

A Novel Hybrid Scheme Using Genetic Algorithms and Deep Learning for the Reconstruction of Portuguese Tile Panels

Daniel Rika
Dept. of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
danielrika@gmail.com

Eli (Omid) David
Dept. of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
mail@elidavid.com

Dror Sholomon
Dept. of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
dror.sholomon@gmail.com

Nathan S. Netanyahu*
Dept. of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
nathan@cs.biu.ac.il

ABSTRACT

This paper presents a novel scheme, based on a unique combination of *genetic algorithms* (GAs) and *deep learning* (DL), for the automatic reconstruction of *Portuguese tile panels*, a challenging real-world variant of the *jigsaw puzzle problem* (JPP) with important national heritage implications. Specifically, we introduce an enhanced GA-based puzzle solver, whose integration with a novel DL-based *compatibility measure* (DLCM) yields state-of-the-art performance, regarding the above application. Current compatibility measures consider typically (the chromatic information of) edge pixels (between adjacent tiles), and help achieve high accuracy for the synthetic JPP variant. However, such measures exhibit rather poor performance when applied to the Portuguese tile panels, which are susceptible to various real-world effects, e.g., monochromatic panels, non-squared tiles, edge degradation, etc. To overcome such difficulties, we have developed a novel DLCM to extract high-level texture/color statistics from the entire tile information.

Integrating this measure with our enhanced GA-based puzzle solver, we have demonstrated, for the first time, how to deal most effectively with large-scale real-world problems, such as the Portuguese tile problem. Specifically, we have achieved 82% accuracy for the reconstruction of Portuguese tile panels with unknown piece rotation and puzzle dimension (compared to merely 3.5% average accuracy achieved by the best method known for solving this problem variant). The proposed method outperforms even human experts in several cases, correcting their mistakes in the manual tile assembly.

*Also affiliated with the Gonda Brain Research Center at Bar-Ilan University, and the Center for Automation Research, University of Maryland, College Park, MD 20742.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '19, July 13–17, 2019, Prague, Czech Republic

© 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-6111-8/19/07...\$15.00
<https://doi.org/10.1145/3321707.3321821>

KEYWORDS

Real World Applications, Genetic Algorithm, Reconstruction, Deep Learning, Jigsaw Puzzle Problem, Portuguese Tile Panels

ACM Reference Format:

Daniel Rika, Dror Sholomon, Eli (Omid) David, and Nathan S. Netanyahu. 2019. A Novel Hybrid Scheme Using Genetic Algorithms and Deep Learning for the Reconstruction of Portuguese Tile Panels. In *Genetic and Evolutionary Computation Conference (GECCO '19)*, July 13–17, 2019, Prague, Czech Republic. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3321707.3321821>

1 INTRODUCTION

Object reconstruction from numerous fragments is a pervasive, important task that has been encountered in many areas throughout human civilization. Piecing together broken pottery, ancient frescoes, or shredded documents from their artifacts are merely a few examples. The most basic generic version of the problem is to assemble an object from its n different (non-overlapping) pieces as accurately and efficiently as possible. (To automate this challenging task, processing is often applied to colored images acquired from these pieces.) The basic problem definition is very similar to the popular *jigsaw puzzle problem* (JPP), which is known to be NP complete [1, 11]. The JPP has been pursued by many researchers, as it is a special instance of a broad class of challenging real-world problems, such as image editing [6], the recovery of shredded documents or photographs [10, 19, 24, 27], art conservation [2, 4, 21], speech descrambling [7, 51], etc., as well as additional problems in areas like biology [26], chemistry [46], literature [28], and more. Obviously, there are notable differences, in practice, between a pure JPP setting and the above real-world problems (e.g. unknown dimensions, missing pieces, gaps between pieces due to degradation over time, pieces from multiple puzzles, etc.). Nevertheless, the JPP serves as a testbed for developing ground-breaking methods for these important challenges.

Every reconstruction procedure requires a *compatibility measure* to estimate the likelihood that two given pieces are adjacent and a strategy for placing the pieces as “accurately” as possible with respect to some global objective function. Although much effort has




Figure 1: Reconstruction of Portuguese tile panels with unknown piece orientation and panel dimensions, due to our proposed system. Left: Input images of Portuguese tile panels, containing 256 (top) and 150 (bottom) pieces. Right: Perfectly reconstructed images due to our novel compatibility measure coupled with the enhanced version of a "kernel-growth" GA.

been devoted to devising reliable compatibility measures for jigsaw-like problems, they may not always be consistent¹; if they were, the problem would not be NP-hard. More importantly, the typical dependence of current compatibility measures on correlations between low-level color/texture statistics in the proximity of tile boundaries, renders jigsaw puzzle solvers based on such measures virtually ineffective for real-world problems, such as the reconstruction of archaeological fragments and shredded documents (where often the information is severely degraded near the points of fraction), or that of Portuguese tile panels, whose image content is not necessarily color-rich and where chromatic information near tile boundaries might be severely corrupted. In addition, many methods for solving optimally the piece placement problem resort to greedy strategies, which are problematic in encountering local optima. Moreover, they usually cannot recover from erroneous placements made early on (as a result of a greedy, locally optimal choice). To meet these challenges, we employ in this paper a *computational intelligence* (CI)

¹In the sense that the most compatible piece to a given piece A , with respect to a compatibility measure in question, may not necessarily be adjacent to A in the "correct" puzzle configuration.

