Computational Geometry 47 (2014) 164–174

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

A Fourier-theoretic approach for inferring symmetries

Xiaoye Jiang [a,*], Jian Sun [b], Leonidas Guibas [c]

[a] Institute for Computational and Mathematical Engineering, Stanford University, USA
[b] Mathematical Sciences Center, Tsinghua University, China
[c] Computer Science Department, Stanford University, USA

ARTICLE INFO

Article history:
Received 5 October 2011
Accepted 1 October 2012
Available online 9 October 2012
Communicated by David Bremner and Greg Aloupis

Keywords:
Symmetry
Fourier theory
Permutation group
Geometric processing

ABSTRACT

In this paper, we propose a novel Fourier-theoretic approach for estimating the symmetry group $\mathbb{G}$ of a geometric object X. Our approach takes as input a geometric similarity matrix between low-order combinations of features of X and then searches within the tree of all feature permutations to detect the sparse subset that defines the symmetry group $\mathbb{G}$ of X. Using the Fourier-theoretic approach, we construct an efficient marginal-based search strategy, which can recover the symmetry group $\mathbb{G}$ effectively. The framework introduced in this paper can be used to discover symmetries of more abstract geometric spaces and is robust to deformation noise. Experimental results show that our approach can fully determine the symmetries of various geometric objects.

1. Introduction

Symmetries are extremely common in both man-made and natural objects. They are especially important when decorative arts are involved, as in architecture, interior or ornament design. Arguably, symmetry detection is one of the fundamental ways in which we understand the world around us and it has been of interest to scientists as well as artists for centuries. In the modern era, a mathematical study of symmetries has led to a theory integrating algebra and geometry. In the context of computational geometry, we often consider the *symmetry group* of a geometric object with a pre-defined metric. Such a symmetry group is the group of all isometric transformations, which preserve the metric. In a wider context, a symmetry group may be any kind of transformation group, as long as we know what kind of mathematical structures we are concerned with so that we can decide which mappings preserve the structure. Conversely, specifying the symmetry group can also define the structure, or at least clarify what we mean by an invariant; this is one way of looking at the *Erlangen Program* of Felix Klein [10].

One easy way of describing symmetries of geometric objects is to look at *group actions*, where we use a set to represent the object, and the symmetries of the object are described by bijective mappings on the set. In this paper, we assume we are given a discrete set $X = \{x_i\}_{i=1}^n$ which describes a geometric object. As shown in Fig. 1(a), the five tip points on a star can be used as a discrete set X to study the symmetry of such a 3D star model. This is because each symmetry of the star model can be identified with a permutation of the elements in X.[1] For convenience, we say a permutation is "good" if it can be identified as a symmetry of the geometric object. It can be shown that all good permutations of X consist of a group $\mathbb{G}$, which we call the *symmetry group of* X.

* Corresponding author.
E-mail addresses: xiaoyej@stanford.edu (X. Jiang), jsun@math.tsinghua.edu.cn (J. Sun), guibas@cs.stanford.edu (L. Guibas).

[1] Choosing different set X can result in different group actions. In this paper, however, we assume such a set X is given where each symmetry can be identified as permuting X.

0925-7721/$ – see front matter
http://dx.doi.org/10.1016/j.comgeo.2012.10.001

In practice, we are often limited to verifying low-order information about X, such as the similarity of curvatures between pairs of points (first order), or the consistency of distances between pairs of points (second order). But how can we integrate such low-order pieces of symmetry evidence together, so that we can identify all the good permutations of X and derive its symmetry group? This question is quite challenging since the space of all permutations grows factorially with the number of elements in X so that directly searching among all permutations is impossible, unless n is small. In this paper, we propose a Fourier-theoretic approach to address this problem, based on low-order similarities in X.

The symmetry group of X is a subgroup of the permutation group $\mathbb{S}_n$, where $n = |X|$. To search for $\mathbb{G}$, we naturally have the following simple strategy: we organize the elements of $\mathbb{S}_n$ in a tree where each leaf represents a particular permutation, and then search for those in $\mathbb{G}$ within the tree, see Fig. 1(b) for an illustration of the tree. Whenever we reach a permutation of X, we check whether it is a good one. However, such a brute force strategy would be computationally intractable for all but very small n.

