Automatic reconstruction of ancient Portuguese tile panels

Fernanda A. Andaló, Gustavo Carneiro, Gabriel Taubin, Siome Goldenstein, Luiz Velho

***Abstract*—Portuguese tile panels, or *azulejos*, are one of Portugal's cultural icons, and a representative cultural heritage of this country. Portugal's museums currently have a large collection of loose tiles that are often reassembled manually, which represents a laborious and challenging work. In this article, we explore the problem of automatically reconstructing ancient tile panels, mapping this problem to the reconstruction of an image from an unordered collection of rectangular non-overlapping tiles, an interest and important formulation of the jigsaw puzzle problem. Here we analyze, in a preliminary study, the application of image puzzle solvers in the assembling of ancient tile panels provided by the National Tile Museum. We compare the obtained results in different formulations of the problem, depending on the prior knowledge – known or unknown panel dimension and tile orientation –, and with missing tiles.**

***Keywords*—Portuguese tiles, Image puzzle, tile panel assembly.**

I. INTRODUCTION

Tile panels (or *azulejos* in Portuguese) is a form of Portuguese (and also Spanish) painted, tin-glazed, ceramic tilework that has become one of the most important forms of artistic representation in Portugal for the last five centuries.

The artists of azulejos were often inspired by or used to copy famous paintings or prints of those paintings [1]. As a result, thousands upon thousands of tiles were produced. Not only in Portugal, but also in other Portuguese colonies (e.g., Brazil), *azulejos* are still commonly found in the interior and exterior of churches, palaces, castles, houses, restaurants, and railway stations. These tile panels usually cover large sections of walls, floors, or even ceilings for decorative purposes. When one of these buildings needs to be renovated or demolished, these *azulejos* can also be destroyed unless they are carefully removed from the building structure.

The National Tile Museum (*Museu National do Azulejo*, MNAz), in Lisbon, Portugal, currently stores a large collection of tiles that have been removed from several buildings in Portugal. In order to study these tiles, several works have been devoted to their characterization, treatment, and conservation [2, 3]. In a work of heritage preservation, the MNAz tries to reassemble the panels back together, in a program called *Devolver ao Olhar* (Giving Back to the View) [4].

The challenges involved in reassembling tile panels are huge because some of them have been removed in such a way that mounting instructions are non-existing, there can be missing tiles, and information about the shape and size of the panels is not always available. Furthermore, a single box of loose tiles can actually contain several (incomplete) panels.

Figure 1 shows an art historian of MNAz working on the reassembling of a tile panel. The tasks involved in reassembling a tile panel comprise placing the tiles from a single origin on the floor and cataloging the position, orientation, and panel identity of each tile. These tasks usually involve the investigation of hundreds of tiles and can become quite time consuming, even for an expert in *azulejos*.

A tile panel assembling process can be seen as the assembly of a jigsaw puzzle, where each tile corresponds to one piece in the puzzle. However, the level of difficulty can be higher than in a usual puzzle because sometimes no information about the final appearance is available, which means that it may be necessary to determine the position and orientation of each tile. Another difficulty is that all the pieces of this puzzle are equal and roughly square, thus failing to provide any information about its orientation and neighboring tiles, except from color continuity. In addition, there is still the problem that many of the tiles are in an advanced state of degradation, making the pairing of them based on appearance a challenging task.

Similar cultural heritage problems have been studied previously, with two important examples being the Thera Frescoes project [6] and the Digital Forma Urbis Romae project [7]. The Thera Frescoes project aims at the automated digitization and matching of free-form fragments of wall paintings (frescoes) recovered from the archaeological site of Akrotiri on the island of Thera. It has three main components: acquisition, matching algorithms that compute candidate matches between fragments, and a user interface that allows users to evaluate the proposed matches. The Digital Forma Urbis Romae project aims at reconstructing the Severan Marble Plan of Rome, an enormous map, carved between 203–211 CE, that covered an entire wall inside the Templum Pacis in Rome. It employs digital technologies to try to reconstruct the map, creating digital photographs and 3D models of all $1,186$ fragments, and building a fully searchable database.

In this work we analyze the application of image puzzle solvers to the automatic reconstruction of Portuguese tile panels provided by the MNAz. These solvers address the problem of reconstructing images from rectangular non-overlapping puzzle pieces of identical shape and size. This type of application has been explored before [5], but with panels very limited in size. We extend it to other scenarios, yet in a preliminary study, by considering larger and mixed panels, and missing tiles.

The application of image puzzle solvers to the problem of

Fernanda A. Andaló is currently with the University of Campinas (Unicamp), but was with the Instituto Nacional de Matemática Pura e Aplicada (IMPA) at the time of the presented research.

Siome Goldenstein is with Unicamp.

Gustavo Carneiro is with the University of Adelaide.