approach in dealing effectively with *both* components of the problem (*i.e.* the search and the compatibility measure). Specifically, we present a **unique combination** of: (1) An enhanced *genetic algorithm* (GA)-based scheme for finding promising (partial) solutions (*i.e.* *fittest chromosomes*), at each iterative stage, as a strategy for optimal piece placement, and (2) a novel *deep learning* (DL) model for learning piece compatibility by directly training on the raw data (of a fairly small training set), without applying any standard feature selection/extraction techniques,

Our contributions are summarized as follows:

- (1) Provided an enhanced GA solver for the construction of Portuguese tile panels;
- (2) Obtained for the first time a DL-based compatibility measure (DLCM) for a real-world JPP-like task;
- (3) Presented a unique combination of the above GA module and the novel compatibility measure for the reconstruction of Portuguese tile panels on a large-scale basis (see *e.g.* Fig. 1);
- (4) Obtained state-of-the-art-results for the above real world problem; specifically, achieved an average accuracy of 82% on Type 2 puzzles with unknown dimensions (compared to merely 3.5% average accuracy achieved by Gallagher's method [16], which is the best method known for solving this problem variant);
- (5) Compiled a new benchmark for the community, regarding training and test data for the Portuguese tile problem.

The paper is organized as follows. Section 2 provides a brief survey of recent related work. Section 3 and Section 4 describe, respectively, our novel GA-based solver and the DL method for learning a compatibility measure. Section 5 presents the datasets used, and Section 6 provides detailed experimental results. Section 7 makes concluding remarks.

2 RELATED WORK

2.1 Synthetic JPP

2.1.1 Traditional Methods. Freeman and Garder [15] introduced initially in 1964 a computational solver, which handled up to nine-piece puzzles. Subsequent research [17, 22, 35, 48] relied solely on shape cues of the pieces. Kosiba *et al.* [23] were the first to use image content, in addition to boundary shape; their method computes color compatibility along the matching contour, rewarding adjacent jigsaw pieces with similar colors. This trend continued for more than a decade (see, *e.g.* [8, 25, 29, 37, 50]), before the research focus shifted from shape-based to merely color-based solvers of square-tile puzzles with known piece orientation (*i.e.* Type 1 puzzles).

Cho *et al.* [5] used *dissimilarity* (*i.e.* the sum, over all neighboring pixels, of squared color differences over all color bands), as a compatibility measure for their probabilistic puzzle solver, that handles up to 432 pieces, given some a priori knowledge of the puzzle. (The sum of squared differences is referred to as SSD.) Their 2010 paper was followed by Yang *et al.* [49], who reported improved performance due to their *particle filter*-based solver. Shortly after, Pomeranz *et al.* [34] presented, for the first time, a fully-automated jigsaw puzzle solver of puzzles containing up to 3,000 square pieces, using the above defined dissimilarity and their so-called *best-buddies* heuristic. Gallagher [16] advanced further the state-of-the-art by considering a more general variant of the problem, where a piece orientation

is unknown (*i.e.* Type 2 puzzle), as well as the puzzle dimensions. Specifically, he presented the preferable measure of *Mahalanobis gradient compatibility* (MGC), which penalizes changes in intensity gradients (rather than changes in intensity) and learns the covariance of the color channels, using the Mahalanobis distance. He suggested also *dissimilarity ratios* for a more indicative compatibility measure.

Sholomon *et al.* [40–42] pursued a GA-based approach based on a number of innovative *crossover* procedures, and demonstrated the effective performance of their methodology on very large Type 1 and Type 2 puzzles (including two-sided puzzles and a number of mixed puzzles). Son *et al.* [45] imposed so-called *loop constraints*, where the dissimilarity ratio (with respect to the smallest distance from a piece edge in question), for each consecutive pair of pieces along a loop of four or more pieces, is below a certain threshold. They were able to improve the accuracy for both Type 1 and Type 2 puzzles in certain cases. Also, they provided, for the first time, an upper bound on the reconstruction accuracy for various datasets. Paikin and Tal [31] proposed a greedy solver based on an asymmetric L_1 -norm dissimilarity and the best-buddies heuristic. They demonstrated how to handle, among other things, puzzles with missing pieces, and reported improved accuracy results and fast running times. More recently, Andaló *et al.* [3] showed how to map the JPP to the problem of maximizing a constrained quadratic function, and presented a deterministic algorithm for solving it via gradient ascent.

2.1.2 DL Methods. Recently, there have been also a few DL works related to the JPP [12, 13, 30, 38]. However, these works barely provide any practical solutions to even “toy instances” of the JPP, and their main thrust is to “re-purpose” a neural network, trained to solve a simple jigsaw puzzle (without manual labeling), to handle advanced tasks, such as object detection and classification, in an unsupervised manner. Other than the above, a DL-based heuristic called *DNN-buddies* was presented in [43], in an attempt to enhance the accuracy of a GA-based solver. It should be noted, though, that the above heuristic is employed in conjunction with the SSD measure, in a rather restrictive manner, so it is expected to perform rather poorly on real-world JPP-like tasks.

2.2 Real-World Portuguese Tile Panels

The reconstruction of ancient frescoes and wall paintings from numerous large repositories of fragmented artifacts, compiled over time due to natural deterioration, is of utmost importance in preserving world cultural heritage. Various efforts to automate the process (*e.g.* [4, 33, 44]) rely primarily on shape matching (in 2D and 3D) of fragments followed by their assembly. While exhibiting good performance on relatively small datasets (only a few hundred fragments), the scalability of these efforts (in terms of the number of fragments and the number of art works in a given pool) is questionable.

