Pattern Recognition Letters 29 (2008) 1924–1933

Contents lists available at ScienceDirect

Pattern Recognition Letters

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

Solving jigsaw puzzles using image features

Ture R. Nielsen, Peter Drewsen, Klaus Hansen *

Department of Computer Science, University of Copenhagen, Universitetsparken 1, DK-2100 Copenhagen, Denmark

A R T I C L E I N F O

Article history:
Received 10 December 2006
Received in revised form 28 May 2008
Available online 14 June 2008

Communicated by R. Davies

Keywords:
Jigsaw puzzle solver
Edge matching
Piece classification
Border similarity measure
Co-occurrence matrix

A B S T R A C T

In this article, we describe a method for automatic solving of the jigsaw puzzle problem based on using image features instead of the shape of the pieces. The image features are used for obtaining an accurate measure for edge similarity to be used in a new edge matching algorithm. The algorithm is used in a general puzzle solving method which is based on a greedy algorithm previously proved successful. We have been able to solve computer generated puzzles of 320 pieces as well as a real puzzle of 54 pieces by exclusively using image information.

Additionally, we investigate a new scalable algorithm which exploits the divide and conquer paradigm to reduce the combinatorially complex problem by classifying the puzzle pieces and comparing pieces drawn from the same group. The paper includes a brief preliminary investigation of some image features used in the classification.

1. Introduction

Solving puzzles has been a popular pastime for ages and several researchers have recently looked into solving puzzles automatically by a computer. However, the problem is hard to solve for machines as compared to the relative ease with which humans can solve even very large puzzles. The computer typically does not have the same certainty of the correctness when matching a pair of puzzle pieces, but rather a likelihood for the correctness.

While the puzzle problem is a justified field of research in itself, research on puzzles might also be justified by the potential for real life applications of some of its components. The automatic solving of jigsaw puzzles deals with problems such as boundary detection, shape matching and texture comparison, which are of general interest in the fields of computer vision and pattern recognition. This paper is based on the assumptions that pieces have rigid edge curves constrained to a plane, and that the puzzle has no missing parts and has the topology of a plane. These assumptions are, however, not crucial for all sub-problems and the edge matching can for instance be extended for use in piecing together fragments of pictures or objects such as archaeological artifacts.

Previous works have all focused on edge matching using shape rather than image features. We therefore feel that using image features in solving jigsaw puzzles has not really been investigated and have chosen to focus on this aspect in our work. Using both shape and image information is likely to provide the best matching results, as image features may improve the precision of the edge matching measure and may be used in an initial grouping of the pieces to reduce the complexity when comparing pieces. We have, however, chosen to exclude entirely the shape aspect to show the strength of using image features and to prove that it is possible to solve jigsaw puzzles without shape information. This also generalizes the puzzle problem to include puzzles where an edge description is either not available or non-unique. In our specific case it allows us to solve jigsaw puzzles which only employ a small set of different shapes, as well as puzzles with rectangular pieces.

It should be noted that our edge matching can be expanded to include shape information, and we expect this will improve the performance on typical puzzles.

This article is based on our previous work (Nielsen et al., 2006) (in Danish) on the same subject. An important addition is the testing of our algorithm on real puzzle pieces.

1.1. Problem definition

The puzzle problem is the problem of assembling a jigsaw puzzle so that all pieces fit together forming a picture. We assume that the pieces each have four sides and are arranged in a rectangular grid. Each side of a piece can either be concave, convex or have a straight edge. Border pieces have one or two straight edges. Unlike previous definitions, however, we do not require non-straight sides to be unique.

To demonstrate the power of being independent of the uniqueness of the shape, we also expand our definition to encompass puzzles with perfectly rectangular pieces. Such puzzles have the additional property that they are easy to generate when we are

* Corresponding author. Tel.: +45 35 32 14 00; fax: +45 35 32 14 01.
E-mail address: khan@diku.dk (K. Hansen).

0167-8655/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2008.05.027

testing our methods. However, rectangular pieces have no distinctive edge, which makes it practically impossible to identify border pieces. For this reason, we instead require that border pieces be identifiable from image features, for instance by having an outline or frame in the image, or that a person identify them. The final tests of our algorithm will, however, be performed on classical pieces of both real and computer generated puzzles.

1.2. Related work

The first article concerning the puzzle problem is Freeman and Gardner (1964), who worked with puzzles without pictures. Despite not actually being implemented, their work laid the foundation for many subsequent papers. Their method attempts to identify critical points along the edge that can be used for segmentation. In this way a measure is calculated for how well the pieces fit together, also known as *partial boundary curve matching*. Radack and Badler (1982) also employed partial boundary curve matching, this time using polar coordinates. Wolfson et al. (1988) developed an algorithm for solving puzzles using two-dimensional *Schwartz–Sharir curve matching* and optimized search trees, and managed to solve puzzles of up to 104 scanned pieces.