In this paper, we propose a novel search strategy within the tree of $\mathbb{S}_n$, called the *marginal probability search*, which fully exploits the algebraic structure of the groups $\mathbb{G}$ and $\mathbb{S}_n$. We note that similar branch-and-bound search strategies for identifying symmetries have already been studied in the literature [19]. However, our paper contributes more in the following two ways of making use of algebraic structures to facilitate the search of symmetries.

Firstly, we consider the symmetry group $\mathbb{G}$ as an *indicator distribution* (see Theorem 2.1 in Section 2) over the permutation group $\mathbb{S}_n$. This novel point of view enables us to utilize techniques from the group representation theory to convert low-order information into a set of Fourier coefficients which characterizes the low-frequency components of the distribution over $\mathbb{S}_n$. Those Fourier coefficients can thus be used to efficiently estimate the marginal probability of the permutations represented by an internal node, which serves as the criterion for pruning the sub-tree rooted at that node. Unlike other traditional pruning criteria [5], the marginal probability is much more informative as it not only evaluates the part of the permutation which is already determined but also summarizes the remaining part which is undecided, and thus provides a more efficient pruning.

Secondly, we exploit the group structure of $\mathbb{G}$ and show that the internal nodes on the same level have either the same marginal probability as the node containing the identity permutation or 0 marginal probability for the indicator distribution $\mathbb{G}$. This ensures the correctness of taking the marginal probability of the internal nodes as the reference for pruning.

The approach proposed in this paper generalizes the existing work on graph automorphism, in the sense that we can deal with noisy similarity information and robustly estimate the symmetry group $\mathbb{G}$ from the input. Moreover, any arbitrary order of similarities, e.g., triple-wise or even high-order similarities can be taken as input in our framework. Different orders of similarities can be combined together without too much difficulty. This is due to the fact that the Fourier analysis can fully disentangle the information over permutations of different orders into orthogonal components. In addition, our approach does not require a concrete realization of the geometric object. As long as a discrete set X, which characterizes the symmetry of the geometric object, can be effectively extracted, our approach can be used for inferring the symmetry group $\mathbb{G}$. Symmetries of objects with heavy distortions or deformations can also be considered if self-similarity between their parts can be assessed.

Related work. In the geometry processing community, a great amount of research has been done on Euclidean symmetry detection [14,17,22], and even on general isometry detection [14,19]. Among them, voting is a commonly used strategy to detect symmetries. However, voting suffers from the curse of dimensionality, i.e., when the number of parameters needed to describe the symmetry is large, it becomes extremely inefficient. Recently in [5], a search strategy is proposed to find shape correspondences where a heuristic similar to the one used in beam search is employed to reduce the exponential complexity.

The problem of inferring the global symmetry from low-order similarities, especially from pairwise similarities, is closely related to the graph automorphism problem, or more generally, the colored graph automorphism problem. However, there are no known polynomial time algorithms for finding the automorphism group of a general graph except for certain special cases. For example, [25] developed a polynomial time algorithm for determine the automorphism group of a triply connected planar graph, which leads to a polynomial time algorithm for determine the planar graph isomorphism [3]; [1,11] showed that there exists a polynomial time algorithm for estimating a structural generating set for the automorphism group of a circulant graph, while the elements of the group still cannot be decided. Another related well-studied problem is the automorphism partitioning where the goal is to determine whether two vertices lie in the same orbits of the automorphism group. It is known that this problem is algorithmically equivalent to the problem of graph isomorphism [20]; A more generalized problem is the so-called 2-orbits automorphism partitioning which has also received great attention in the mathematics literature. In this problem, the goal is to decide whether different pairs of vertices lie in the same orbits. When the associated cellular algebras are Schurian [18], the 2-orbits automorphism partitioning can be immediately deduced from the cellular algebra. Though there are polynomial time algorithms for computing the cellular algebra of the *coherent configurations* or *association schemes* [2,26], it is not guaranteed that the correct orbits can be computed for an arbitrary graph.

In this paper, we take a novel view on the given low-order similarity information, that is to relate it to a marginal probability matrix of a distribution over the permutation group $\mathbb{S}_n$. Based on the Fourier analysis, we can completely disentangle and interpret the information within such a marginal matrix. This approach is quite general and can deal with any order of the input matrix in a unified analytical framework. Moreover, it can naturally deal with noisy input so that a