Our focus in this paper is on the reconstruction of the Portuguese tile panels [9], which concerns the assembly of ancient panels of 2D square tiles that have been removed from many buildings and landmarks in Portugal (see Figure 2). Currently, over one hundred thousand such tiles are stored at the Portuguese National Tile Museum (Museu Nacional do Azulejo) in Lisbon, and are awaiting manual assembly by human experts. In view of the extremely challenging nature of the problem, it would take decades, at the current

pace, before all these “jigsaw puzzles” are solved, *i.e.* before the panels are assembled by the human experts [32].

Fonseca [14] acquired tile images and adapted their shape to squares; he then applied an augmented Lagrange multipliers technique to an equivalent optimization problem and a greedy approach for Type 1 and Type 2 variants, respectively. He obtained 57.8% and 39.1% accuracy for these cases, respectively, on panels containing only a few dozen tiles. In comparison, Gallagher’s method [16] achieves corresponding accuracy levels of 64.5% and 49.4%. Andaló *et al.* [2] reported perfect reconstruction (of 4 mixed tile panels) using their PSQP method [3] for known tile orientation. However, their method does not handle the Type 2 variant, and its preliminary results were obtained for panels containing a fairly small number of, presumably, high-resolution tiles.



Figure 2: Manual assembling of a panel of Portuguese tiles at the National Tile Museum (Museu Nacional do Azulejo, MNAz), Lisbon, Portugal: Source [14].

3 GA SOLVER

We seek a global optimizer that can exploit the relative accurate piece adjacency prediction capability, but that can also overcome its inaccuracies. Previous solvers rely typically on some specialized criterion, which implies a subset of edge adjacencies that are likely to be correct. To avoid searching for such a specific criterion, we pursue a GA approach [18] for tile placement, in the spirit of the *kernel-growth* scheme presented in Sholomon *et al.* [40, 41]. Since the proposed GA solver is of a random nature, it could correct, potentially, wrong adjacencies during the global optimization.

Following [40], we describe here the new hierarchical phases of our modified *crossover* operator. In a nutshell, a chromosome is associated with a puzzle configuration (or a “solution”), and its fitness function is defined by the overall sum of pairwise, adjacent tile compatibilities (see below). The principle of hierarchical phases is that a piece is added to the growing kernel at each phase only if the previous phases have been exhausted (*i.e.* no further pieces can be added due to these phases); the crossover terminates once the kernel contains all the pieces. Our proposed phases and their hierarchical arrangement are as follows.

- **Phase I:** If there is a free (piece) boundary in the kernel, which has a neighboring piece in a chromosome parent, such that the score of each of these adjacent pieces is greater than $\max(0.8, C_{\text{mean}})$, where C_{mean} is the chromosome’s average compatibility across all boundaries, then add the neighboring

piece to the kernel. We define the score of a piece as the average compatibility measure between the piece and all of its neighbors. This phase gives priority to the chromosome parent with the higher fitness, assuming that it would yield a more accurate reconstruction rate.

- **Phase II:** Similar to Phase I, except that this phase selects the chromosome parent with the lower fitness.
- **Phase III:** If there is a free (piece) boundary in the kernel, such that the two chromosome parents agree on the adjacent piece, place this piece next to the boundary in question.
- **Phase IV:** If there is a free (piece) boundary in the kernel, such that its most compatible piece is available (*i.e.* is not placed already in the kernel), then add that compatible piece to the kernel.
- **Phase V:** If there is a free (piece) boundary in the kernel, such that its second-most compatible piece is available, then add the latter piece to the kernel.
- **Phase VI:** Pick randomly one of the remaining pieces, and place it randomly at one of the free boundaries of the kernel.

We introduce a certain degree of randomness to the process (known as *mutation*), in order to avoid local maxima, by skipping some of the crossover phases, with small probability. Specifically, we skip the first and second phases with 10% probability and the third phase with 20% probability. The other phases are always executed.

Other hyper-parameters of our modified GA solver (which were arrived at after exhaustive experimentation) are as follows: Chromosomes are chosen for the crossover operation according to the *roulette wheel selection*, the population consists of 100 chromosomes, and the GA runs for 500 generations.

3.1 Rationale

Before explaining the rationale behind the above phases, we note that our proposed crossover does not draw on the notion of *best-buddies*, as was defined and used *e.g.* in [34] and [40]. The reason for that is that in contrast to the (synthetic) JPP, where best-buddy pairs were found to be adjacent with 95% probability, our experience with the Portuguese tile panels shows that best-buddy pairs, with respect to our state-of-the-art DLCM (described in the next section), are correct with only 70% probability.

Regarding our modified crossover operator, note that the objective of the first and second phase is to inherit correctly-reconstructed segments from the parents. We constrain the score of each of the two pieces in question to be at least 0.8, as a good starting threshold. (Note that the score is in the range between 0 to 1, due to the normalization of the compatibility measure as explained in Sec. 4.3.) Furthermore, since the algorithm improves as the number of generations goes up, (*i.e.* chromosome fitness increases), the resulting threshold becomes greater than 0.8. The idea behind the dynamic threshold, is to overcome errors made in previous generations. Phase III is carried out if the two chromosomes agree on the same pair of pieces, *i.e.* they are likely to be correct, with high probability.

In the first three phases the crossover inherits adjacent pieces from the parents; however, these phases might not necessarily result in a successful addition of a new piece to the kernel. Thus, Phases IV and V, which rely solely on our proposed DLCM, could be used

alternatively by considering the most compatible and second-most compatible pieces.

If Phases IV and V still fail to add one more piece to the kernel (because the pieces considered are already placed in the kernel), Phase VI is invoked to complete the puzzle configuration, by placing randomly a free piece at an open boundary.