Gabriel Taubin is with Brown University.

Luiz Velho is with IMPA.

Fig. 1. A person assembling a panel made of Portuguese tiles at the National Tile Museum (*Museu National do Azulejo*, MNAz), in Lisbon, Portugal. Image provided by [5].

reconstructing tile panels has the advantage of not requiring complicated pipelines and equipment. With a standard camera, one has to simply digitize the tiles. Each tile is automatically corrected by the adjustment of its shape and size and then, using all the prior knowledge available, a solver can reconstruct the entire panel or pieces of it. The tiles used in this work have not received any treatment due to their deterioration, yet it is possible to reconstruct the panels entirely or several parts of them.

In Section II we cover the literature on image puzzle solvers and in Section III we present the formulation of relevant solvers and their comparison with other methods. Section IV shows the application of such solvers to the task of reassembling Portuguese tile panels, and finally Section V concludes our work.

II. Background on Image Puzzle Solvers

The problem of automatically reconstructing an image from a collection of unordered non-overlapping pieces is computationally complex in a sense that no efficient algorithm is known capable of solving it in a deterministic manner when the compatibility between the pieces is uncertain, i.e., when it is not possible to determine the adjacent pieces without ambiguities [8].

The problem has an inherent difficulty revealed from its *global nature*. It is hard to construct an entire image puzzle when dealing only with local matches, because no exact measure of similarity between tiles is known to date. Methods have to rely on good initialization for local searches or explore the solution space in search of a good solution [9]. Moreover automatic puzzle solvers have to overcome the *combinatorial nature* of the problem, in which the number of possible solutions increases super-exponentially with the number of available pieces because, in the worst case, every possible permutation of the pieces can be a valid solution.

Generally image puzzle solvers are developed for two kinds of puzzle: *pictorial*, in which the correctly assembled pieces form an image, and *apictorial*, where there is no chromatic difference between the pieces and their distinct shapes, when assembled correctly, form a unique plane.

The first solver was proposed by Freeman and Garder [10] to solve 9-piece apictorial puzzles, and it is considered the basis to many subsequent works. Thirty years later, the method by Kosiba et al. [11] was the first to consider chromatic information, successfully assembling small pictorial puzzles with traditional pieces.

Besides the reconstruction of tile panels, puzzle solvers can generate solutions to other scientific problems: reassembling of broken archaeological artifacts [7, 12, 13, 14], reconstruction of shredded documents [15, 16], speech recognition [17], DNA/RNA modeling [18], image editing [19], among others.

In this work, we consider pictorial puzzles, but formed by identical rectangular pieces, or *tiles*. The literature for this kind of puzzle is somewhat recent [5, 20, 21, 22, 23, 24, 25, 26] and not every work considers the same *a priori* knowledge of the problem. Works by Cho et al. [20], Pomeranz et al. [21], Andaló et al. [23], Sholomon et al. [24] and Mondal et al. [25] consider that the puzzle dimension and the orientation of each tile are known, opposed to the works by Gallagher [22], Fonseca [5] and Son et al. [26]. All of them accept only square tiles, except the method by Andaló et al. [23] that can solve puzzles with arbitrary rectangular tiles, a useful characteristic when assembling shredded documents, for example.

Cho et al. [20] obtained an approximate reconstruction of the original image using graphical models and a global probabilistic function. However the method needs information about the layout of the original image, such as the correct location of some tiles informed by the user. Although being semi-automatic, this strategy allows the assembling of puzzles up to 432 tiles.

Pomeranz et al. [21] presented a method that does not need user intervention. It is based in a greedy approach, in which a compatibility function is computed to measure the affinity between the tiles, and then the method solves three problems:

positioning, segmentation, and translation. The positioning module put all tiles on the grid following a predetermined logic and considering randomly selected seeds; the segmentation module identifies the regions that are more likely to be assembled correctly; and the translation module reallocates regions and tiles to produce the final result. With this greedy strategy, they achieved the considerable improvement of solving puzzle with up to 3300 tiles.

Sholomon et al. [24] proposed a genetic algorithm to solve very large puzzles up to 22,834-puzzle pieces with known tile orientation and puzzle dimension.

Mondal et al. [25], instead of trying to find enhanced compatibility metrics across piece boundaries, combined existing techniques to achieve higher accuracy and robustness.

Son et al. [26] presented an algorithm based on “loop constraints” to solve puzzles with unknown panel dimension and orientation. The algorithm finds loops of pieces which form cycles and then aggregate these loops into higher order “loops of loops”.

The other works by Andaló et al. [23], Fonseca [5], and Gallagher [22] are described in more details in the next section. We evaluate these methods using a standard dataset of natural images, showing that they provide good results in comparison with the other methods, considering the tested metrics and puzzle dimensions.