Kosiba et al. (1994) proposed the first method using both the image on and the shape of the pieces, thus managing to correctly solve small puzzles as well as puzzles of up to 54 pieces with "satisfactory" results. The idea of using both image and shape was further employed by Chung et al. (1998), who also managed to solve puzzles of 54 pieces, and later by Yao and Shao (2003), who managed to solve puzzles with "dozens" of pieces.

To the best of our knowledge the largest puzzle solved to date by a computer has 200 pieces, and was solved using the method introduced by Goldberg et al. (2004). Their method employs only the shape of the pieces, and instead of a search tree they use a greedy approach, building the puzzle from the corners of the solved border. This heuristic solves the puzzle problem in polynomial time, thus greatly reducing the complexity.

1.3. Solution outline

The automatic assembly of a puzzle can be divided into five separate problems:

- preprocessing,
- edge matching,
- solving the border,
- solving the interior,
- laying the pieces and presenting the solution.

Preprocessing generally consists of scanning real puzzles, separating the individual pieces and then rotating them to achieve the proper alignment. There are several challenges to this process for real puzzles, which we discuss in Section 4.

Edge matching is performed by matching an edge of one piece to an edge of another piece, thereby obtaining the likelihood of the two pieces fitting together at those edges. This is described in Section 2, where we propose a new method solely using image features. This local match can then be used in assembling the border as well as the interior pieces.

Having devised a method of matching pieces, the *solution* to the puzzle problem lies in the ordering of the pieces to obtain the highest global match. This is a combinatorial problem whose complexity is $O(n!)$, when searching through the entire solution space. Due to the high complexity, the problem is often split into the separate problems of solving the border and solving the interior. Because the complexity of the interior puzzle remains the same, a heuristic is used to reduce the problem. In Section 3 we will briefly discuss a heuristic for solving the border as discovered by Wolfson et al. (1988), as well as a highly successful heuristic for solving the interior by Goldberg et al. (2004).

Finally, the puzzle solver presents the solution by *displaying the assembled pieces*, and this includes minimizing overlaps and holes through corrections to the orientation of the pieces. While some algorithms perform this step during the solution stage to further evaluate edge similarity, it can be considered a separate sub-problem and be performed after the topological solution is known. As this problem is not the focus of our work, we will use a simple algorithm, which only rotates pieces in increments of 90° (given by the topological solution) and does not try to minimize the error through corrections to the orientation.

1.4. A brief remark on the test data

We use two sets of puzzles for testing. Two images are used to create synthetic puzzles for testing edge matches. One image (`landscape`) has groups of relatively homogeneous pieces and the other (`construction`) has varying colors and textures. The images are shown in Fig. 1.

During our edge matching tests we divide the images into puzzles of rectangular and classical non-rectangular pieces ranging in numbers from 56 to 504. The pattern used for cutting these pieces is shown in Fig. 2, and is repeated vertically and horizontally, thus creating at most 16 unique shapes for interior pieces and at most 32 unique edges.

In addition, two real puzzles of 24 and 54 pieces have been scanned (see Section 4 for details). Fig. 10 shows part of the solution found by our algorithm for the 54-piece puzzle (`Benjamin`). The complete set of test puzzles may be found at the DIKU Image Group's FTP server (2008).

2. Edge matching

How well two puzzle pieces match along a common edge is expressed as a *similarity measure*. As mentioned above, there have only been a few experiments on using image features for matching pieces and none which rely only on this. Kosiba et al. (1994) and Chung et al. (1998) both sampled small areas at intervals along the edge using local image features such as mean, variance and histogram difference to calculate a similarity measure. Yao and Shao (2003) used the integration degree[1] on subdivided strips all along the edge for their solution.

Common for these methods is that shape matching is used as well. In this article, we propose a method that solely uses image features. The method is different from previous methods in that it only considers a single-pixel wide continuous strip for each edge as described in Sections 2.1 and 2.2. It is closely related to directional edge detection in that two pieces are considered a likely match if there is little or no gradient at the common edge between them.

The sampling method and the similarity measure are interconnected, and we first describe the similarity measure (Section 2.1), assuming we have derived straight image strips from each edge. In Section 2.2 we then define a way of sampling the edge of the pieces to generate these strips.

2.1. Similarity measure

Imagine two vertical image strips representing the edges of two pieces that are to be compared. By placing them side by side and

[1] The integration degree is a form of variance test, and is defined by Yao and Shao (2003) from the separability of image features, which in turn was described by Otsu (1979).