4 TRAINING A COMPATIBILITY MEASURE

We have striven to develop a DL model for learning automatically a compatibility measure, such that given two puzzle pieces, it would distinguish between adjacent and non-adjacent pieces. The proposed method is based loosely on ideas from the field of *metric embedding learning*. The goal of metric embedding learning is to learn a function $f_{\theta}(x) : \mathbb{R}^F \rightarrow \mathbb{R}^D$, which maps semantically similar points from the data manifold \mathbb{R}^F onto metrically close points in \mathbb{R}^D . This approach was first presented by Weinberger and Saul [47], in the context of nearest-neighbor classification. Schroff *et al.* [39] subsequently proposed using a deep *convolutional neural network* (CNN)-based embedding of human faces, which is trained via a so-called *triplet-loss* described below.

We propose to formulate the problem of learning a compatibility measure as learning a single-dimensional embedding $\mathbb{E} \times \mathbb{E} \rightarrow \mathbb{R}$, where \mathbb{E} is the group of all puzzle piece edges. Here we want to ensure that given a piece-edge e_i (*anchor*) and its adjacent piece-edge e_j in the original image, the score of the *positive pair* (e_i, e_j) will be higher than any *negative pair* (e_i, e_k). This can be achieved by minimizing the loss

$$L = \sum_{e_i, e_j, e_k \in T} \max(0, 1 - f(e_i, e_j) + f(e_i, e_k)),$$

where f is a deep convolutional neural network and T is the training set.

4.1 Triplet Selection

Since the number of possible triplets in the training set is quite large, we generated the training triplets online. Specifically, we selected, for each epoch, 25 pieces at random from every puzzle in our dataset. We used the edges of each piece as anchors, generating positive pairs from each edge and its neighboring pieces. (Usually this results in four pairs, but could also result in three or two pairs only, for pieces along the puzzle boundaries and the four corner pieces, respectively.) For each such positive pair, we randomly select a non-adjacent piece edge and create its accompanying negative pair to form a triplet.

Next, we randomly augmented each piece in each pair, using either *degradation* or *shifting*. Degradation replaces randomly the outermost pixel frame of the piece with zeros. With uniform probability, we may replace no pixel, replace a pixel-wide frame, or replace a double pixel frame. This should aid the network in learning more than only near-border textures. For shifting, we randomly shift the piece anywhere between zero to two pixels horizontally or vertically (filling with zeros empty locations). Figure 3 demonstrates some possible outcomes.

4.2 Deep Convolutional Neural Networks

We trained a deep *convolutional neural network* (CNN), which receives as input a pair of puzzle pieces and returns a real number



Figure 3: Illustration of tile augmentation via degradation and shifting. From left to right: Degraded tile by removing 2 pixels from its outer frame; shifted tile by one pixel to the left, and one pixel up; augmented tile with degradation and shifting.

score. All pieces are of size 50×50 pixels. Although most actual puzzle pieces are larger, we downscaled them to better fit in memory and speed up the training phase. Always taking the anchor piece to be on the left, we rotate the pieces accordingly. For example, to compare the left edge of an anchor piece with the right edge of piece p , we would rotate both pieces by 180° , so that the anchor piece will still be on the left, but its left edge now points to the right.

During training we noticed that determining the degree of compatibility for some pairs could be rather difficult for both the network and human experts, but it becomes quite easier for humans when looking only at a single color channel. Drawing on this observation, we trained the following networks: Red-Net, Green-Net, and Blue-Net (named after the color channels each receives as input), as well as a fourth network, RGB-Net (which receives all three channels as input). All networks share the exact same architecture, as depicted in Figure 4. During training we presented all networks with the same batch (*i.e.* same training samples); each network’s loss was calculated separately, so as not to affect the other networks. Table 2 gives a performance comparison, regarding the above individual networks and their proposed combined scheme.

We trained all networks using *stochastic gradient descent* (SGD) with standard backpropagation [36] and Adam [20], using a learning rate of 0.0001. We used a batch size of 64 and ran for a total of 850 epochs. For training we used a modern PC with 3.5GHz CPU, 32GB RAM, and a single GPU with 11GB memory.

4.3 Post-processing

To enhance the global optimization, we first apply a per-edge normalization of all compatibility scores to the range between 0 and 1, using the min-max normalization. Namely, for each piece edge e_i we calculate its compatibility with every other edge, extract the minimum and maximum across all compatibility scores, and normalize according to

$$C'(e_i, e_j) = \frac{C(e_i, e_j) - \min(C(e_i, *))}{\max(C(e_i, *)) - \min(C(e_i, *))}.$$

Next, we note that the framework described above offers no symmetry guarantee, *i.e.* that for any two piece edges e_i and e_j , $C(e_i, e_j) = C(e_j, e_i)$. Assuming that any deviation in symmetry is mostly erroneous, we manually enforce symmetry by averaging the two scores and defining the following symmetric compatibility measure

