transformed spaces. The main result is the Graph Representation Theorem stating that a graph is a point in some geometric space, called orbit space. This result shows that the max-operation transforms the feature space to a quotient by a group action. Orbit spaces are well investigated and easier to explore than the original graph space. We derive a number of geometric results, translate them to graph spaces, and show how the proposed results can be applied to statistical pattern recognition.

1. Introduction

In different fields of structural pattern recognition, such as computer vision, chemo- and bioinformatics, objects are naturally represented by attributed graphs [6,12,16,40]. One persistent problem is the gap between structural and statistical methods in pattern recognition [5,8]. This gap refers to a shortcoming of powerful mathematical methods that combine the advantages of structural representations with the advantages of statistical methods defined on Euclidean spaces.

One reason for this gap is an insufficient understanding of the geometric structure of graph spaces. The geometry of any space depends on the choice of the underlying distance function. For example, spaces endowed with an intrinsic metric have a richer geometric structure than spaces endowed with an arbitrary distance. The prime example are Euclidean spaces, which have a fundamental status in mathematics. The abundant geometric structure of Euclidean spaces is introduced by an inner product that in turn induces the Euclidean metric.

Similarly as inner products, optimal alignment kernels have the potential to introduce a less abundant but still well-developed geometric structure into graph spaces. The potential arises from the fact that the definition of optimal alignment kernels implicitly involves inner products. An optimal alignment kernel is a graph similarity function defined as a pointwise maximizer of a set of positive-definite kernels, where the maximum is taken over all possible one-to-one correspondences (alignments) between the nodes of both graphs. Since positive-definite kernels are not closed under the max-operation [1,44], optimal alignment kernels are indefinite graph kernels. Consequently, the geometric structure introduced by optimal alignment kernels is limited.

Although it is well-known that graph alignment kernels are indefinite graph kernels, it is unclear how the max-operation transforms the geometry of the feature space associated to the underlying positive-definite kernels. In addition, it is also unclear how pattern recognition methods like linear classifiers can be extended to those transformed feature spaces.

E-mail address: brijnesh.jain@gmail.com.

http://dx.doi.org/10.1016/j.dam.2016.06.027

In this contribution, we study the geometry of graph spaces endowed with an optimal alignment kernel. The main result is the Graph Representation Theorem stating that a graph is a point in some geometric space, called orbit space. This result shows that the max-operation transforms the associated feature space to a quotient of a Euclidean space by a group action. Orbit spaces are well investigated, provide access to a plethora of existing results, and are much easier to explore than graph spaces. We derive several geometric concepts and results, translate them to graph spaces, and show how the proposed results can be applied to generalize linear classifiers and other statistical pattern recognition methods to graph spaces. The proposed orbit space framework provides a mathematical foundation for narrowing the gap between structural and statistical pattern recognition and places existing learning methods in graph spaces on a theoretically sound basis.

The rest of this paper is structured as follows: Section 2 discusses related work. Section 3 introduces graph alignment spaces. In Sections 4 and 5, we present the Graph Representation Theorem and derive general geometric results. Section 6 places the results into the context of statistical pattern recognition. Finally, Section 7 concludes with a summary of the main results.

2. Related work

Graph alignment kernels and their induced distances form a common and widely applied class of graph (dis-)similarity functions [46]. Though graph alignment kernels are indefinite graph kernels, they have been applied to support vector learning by Geibel et al. [18,28] and shortly after by Fröhlich et al. [14].

The problem of computing such graph (dis-)similarities is referred to as the graph matching problem [6,12,40,46]. The graph matching problem can be formulated as a quadratic assignment problem [41], which is known to be NP-complete [17]. Consequently, the majority of work on graph alignment distances is devoted to devise efficient algorithms and heuristics for solving the underlying graph matching problem [46].

In contrast to the amount of work on practical and computational aspects of the graph matching problem, there is only few work on theoretical properties of graph spaces [23,31]. Hurshman and Janssen showed that graph spaces endowed with distances based on the maximum common subgraph induce a discrete topology on the set of isomorphism classes of graphs [23]. Then they introduced and studied the concept of uniform continuity of graph parameters. Graph spaces endowed with an optimal alignment kernel have been first investigated by [31], where the focus was on analytical properties for the differential machinery.

A similar approach as in [31] has been followed by Feragen et al. [9,11,10] for tree shapes. Tree-shape spaces are also orbit spaces, but additionally include continuous transitions in the tree topology. This additional requirement leads to a substantially different geometric structure. Feragen et al. studied the geodesic structure of tree-shape spaces that transfers to graph spaces as correctly hypothesized by Feragen et al. (see Theorem 4.7). More generally, this contribution can be placed into the context of statistical analysis of non-Euclidean spaces in the spirit of [42,45] for complex objects, for tree-structured data [10,45], and for shapes [7,37]. Examples of statistical methods adapted to graphs are the sample mean [19,29,30,36], central clustering algorithms [20,21,35,29], learning vector quantization [32,34], and generalized linear classifiers [25,24].

3. Graph alignment spaces

3.1. Attributed graphs

Let $\mathcal{A}$ be the set of node and edge attributes. We assume that $\mathcal{A}$ contains a distinguished symbol ν denoting the null-attribute.

Definition 3.1. An attributed graph is a triple $X = (\mathcal{V}, \mathcal{E}, \alpha)$ consisting of a finite set $\mathcal{V} \neq \emptyset$ of nodes, a set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ of edges, and an attribute function $\alpha : \mathcal{V} \times \mathcal{V} \to \mathcal{A}$ such that $\alpha(i, j) = \nu$ if and only if $(i, j) \notin \mathcal{E}$ for all distinct nodes $i, j \in \mathcal{V}$.

The attribute function α assigns an attribute to each pair of nodes. Node attributes may take any value from $\mathcal{A}$, edges have non-null attributes, and non-edges are labeled with null-attribute ν. The node set of a graph X is occasionally referred to as $\mathcal{V}_X$, its edge set as $\mathcal{E}_X$, and its attribute function as α_X. By $\mathcal{G}_{\mathcal{A}}$ we denote the set of attributed graphs with attributes from $\mathcal{A}$. The order $|X|$ of graph X is its number $|\mathcal{V}_X|$ of nodes.

Graphs can be directed and undirected. Attributes can take any values including binary values, discrete symbols, numerical values, vectors, strings, and combinations thereof. Thus, the definition of attributed graph is sufficiently general to cover a wide class of graphs such as binary graphs from graph theory, weighted graphs, molecular graphs, protein structures, and many others.

3.2. Optimal alignment kernels

The inner product induces a rich geometric structure into Euclidean spaces. In a similar way, we want to induce a geometric structure into graph spaces by adapting the inner product to graphs. The challenge is that neither the graph space $\mathcal{G}_{\mathcal{A}}$ nor the attribute set $\mathcal{A}$ are vector spaces, in general. A solution to this problem are optimal alignment kernels. An