III. Image puzzle solvers

In this section, we describe three image puzzle solvers. The solver by Andaló et al. [23, 27] can be applied to puzzles with arbitrary rectangular tiles, with known panel dimension and tile orientation. The solvers by Fonseca [5] and Gallagher [22] can be applied to puzzles with square tiles and unknown panel dimension and tile orientation. The following subsections briefly describe each solver. For a more detailed explanation, please refer to the original publications.

A. Method by Fonseca [5]

The work by Fonseca [5] was the first to apply the idea of image puzzles to panels of Portuguese tiles, although it was developed only for small panels.

The greedy method tries to minimize the distance between tile appearances at each iteration of the algorithm, as tiles are connected to the final solution. It begins by computing a Global Distance Matrix S, of size $4N \times 4N$, where N is the number of tiles, that encapsulates the distance between all tiles in every possible tile orientation. The lowest value is chosen and the corresponding tiles are put together in the final solution as neighbors.

At this point, there are six available borders in the solution, so that a new tile can be connected, and $N-2$ possible connections (tiles that have not been used yet). A $4N \times 4N$ mask is created and an element-wise product between the mask and S provides the minimum value corresponding to the best tile connection. The purpose of this mask is to disallow new connections with tiles that have already been used in the final solution.

This procedure is repeated until all tiles have been connected to the final solution.

To ensure a good quality result, an heuristic called *Lowe Scores* is also employed. When there are two tile candidates to be connected in the final solution, with close distance values according to a threshold, the connection is rejected. This heuristics suggests that a connection is not meaningful if the tile candidates have almost the same distance.

B. PSQP – Puzzle Solving by Quadratic Programming [23]

The method presented by Andaló et al. [23, 27], named PSQP (Puzzle Solving by Quadratic Programming), is based in maximizing a global matching function which calculates the overall compatibility of a certain tile permutation.

Consider an image partitioned into a regular 2D grid, forming N tiles of identical dimensions; and an empty grid of the same size as the previous one with N locations. The problem is to determine a one-to-one correspondence between the N tiles and the N locations, optimal with respect to a global compatibility function $\varepsilon(P)$ that sums up the compatibility of the neighboring tiles, considering P as a permutation matrix that assigns tiles to locations (Figure 2). Briefly, the compatibility between two tiles can be though of as a measure based on the color difference of the touching borders, when the tiles are considered as neighbors in a solution.

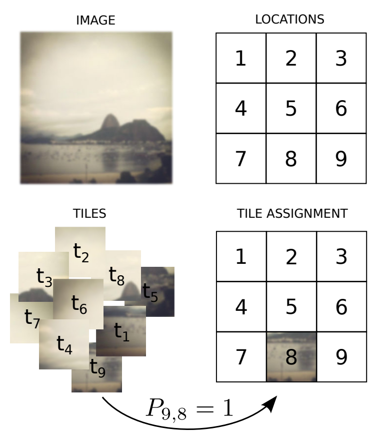

Fig. 2. PSQP problem formulation. Each tile t_i is assigned to a location j. The solution is represent by a permutation matrix P. Image adapted from [23].

The goal is to maximize $\varepsilon(P)$ over all permutation matrices P of size $N \times N$. Since this is a hard combinatorial optimization problem, it is necessary to extend the domain of the global compatibility function to the set of doubly stochastic matrices by relaxing the binary constraint of matrix P. The problem is then reformulated as a constrained continuous optimization problem, which can be solved by numerical methods.

The final global compatibility function to be maximized is

$$\begin{aligned}&\text{Maximize } f(p) = p^\top A p,\\ &\text{subject to } P\mathbb{1} = \mathbb{1}, P^\top \mathbb{1} = \mathbb{1}, \text{ and } p_{ij} \geq 0,\end{aligned} \tag{1}$$

where p is the column concatenation of P, A is the Hessian of $\varepsilon(P)$, $\mathbb{1}$ is a column vector of size N with all elements equal to one.

To search for the local maxima of the problem, which in practice is a permutation matrix representing a possible solution, the authors proposed a modified constrained gradient ascent algorithm, with gradient projection.

C. Method by Gallagher [22]

The work by Gallagher [22] was the first to introduce puzzles in which the orientation of the pieces is unknown. To solve this kind of puzzle, Gallagher proposed a new compatibility metric called the Mahalanobis Gradient Compatibility (MGC), that describes the local gradients near the boundary of a piece. This metric penalizes changes in intensity gradient, learns the covariance between the color channels and then uses the Mahalanobis distance.

The proposed method to assemble puzzles is inspired by the Minimum Spanning Tree (MST) algorithm for graphs. The problem is formulated as a graph where each piece is a vertex, and edge weights are the MGC computed for the corresponding pieces. The MST is the cheapest possible configuration that could be used to assemble the pieces into a single connected component, but some geometric constraints need to be applied so that the MST does not result in a puzzle that overlaps itself.