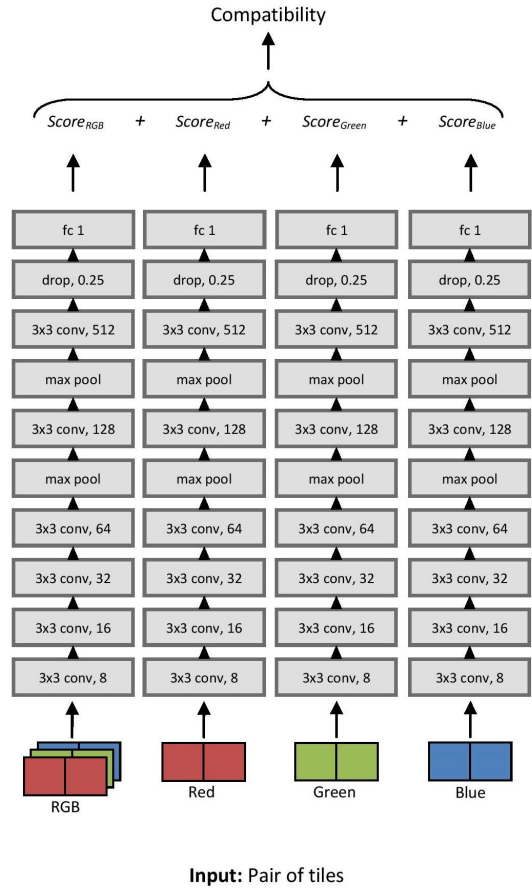


Figure 4: DLCCM architecture. The input is a pair of two squared tiles of size 50×50 pixels (*i.e.* input dimension is $50 \times 100 \times 3$). The DLCCM network contains 3.4M parameters, and uses non-linear ReLU activation functions with no bias.

$$C''(e_i, e_j) = C''(e_j, e_i) = \frac{C'(e_i, e_j) + C'(e_j, e_i)}{2}.$$

5 DATASETS

We acquired eight high-resolution images from the National Tile Museum (Museu Nacional do Azulejo, MNAZ), Lisbon, Portugal, which were kept as test data for the final reconstruction, *i.e.* they were not used during the CNN training. The size of each image is given in Table 1.

We acquired nine additional images of smaller size from the MNAZ: Five images of 25 pieces each and four images of 40, 48, 60, and 72 pieces, respectively. Due to the relatively small number of pieces per image, these images might not be adequately representative of the actual reconstruction problem. Nevertheless, given their acquisition from the museum, we regard them as sufficiently representative, in terms of content, and thus use them as a held-out validation set during the CNN training.

Image	Rows	Columns	Total Pieces	Piece Size (pixels)
Image 0	8	18	144	650
Image 1	16	16	256	150
Image 2	9	12	108	240
Image 3	11	29	319	100
Image 4	15	10	150	165
Image 5	9	23	207	225
Image 6	9	18	162	280
Image 7	12	10	120	240

Table 1: Image details of test set received from the MNAz.



Figure 5: Training set images downloaded from the Internet.

In addition to the above datasets, we also downloaded 89 images of Portuguese tile panels from the Internet, some of which were photographed by tourists. Figure 5 depicts a few downloaded images. We manually went through each puzzle, and counted the number of pieces per row and column. Knowing also the image dimensions, we could easily resize each image to 50×50 pixels. We picked nine of these images, “cut” them manually to pieces along tile lines, and added the resulting images to the validation set. The other 80 images were used as a training set for the CNN; cutting automatically the images, we gathered a total of 9,031 pieces. The automatic cutting may not always overlap fully with the actual piece boundaries, but this does not occur too often and might even aid in avoiding overfitting.

To summarize, for the training of the compatibility measure, we used a training set of 80 images and a validation set of 18 images (nine from the MNAz and nine from the Internet). For the evaluation of the compatibility measure and the overall solver’s reconstruction capability, we use a test set of the eight high-resolution images acquired from the MNAz.



Figure 6: Test set images received from the MNAz.

6 EXPERIMENTAL RESULTS

6.1 Compatibility Measure Evaluation

Previous works [34] evaluate compatibility measures by their accuracy. For each piece edge, we rank all other piece edges according to the measure in question, and report the frequency of occurrence (in percentage) that the piece edge ranked as the most compatible was indeed the correct edge. We used a generalized metric, which we call $rank_\alpha$ score, to report the percentage of actual neighboring edges found at each location of the sorted array. In other words, we define $rank_i$ of a measure as the ground truth fraction of adjacent edges which were ranked i -th most compatible according to the measure. Thus, the standard accuracy criterion for a given measure would be $rank_1$, since a perfect measure should have $rank_1 = 100\%$ and $rank_i = 0\%$ for all $i > 1$.

We trained our CNN-based compatibility measure as previously described, and evaluated it on our test set images, which were not used at all during the training phase. Our compatibility measure achieves $rank_1$ of 68.45%, assuming known piece orientation (Type 1 variant) and $rank_1$ of 56.9%, relaxing this assumption (Type 2 variant). We compared our results to the SSD [34] measure, which achieves 12.7% and 7.3%, respectively, and the MGC measure [16], which achieves 17.4% and 9.1%, respectively. Also, we compared between the performances of the individual sub-networks of our CNN model. The entire comparison is summarized in Table 2.

Next, we compared the different $rank_\alpha$ scores of our measure versus those obtained for SSD and MGC. Figure 7 presents these scores for a single test image and the average of these scores over the entire test set. The plots obtained attest to the relatively high quality of the learned measure, having the highest $rank_1$ score and