The proposed algorithm has three stages:

1) **Constrained tree:** the method applies a constrained version of the MST algorithm to find a tree in the input graph, according to the geometric constraints of the problem.
2) **Trimming:** if the resulting tree does not fit into a regular frame and the dimension of the puzzle is known, then the assembled tree is trimmed.
3) **Filling:** after trimming, the puzzle frame can have unoccupied holes. At this stage, holes are filled by order of the number of occupied adjacent neighbors and, for each hole, the candidate piece is the one with the minimum total dissimilarity score across all neighbors.

D. Image puzzle solvers comparison

In order to understand the applicability and accuracy of the detailed image puzzle solvers, we compare them in a standard dataset of images. This dataset is composed of twenty natural images provided by [20]. Each puzzle consists of 432 tiles of size 28×28 pixels.

The accuracy of the solutions are measured according to two different metrics previously proposed by Cho et al. [20] and Gallagher [22]:

Neighbor comparison: for each tile, this metric computes the fraction of its neighboring tiles that are also its neighbors in the correct solution. The accuracy is the mean fraction of correctly assigned neighbors.

Perfect reconstruction: binary indication of whether every tile is assigned to the correct location in a puzzle.

Note that directly comparing the resulting puzzle with the ground-truth image is not a good metric because it is unable to cope with slightly shifted solutions [24].

First we present the results for the methods that can work with known puzzle dimension and tile orientation [20, 21, 22, 23, 24, 25, 26].

Table I summarizes the mean accuracy for each method in the dataset of 20 images. Methods by Mondal et al. [25], which combine several previously proposed techniques, Son et al. [26], which employs “loop constraints”, and PSQP [23] attained the highest accuracy among all methods and also more perfect reconstructions. The method by Sholomon et al. [24], by employing a greedy method, has high accuracy but is not able to perfectly reconstruct many puzzles.

TABLE I. ACCURACY FOR EACH METHOD, CONSIDERING KNOWN PANEL DIMENSION AND TILE ORIENTATION.

Methods/Metrics	**Neighbor (%)**	**# Perfect**
Cho et al. [20]	55	0
Pomeranz et al. [21]	94	13
PSQP [23]	96	13
Gallagher [22]	95	12
Sholomon et al. [24]	96	7
Mondal et al. [25]	97	13
Son et al. [26]	96	13

As PSQP, Cho et al. [20] also employs a global approach, but by maximizing a probabilistic function via Loopy Belief Propagation. It needs some tiles to be fixed in their right position, however it is not able to perform perfect reconstructions. Methods by Pomeranz et al. [21], and Gallagher [22] have similar accuracy.

Figure 3 shows results comparing PSQP [23] with the method by Pomeranz et al. [21], and comparing PSQP [23] with the method by Sholomon et al. [24].

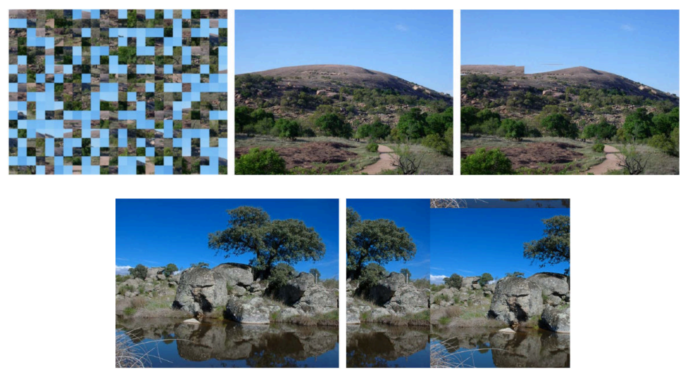

Fig. 3. Puzzle solvers applied to natural images, considering known panel dimension and tile orientation. Top: initial permutation, result obtained with PSQP, and result with the method by Pomeranz et al. [21]. Bottom: result obtained with PSQP (100% accurate), and result with the method by Sholomon et al. [24].

Considering another puzzle formulation, with unknown tile orientation, not all previously described methods can be applied, because their formulations assume that each tile is informed in its upright orientation. We compare three methods that allow this new formulation [5, 22, 26]. Note that these methods can solve puzzles with unknown panel dimension.

Table II summarizes the results. The method by Son et al. [26] has better accuracy because, in contrast to the other methods which avoid or ignore puzzle cycles, it exploits these loops as a form of outlier rejection. Nevertheless, all method attain good accuracy.

TABLE II. ACCURACY FOR EACH METHOD, CONSIDERING UNKNOWN PANEL DIMENSION AND TILE ORIENTATION.

Methods/Metrics	Neighbor (%)	# Perfect
Gallagher [22]	90.4	9
Fonseca [5]	92.6	9
Son et al. [26]	94.9	12

Fig. 4. Puzzle solvers applied to natural images, considering unknown panel dimension and tile orientation. From left to right: original image, result obtained with the method by Fonseca [5], and result with the method by Gallagher [22].

Figure 4 shows one of the obtained results comparing the method by Fonseca [5] with the method by Gallagher [22] with unknown panel dimension.

IV. ASSEMBLY OF ANCIENT PORTUGUESE TILE PANELS

Despite the previously discussed difficulties inherent to image puzzles, the assembling of Portuguese tiles poses a new set of challenges: the advanced degradation state of some tiles. Other factors interfere in the process of comparing the tiles: variations in colors, cracks along the borders, non-continuous strokes from tile to tile, texture created by typically thorough strokes, etc.

The experiments were conducted with tile panels provided by the MNAz. The tiles were acquired by a simple tool [5] that facilitates the acquisition process and the shape correction of the tiles. Each tile is captured by a camera and the corresponding image is automatically corrected: by using Hough transforms, the corners of the acquired tile are found; and then the homography is calculated and applied to correct the perspective and size to match a square (Figure 5).

Fig. 5. Tile acquisition and correction. Left: tile acquired by a camera. Right: corrected tile, considering perspective and size. Image provided by Fonseca [5].

To reconstruct the tile panels, we considered the three image puzzle solvers detailed in Section III. Except for the method by Fonseca [5], the other two solvers [22, 23] were chosen because of their good accuracy and available implementation. Results for the first solver are from [5]. Because of the discussed issues inherent to Portuguese tile panels, the accuracy of their reconstruction is expected to be lower than in the previous experiment, considering more unconstrained scenarios.

We considered three sets of panels provided by the MNAz:

- Twelve subsets of panels, with 25 tiles each;
- Four large panels, with 40, 48, 60 and 72 tiles each;
- Four mixes of 4 different tile panels, with 100 tiles each.

Experiments with each of the panel sets are described in the next subsections.

A. *Experiment with small panels*

In this first experiment, we used twelve subsets of panels provided by the MNAz, with 25 tiles each.

We consider unknown panel dimension and three conditioning scenarios: known and unknown tile orientation, and missing tiles.

Known tile orientation

In this scenario the three methods [5, 22, 23] can be used. Note that, because the panels are small (25 tiles each), PSQP can be applied to the possible three configurations and the one that yields the highest global compatibility is chosen as the solution.

Table III summarizes the results. PSQP was able to reconstruct all the panels with 100% accuracy. Differently from PSQP, the other two methods assemble the panel taking into account the unknown dimensions and, for this reason, their lower accuracy is expected.

Unknown tile orientation

In this case only the methods by Fonseca [5] and Gallagher [22] can be applied, since PSQP needs the correct tile orientations to provide a solution.

Table III summarizes the results. Although the accuracy in this scenario is low, note that the method by Gallagher [22] is still able to correctly assign half of the neighboring tiles. In a real scenario, these results could help the restorers to assemble the entire panel.

TABLE III. ACCURACY FOR EACH METHOD, CONSIDERING UNKNOWN PANEL.

Methods/Metrics	Tile orientation	Neighbor (%)	# Perfect
PSQP [23]	Known	100.0	12
Gallagher [22]	Known	64.5	4
Fonseca [5]	Known	57.8	0
Gallagher [22]	Unknown	49.4	3
Fonseca [5]	Unknown	35.9	0

Missing tiles

To test the methods in solving panels with missing tiles, we conduct the same experiments but removing up to 30% of the tiles, and considering that the total number of tiles in the original panel is known. Figure 6 shows the resulting accuracy for growing quantities of missing tiles, and Figure 7 illustrates some of the results.

Although PSQP provides better results, it is more affected by the growing number of missing tiles. Considering the scenario where the orientation of the tiles is

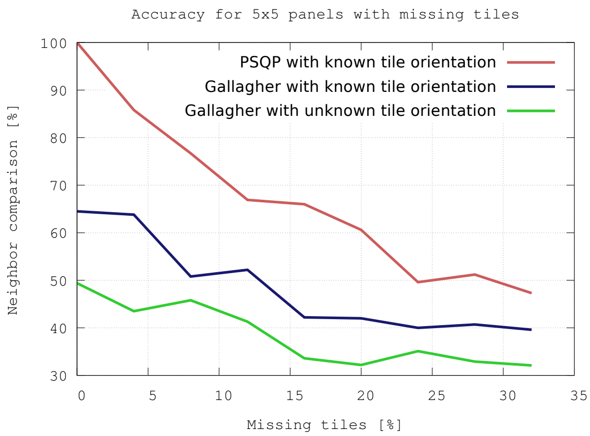

Fig. 6. Accuracy for 5 × 5 panels, when considering a growing number of missing tiles, and unknown panel dimension. Although PSQP is more accurate, it is more affected by the growing number of missing tiles.