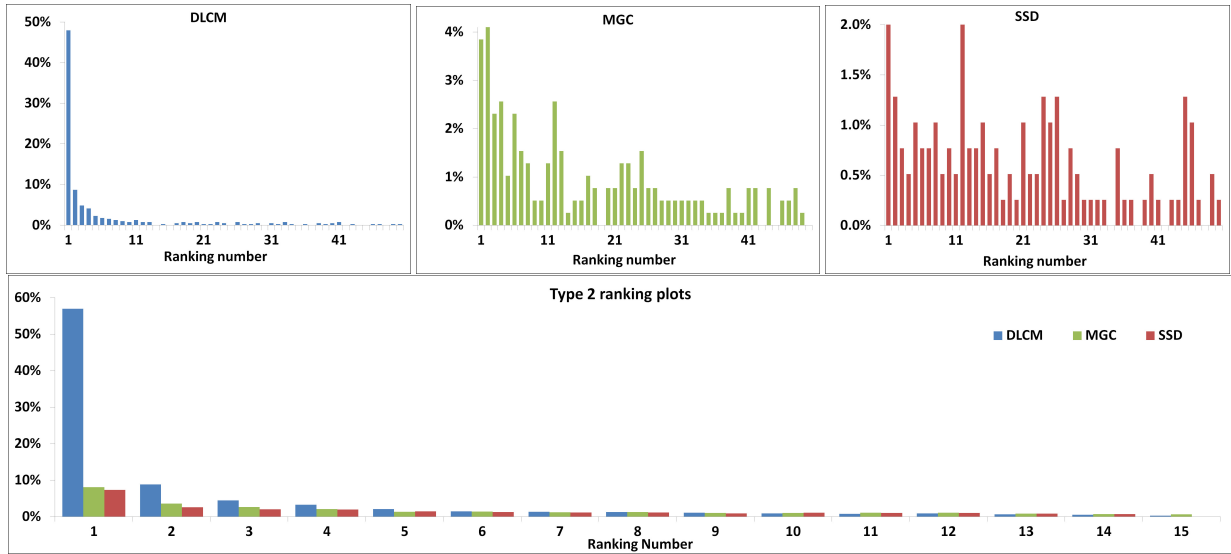


Figure 7: Rank percentages using our DLCM vs. SSD and the MGC measures for Type 2 puzzles. Top three plots correspond to a single test image (with unknown piece orientation). Bottom plot corresponds to average ranking percentage over all eight test images (with unknown piece orientation). Note the clear-cut superior performance of DLCM. Interestingly, $rank_2$ percentage of our CNN model is greater than the $rank_1$ percentage obtained for the SSD and MGC measures.

	Type 1	Type 2
SSD [34]	12.7%	7.3%
MGC [16]	17.4%	9.1%
Red-Net	56.9%	44.1%
Green-Net	57.2%	45.1%
Blue-Net	53.4%	40.8%
RGB-Net	59.5%	47.5%
DLCM	68.4%	56.9%

Table 2: Comparison of $rank_1$ scores of our DLCM with those for the SSD and MGC measures; also included are $rank_1$ scores of the DLCM’s four sub-networks (i.e. Red-Net, Green-Net, Blue-Net, and RGB-Net), demonstrating the added value of their combination.

monotonically-decreasing lower ranks, unlike the more uniform distribution obtained for the other measures.

Also, to verify the assumption that led to the post-processing steps described in Section 4.3, we evaluated the raw measure obtained by the CNN. The values obtained for this measure were 62.8% and 50.6%, respectively, for the Type 1 and Type 2 problem variants. These results strongly support the use of the post-processing step, according to Subsection 4.3.

The results clearly indicate that our trained measure is by far superior to other established compatibility measures, both quantitatively, in terms of higher accuracy, as well as qualitatively in terms of a smoother distribution.

6.2 Puzzle Reconstruction

We incorporated our newly trained compatibility measure into our enhanced GA framework, in an attempt to reconstruct each of the test set images. We report the reconstruction accuracy, according to the *neighbor comparison* definition applied in previous works, namely the fraction of correctly assigned neighbors, i.e. the fraction of ground truth adjacent edges in our solution.

We attempted reconstruction under four different variants of the problem. In all variants we assumed an unknown location of the different pieces. The variants differ with respect to a priori knowledge of piece orientation and puzzle dimensions. Obviously, the hardest variant, which is most reflective of a real-world scenario, is the one for which both piece orientation and puzzle dimensions are unknown.

We ran our GA version ten times on each image, and reported the best result. For comparison, we also tried reconstructing the images using the solver proposed by Gallagher [16]. We chose to compare against this solver, because it is one of the few solvers that supports all of the different variants and whose reported performance is still competitive relatively to state-of-the-art on available JPP benchmarks and the Portuguese tile panels in [2]. To justify the net added value of our proposed *kernel-growth* GA solver, we compared also its performance (using our DLCM) with that of the GA solvers [40–42]. The comparative results for all four cases are reported in Table 3. Examples of reconstructed panels are shown in Figure 1.

Interestingly, while inspecting the reconstructed puzzles, we noticed three puzzles that were reported as not perfectly solved, despite the fact that their overall global score was greater than ground truth. Further manual inspection revealed that apparently, the image was not assembled correctly by the museum staff, and that the solution

Method	Type 1		Type 2	
	Known dims.	Unknown dims.	Known dims.	Unknown dims.
Gallagher+MGC	—	13.0%	—	3.5%
Kernel-growth [40, 41]+symmetric DLCM	84.5%	—	58.6%	—
Multi-segment [42]+symmetric DLCM	—	—	—	62.9%
Our kernel-growth+DLCM	96.9%	96.2%	66.5%	70.6%
Our kernel-growth+symmetric DLCM	96.3%	96.0%	86.8%	82.2%

Table 3: Reconstruction comparison (from top to bottom): Gallagher’s greedy solver, using the MGC compatibility measure [16]; kernel-growth GA (due to Sholomon *et al.*) with our proposed (symmetric) DLCM; multi-segment GA (due to Sholomon *et al.*) with our (symmetric) DLCM; our proposed kernel-growth GA with (non-symmetric) DLCM, and same hybrid scheme with symmetric post-processing.

suggested by our algorithm was indeed the correct one. Figure 8 shows these segments in question.




Figure 8: Left: Images with human errors (highlighted by red), received from the MNAz. Right: Correct assembly by our system for Type 2 puzzle with known dimensions.

7 CONCLUSIONS

We presented in this paper a novel hybrid scheme, based on an enhanced GA solver and a novel DL compatibility measure, for solving the challenging, real-world task of the reconstruction of Portuguese tile panels, which is a high-profile national endeavor of significant importance to Portugal’s cultural heritage. Specifically, we demonstrated how to integrate successfully the above innovative components to achieve ground-breaking performance (over 96% accuracy for Type 1 variant and roughly 87% and 82% accuracies, for Type 2 variant with known and unknown dimensions, respectively), for tile panels containing hundreds of relatively low-resolution tiles.

Finally, we have compiled a decent benchmark of Portuguese tile panels, to be used by the Computer Vision and Evolutionary Computation communities for training and testing.

With regards to future work, we intend to improve our DL-based compatibility (by considering, for example, additional training data), in an attempt to enhance the overall performance of our GA solver. In addition, we intend to extend the capabilities of our system to handle also missing tiles and mixed panels of tiles, to meet as many practical challenges as possible associated with the Portuguese tile problem.

REFERENCES

- [1] T. Altman. 1989. Solving the jigsaw puzzle problem in linear time. *Applied Artificial Intelligence an International Journal* 3, 4 (1989), 453–462.
- [2] F. A. Andaló, G. Carneiro, G. Taubin, S. Goldenstein, and L. Velho. 2016. *Automatic reconstruction of ancient Portuguese tile panels*. Technical Report A773/2016. Instituto Nacional de Matemática Pura e Aplicada.
- [3] F. A. Andaló, G. G. Taubin, and S. Goldenstein. 2017. PSQP – Puzzle Solving by Quadratic Programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 2 (2017), 385–396.
- [4] B. J. Brown, C. Toler-Franklin, D. Nehab, M. Burns, D. Dobkin, A. Vlachopoulos, C. Doumas, S. Rusinkiewicz, and T. Weyrich. 2008. A system for high-volume acquisition and matching of fresco fragments: Reassembling Theran wall paintings. *ACM Transactions on Graphics* 27, 3 (2008), 84.
- [5] T. S. Cho, S. Avidan, and W. T. Freeman. 2010. A probabilistic image jigsaw puzzle solver. In *IEEE Conference on Computer Vision and Pattern Recognition*. 183–190.
- [6] T. S. Cho, M. Butman, S. Avidan, and W. T. Freeman. 2008. The patch transform and its applications to image editing. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1–8.
- [7] T. Chuman, K. Kurihara, and H. Kiya. 2017. On the security of block scrambling-based ETC systems against jigsaw puzzle solver attacks. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*. 2157–2161.
- [8] M. G. Chung, M. M. Fleck, and D. A. Forsyth. 1998. Jigsaw puzzle solver using shape and color. In *Proceedings of the Fourth IEEE International Conference Signal Processing*, Vol. 2. 877–880.
- [9] M. A. P. de Matos and Museu Nacional do Azulejo. 2011. *Azulejos: Masterpieces of the National Tile Museum of Lisbon*. Chandeigne.
- [10] A. Deever and A. Gallagher. 2012. Semi-automatic assembly of real cross-cut shredded documents. In *Proceedings of the International Conference on Image Processing*. 233–236.
- [11] E. D. Demaine and M. L. Demaine. 2007. Jigsaw puzzles, edge matching, and polyomino packing: Connections and complexity. *Graphs and Combinatorics* 23 (2007), 195–208.