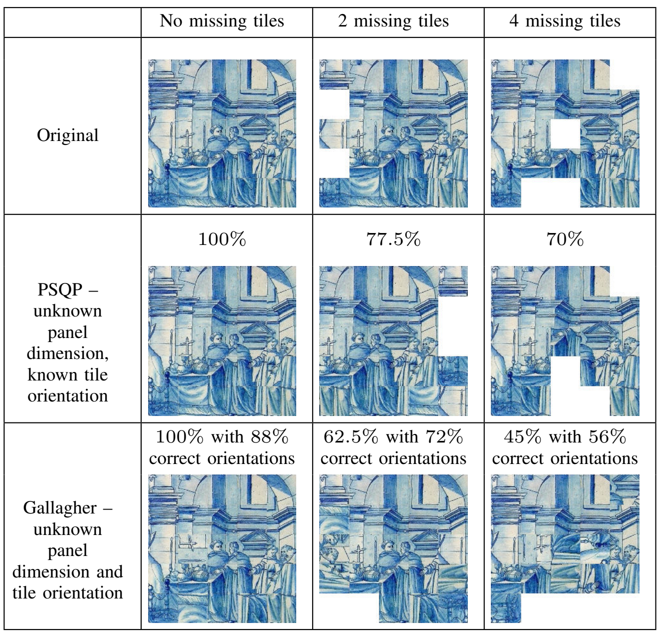

Fig. 7. Results with a growing number of missing tiles (none, two, and four missing tiles), using methods PSQP and by Gallagher [22], in different scenarios.

not known, the method by [22] can achieve the mean accuracy of 32% with 30% of missing tiles.

The observable errors are generated mainly by the metrics used to compare the tiles. Puzzle solvers usually consider characteristics measured at the border of the tiles, like color and gradient. It is essential that these properties have continuity between the tiles, but this is not always the case, specially when there are several missing tiles. Methods that explore the solution space, as PSQP, tend to provide better results in such cases, but it is not trivial to extend them to unconstrained scenarios. It is also important to note that not so highly accurate results, such as the ones achieved with the method by Gallagher [22] with unknown panel dimension and tile orientation, are also important to aid the reconstruction of the entire panel, as can be observed in Figure 7. In such cases, the overall appearance is still captured in the resulting panel.

B. Experiment with larger panels

Considering the large panels provided by the MNAz, with 40, 48, 60, and 72 tiles each, we could only experiment with PSQP and the method by Gallagher [22], as the method by Fonseca [5] was developed for small panels.

Table IV summarizes the results with known and unknown tile orientation and some examples are shown in Figure 8.

To test the larger panels with missing tiles, we computed the accuracy of the methods after removing up to 30% of the tiles. Figure 9 shows the resulting accuracy for growing quantities of missing tiles

TABLE IV. ACCURACY FOR EACH METHOD, CONSIDERING UNKNOWN PANEL DIMENSION.

Methods/Metrics	**Tile orientation**	**Neighbor (%)**	**# Perfect**
PSQP [23]	Known	96.7	3
Gallagher [22]	Known	51.7	1
Gallagher [22]	Unknown	46.6	1

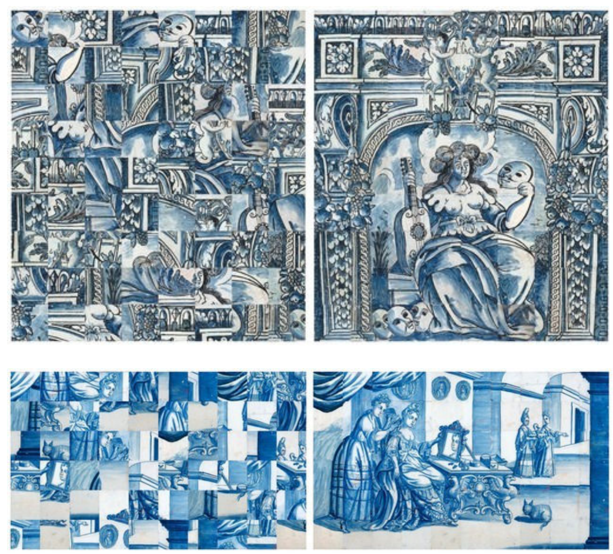

Fig. 8. PSQP applied to larger tile panels. Top: Initial permutation and result for a 72-tile panel. Bottom: Initial permutation and result for a 60-tile panel.

One can note that the methods kept the same behavior observed in smaller panels, despite being less accurate when missing tiles are considered.