- [12] L. Dery, R. Mengistu, and O. Awe. 2017. Neural combinatorial optimization for solving jigsaw puzzles: A step towards unsupervised pre-training. <http://cs231n.stanford.edu/reports/2017/pdfs/110.pdf>. (2017).
- [13] C. Doersch, A. Gupta, and A. A. Efros. 2015. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*. 1422–1430.
- [14] J. T. Fonseca. 2012. *Montagem Automática de Painéis de Azulejos*. M.Sc. Thesis, Instituto Superior Técnico, Universidade Técnica de Lisboa.
- [15] H. Freeman and L. Garder. 1964. Apictorial jigsaw puzzles: The computer solution of a problem in pattern recognition. *IEEE Transactions on Electronic Computers* EC-13, 2 (1964), 118–127.
- [16] A. C. Gallagher. 2012. Jigsaw puzzles with pieces of unknown orientation. In *IEEE Conference on Computer Vision and Pattern Recognition*. 382–389.
- [17] D. Goldberg, C. Malon, and M. Bern. 2002. A global approach to automatic solution of jigsaw puzzles. In *Proceedings of the Eighteenth ACM Annual Symposium on Computational Geometry*. 82–87.
- [18] J. H. Holland. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- [19] E. Justino, L. S. Oliveira, and C. Freitas. 2006. Reconstructing shredded documents through feature matching. *Forensic Science International* 160, 2 (2006), 140–147.
- [20] Diederik P. K. and Jimmy B. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014). [arXiv:1412.6980](http://arxiv.org/abs/1412.6980) <http://arxiv.org/abs/1412.6980>
- [21] D. Koller and M. Levoy. 2006. Computer-aided reconstruction and new matches in the forma urbis Romae. *Bullettino della Commissione Archeologica Comunale di Roma* (2006), 103–125.
- [22] W. Kong and B. B. Kimia. 2001. On solving 2D and 3D puzzles using curve matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Vol. II. 583–590.
- [23] D. A. Kosiba, P. M. Devaux, S. Balasubramanian, T. L. Gandhi, and K. Kasturi. 1994. An automatic jigsaw puzzle solver. In *Proceedings of the 12th IAPR International Conference on Pattern Recognition*, Vol. 1. 616–618.
- [24] H. Liu, S. Cao, and S. Yan. 2011. Automated assembly of shredded pieces from multiple photos. 13, 5 (2011), 1154–1162.
- [25] M. Makridis and N. Papamarkos. 2006. A new technique for solving a jigsaw puzzle. In *Proceedings of the International Conference on Image Processing*. 2001–2004.
- [26] W. Marande and G. Burger. 2007. Mitochondrial DNA as a genomic jigsaw puzzle. *Science* 318, 5849 (2007), 415–415.
- [27] M. A. O. Marques and C. O. A. Freitas. 2009. Reconstructing strip-shredded documents using color as feature matching. In *Proceedings of the ACM Symposium on Applied Computing*. 893–894.
- [28] A. Q. Morton and M. Levison. 1968. The computer in literary studies. In *Proceedings of the IFIP Congress*. 1072–1081.
- [29] T. R. Nielsen, P. Drewsen, and K. Hansen. 2008. Solving jigsaw puzzles using image features. *Pattern Recognition Letters* 29, 14 (2008), 1924–1933.
- [30] M. Noroozi and P. Favaro. 2016. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. *arXiv preprint arXiv:1603.09246* (2016).
- [31] G. Paikin and A. Tal. 2015. Solving multiple square jigsaw puzzles with missing pieces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4832–4839.
- [32] A. N. Pais. 2018. Director of Museu Nacional do Azulejo, *personal communication*. (2018).
- [33] C. Papaodysseus, T. Panagopoulos, M. Exarhos, C. Triantafillou, D. Fragoulis, and C. Doumas. 2002. Contour-Shape Based Reconstruction of Fragmented, 1600 B.C. Wall Paintings. *IEEE Transactions on Signal Processing* 50, 6 (2002), 1277–1288.
- [34] D. Pomeranz, M. Shemesh, and O. Ben-Shahar. 2011. A fully automated greedy square jigsaw puzzle solver. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9–16.
- [35] G. M. Radack and N. I. Badler. 1982. Jigsaw puzzle matching using a boundary-centered polar encoding. *Computer Graphics and Image Processing* 19, 1 (1982), 1–17.
- [36] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. 1986. Learning representations by back-propagating errors. *Nature* 323, 6088 (1986), 533.
- [37] M. S. Sagiroglu and A. Erçil. 2006. A texture based matching approach for automated assembly of puzzles. In *Proceedings of the 18th IEEE International Conference on Pattern Recognition*, Vol. 3. 1036–1041.
- [38] R. Santa Cruz, B. Fernando, A. Cherian, and S. Gould. 2017. DeepPermNet: Visual Permutation Learning. *arXiv preprint arXiv:1704.02729v1* (2017).
- [39] F. Schroff, D. Kalenichenko, and J. Philbin. 2015. FaceNet: A Unified Embedding for Face Recognition and Clustering. *arXiv preprint arXiv:1503.03832v3* (2015).
- [40] D. Sholomon, O. David, and N. S. Netanyahu. 2013. A Genetic Algorithm-Based Solver for Very Large Jigsaw Puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1767–1774.
- [41] D. Sholomon, O. David, and N. S. Netanyahu. 2014. A Generalized Genetic Algorithm-Based Solver for Very Large Jigsaw Puzzles of Complex Types. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 2839–2845.
- [42] D. Sholomon, O. David, and N. S. Netanyahu. 2014. Genetic algorithm-based solver for very large multiple jigsaw puzzles of unknown dimensions and piece orientation. In *Proceedings of the ACM Conference on Genetic and Evolutionary Computation*. 1191–1198.
- [43] D. Sholomon, O. E. David, and N. S. Netanyahu. 2016. DNN-Buddies: A Deep Neural Network-Based Estimation Metric for the Jigsaw Puzzle Problem. In *Proceedings of the International Conference on Artificial Neural Networks*. Springer International Publishing, 170–178.
- [44] E. Sizikova and T. Funkhouser. 2016. Wall painting reconstruction using a genetic algorithm. In *Proceedings of the EUROGRAPHICS Workshop on Graphics and Cultural Heritage*. 170–178.
- [45] K. Son, J. Hays, and D. B. Cooper. 2014. Solving Square Jigsaw Puzzles with Loop Constraints. In *Proceedings of the European Conference on Computer Vision*. Springer, 32–46.
- [46] C. E. Wang. 2000. *Determining Molecular Conformation from Distance or Density Data*. Ph.D. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.
- [47] K. Q. Weinberger and L. K. Saul. 2009. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research* 10, Feb (2009), 207–244.
- [48] H. Wolfson, E. Schonberg, A. Kalvin, and Y. Lamdan. 1988. Solving jigsaw puzzles by computer. *Annals of Operations Research* 12, 1 (1988), 51–64.
- [49] X. Yang, N. Adluru, and L. J. Latecki. 2011. Particle filter with state permutations for solving image jigsaw puzzles. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2873–2880.
- [50] F. Yao and G. Shao. 2003. A shape and image merging technique to solve jigsaw puzzles. *Pattern Recognition Letters* 24, 12 (2003), 1819–1835.
- [51] Y. X. Zhao, M. C. Su, Z. L. Chou, and J. Lee. 2007. A puzzle solver and its application in speech descrambling. In *Proceedings of the WSEAS International Conference Computer Engineering and Applications*. 171–176.