C. Experiment with mixed panels

In this third experiment, we mixed the tiles of four panels together, resulting in four mixed panels with 100 tiles each.

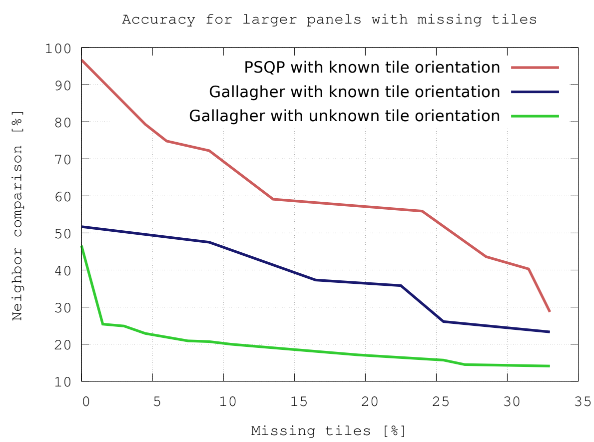

Fig. 9. Accuracy for larger panels, when considering a growing number of missing tiles, and unknown panel dimension. Methods kept the same behavior observed in smaller panels, despite being less accurate when missing tiles are considered.

Considering known tile orientation and unknown panel dimension, PSQP was able to solve them with 100% accuracy. And considering unknown tile orientation, the method by Gallagher [22] solved the panels with the mean accuracy of 43%. Figures 10 shows one of these results.

Fig. 10. Solving 4 mixed tile panels. The first panel was constructed by PSQP, with known tile orientation, and the second one by Gallagher [22], with unknown tile orientation.

A restorer could separate the mixed tiles prior to the assembly, taking into account, for instance, the appearance, color, and subject of the tiles. However, when these characteristics are similar, the automatic solver can separate them better and faster. The mixed panels were assembled in less than 10 seconds each.

D. Discussion

There is a clear trade-off to choose which method should be considered in a real scenario, and it depends manly on the prior knowledge of the panel. If more information is given, then higher accuracy is achieved. It is sometimes possible to have cues such as the panel dimension. For example, as stated before, several panels were produced as copies of original paintings. If these paintings are known *a priori*, then other aspects can be derived.

Concerning the efficiency of the methods, it is important to note that, although PSQP yields the best accuracy in the more restricted scenario, its global and deterministic nature hinders its use for really large panels, when compared to the other approximate methods. Therefore, there is also a trade-off between effectiveness and efficiency that must be considered, depending on the size of the panel. For the tested panel sizes, the efficiency of the methods is similar.

V. Conclusion

In this paper, we have studied three image puzzle solvers when applied to the reconstruction of ancient tile panels, to help preserving the cultural heritage of the *azulejo* (Portuguese tile). There are advantages in the use of image puzzle solvers to reconstruct such panels: no other automatic method exists to date, leaving the laborious task completely to the restorers; the presented algorithms are efficient; and only simple equipment is required.

Preliminary experimental results showed that PSQP [23] is promising in reconstructing panels when the tile orientations are known. Nevertheless, it is important to extend its application to panels with unknown dimensions, because as more tiles are considered, it is impossible to test every configuration.

The method by Gallagher [22] can be applied when no *a priori* information about the panels is available. The application of such image puzzle solvers can aid restorers in reconstructing several parts of the panels (for instance, half of neighboring tiles in the panel), and visualizing the big picture.

The problem of panels with missing tiles has not been incorporated in the formulation of any method in the literature. Nevertheless, we applied the standard solvers in panels with missing tiles and the results were promising.

In a future work, we will consider the study of several characteristics of the tiles, such as color pallete, stroke style, and material, that can aid the separation of mixed tiles prior to the overall assembly.

References

[1] N. da Silva, M. Marques, G. Carneiro, and J. Costeira, “Explaining scene composition using kinematic chains of humans: application to Portuguese tiles history,” in *Conference on Computer Vision and Analysis of Images of Art II*, 2011.

[2] T. Silva, M.-O. Figueiredo, M.-A. Barreiros, and M.-I. Prudêncio, “Decorative 18th century blue-and-white Portuguese tile panels: a type-case of environmental degradation,” *Journal of Materials*, vol. 2013, 2013.

[3] J. Pessoa, J. Antunes, M. Figueiredo, and M. Fortes, “Removal and analysis of soluble salts from ancient tiles,” *Studies in conservation*, vol. 41, no. 3, pp. 153–160, 1996.

[4] R. Temática, “A blind person’s cure – an unknown tile panel from the collection of the National Tile Museum,” http://redeazulejo.fl.ul.pt/redeazulejaria_en/noticias_arquivo.aspx?nid=522&op=detalhes&menu=290&id=332, 2011, accessed: September 2014.

[5] J. Fonseca, "Montagem automática de painéis de azulejos," Master's thesis, Instituto Superior Técnico, Universidade Técnica de Lisboa, 2012.
[6] H. Shin, C. Doumas, T. Funkhouser, S. Rusinkiewicz, K. Steiglitz, A. Vlachopoulos, and T. Weyrich, "Analyzing and simulating fracture patterns of Theran wall paintings," *Journal on Computing and Cultural Heritage*, vol. 5, no. 3, pp. 10:1–10:14, 2012.
[7] D. Koller, J. Trimble, T. Najbjerg, N. Gelfand, and M. Levoy, "Fragments of the city: Stanford's digital Forma Urbis Romae project," in *Proceedings of the Third Williams Symposium on Classical Architecture, Journal of Roman Archaeology Suppl. 61*, 2006, pp. 237–252.
[8] E. Demaine and M. Demaine, "Jigsaw puzzles, edge matching, and polyomino packing: connections and complexity," *Graphs and Combinatorics*, vol. 23, no. 1, pp. 195–208, 2007.
[9] O. van Kaick, H. Zhang, G. Hamarnehn, and D. Cohen-Or, "A survey on shape correspondence," *Computer Graphics Forum*, vol. 30, no. 6, pp. 1681–1707, 2011.
[10] H. Freeman and L. Garder, "Apictorial jigsaw puzzles: The computer solution of a problem in pattern recognition," *IEEE Transactions on Electronic Computers*, no. 2, pp. 118–127, 1964.
[11] D. Kosiba, P. Devaux, S. Balasubramanian, T. Gandhi, and K. Kasturi, "An automatic jigsaw puzzle solver," in *12th International Conference on Pattern Recognition (IAPR)*, vol. 1, 1994, pp. 616–618.
[12] J. McBride and B. Kimia, "Archaeological fragment reconstruction using curve-matching," in *Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2003.
[13] B. Brown, C. Toler-Franklin, D. Nehab, M. Burns, D. Dobkin, A. Vlachopoulos, C. Doumas, S. Rusinkiewicz, and T. Weyrich, "A system for high-volume acquisition and matching of fresco fragments: Reassembling Theran wall paintings," *ACM Transactions on Graphics (Proceeding of SIGGRAPH 2008)*, vol. 27, no. 3, pp. 84:1–84:9, 2008.
[14] Q.-X. Huang, S. Flöry, N. Gelfand, M. Hofer, and H. Pottmann, "Reassembling fractured objects by geometric matching," *ACM Transactions on Graphics (TOG)*, vol. 25, no. 3, pp. 569–578, 2006.
[15] E. Justino, L. Oliveira, and C. Freitas, "Reconstructing shredded documents through feature matching," *Forensic Science International*, vol. 160, no. 2, pp. 140–147, 2006.
[16] L. Zhu, Z. Zhou, and D. Hu, "Globally consistent reconstruction of ripped-up documents," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 1, pp. 1–13, 2008.
[17] Y.-X. Zhao, M.-C. Su, Z.-L. Chou, and J. Lee, "A puzzle solver and its application in speech descrambling," in *Proceedings of the International Conference on Computer Engineering and Applications*, 2007, pp. 171–176.
[18] W. Marande and G. Burger, "Mitochondrial DNA as a genomic jigsaw puzzle," *Science*, vol. 318, no. 5849, pp. 415–415, 2007.
[19] T. Cho, S. Avidan, and W. Freeman, "The patch transform," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 8, pp. 1489—1501, 2010.
[20] ——, "A probabilistic image jigsaw puzzle solver," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 183–190.
[21] D. Pomeranz, M. Shemesh, and O. Ben-Shahar, "A fully automated greedy square jigsaw puzzle solver," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 9–16.
[22] A. Gallagher, "Jigsaw puzzles with pieces of unknown orientation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 382–389.
[23] F. Andaló, G. Taubin, and S. Goldenstein, "PSQP: Puzzle solving by quadratic programming," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to appear.
[24] D. Sholomon, O. David, and N. Nathan, "A genetic algorithm-based solver for very large jigsaw puzzles," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 1767–1774.
[25] D. Mondal, Y. Wang, and S. Durocher, "Robust solvers for square jigsaw puzzles," in *International Conference on Computer and Robot Vision (CRV)*, 2013, pp. 249–256.
[26] K. Son, J. Hays, and D. Cooper, "Solving square jigsaw puzzles with loop constraints," in *European Conference on Computer Vision (ECCV)*, 2014, pp. 32–46.
[27] F. Andaló, G. Taubin, and S. Goldenstein, "Solving image puzzles with a simple quadratic programming formulation," in *XXV Conference on Graphics, Patterns and Images (SIBGRAPI)*, 2012.