Optimization of Fresco Assembly for Accuracy

Nadezhda Shchegoleva[1(✉)], Maria Gladkaya[1], and Gennady Dik[2]

[1] St. Petersburg University, St. Petersburg, Russia
n.shchegoleva@spbu.ru, st102678@student.spbu.ru
[2] St. Petersburg LLC "System Technologies", St. Petersburg, Russia
g.dick@systechnologies.ru

Abstract. This study introduces a novel approach to the puzzle assembly problem, leveraging textural features and geometric constraints. The texture in regions extending beyond the boundaries of puzzle pieces is estimated using inpainting and texture synthesis techniques. Feature descriptors are extracted from both the original and the synthesized images. An affinity metric is defined to quantify the correspondence between puzzle pieces, and the assembly process is formulated as an optimization problem aimed at maximizing the overall affinity score. To accelerate the alignment procedure, an image registration technique based on the Fast Fourier Transform (FFT) is employed. Experiments were conducted using different image features to study the impact of their use on assembly quality. Experimental results are presented on real and artificial data sets.

Keywords: Archeological reconstruction · Puzzle solving · Partial matching

1 Introduction

In the modern world, more and more importance is attached to the preservation of cultural heritage. Archaeological artifacts are unique materials that can tell us about the culture and lifestyle of our ancestors. Their study allows us to understand the features of the technologies used, crafts, religious beliefs and the daily life of people who lived before us.

A significant part of archaeological artifacts comes to the disposal of researchers in a damaged state, which is due to erosion processes, mechanical defects, loss of individual fragments and other destructive factors. The process of reconstruction of such objects is a complex scientific and technical task that requires the development of effective algorithms for automated restoration. The development and improvement of such methods is an urgent problem of great importance for the preservation of the cultural heritage of mankind.

It is worth noting that restoring broken artifacts manually is usually a very time-consuming task that requires a lot of time. The use of algorithms that can automatically "assemble" pieces of artifacts together can significantly speed up research and allow archaeologists to focus on other more important aspects.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
O. Gervasi et al. (Eds.): ICCSA 2025 Workshops, LNCS 15894, pp. 294–308, 2026.
https://doi.org/10.1007/978-3-031-97648-3_20

One of the key problems in the reconstruction of frescoes is the destruction of the edges of the fragments due to erosion and mechanical damage, which leads to their unevenness and deformation. As a result, critical data for the task of reconstruction can be lost under the influence of time, natural factors or acts of vandalism that cause irreparable damage to the historical and cultural heritage.

As noted earlier, the destroyed frescoes are characterized not only by the geometric parameters of the fragments, but also by visual characteristics such as texture, color, continuity and shape of lines. In this regard, the development of effective reconstruction algorithms requires an integrated approach that takes into account not only geometric correspondences and coincidences of images in different fragments, but also their spatial location relative to each other.

All of the above makes the task of developing algorithms that restore archaeological artifacts from several fragments extremely relevant and important for modern society.

2 Algorithms for Automatic Reconstruction of Images from Their Fragments

Over the past few years, a large number of different algorithms have been proposed that use textural and geometric features to automatically reconstruct images from their fragments.

The simplest case of reconstruction of images divided into uniform square fragments is considered by authors in the work [3]. The authors proposed a system capable of simultaneous reconstruction of several objects whose fragments are chaotically mixed.

The absence of some fragments is also not a critical factor for the algorithm, since the proposed method is based on identifying and sequentially combining the most compatible pairs of fragments. In the course of the algorithm, new elements are iteratively attached to the formed pairs based on an assessment of the degree of their compliance. This evaluation takes into account changes in pixel brightness along the boundary of the tiles being joined, as well as derivative information along the adjacent tiles boundaries. The developed algorithm demonstrated a high accuracy of recovery, reaching 95.3%.

Authors of article [8] considered a solution to a problem where the fragments are connected in the form of a "brick wall", which is essentially a set of rectangular parts that can have different sizes and can be placed next to each other with an arbitrary displacement. In their work, they showed that the solution can be reduced to finding the correct displacement between two adjacent parts. The method is based on the use of a number of metrics to assess the difference and compatibility of correlations between all sides of the fragments. The difference score is calculated based on the L_2 norm and color channels of the images. The compatibility score is a key metric when deciding whether to attach a pair of fragments. The compatibility score is calculated by estimating the difference between the ith fragment and the jth fragment and estimating the difference between the ith fragment and the kth fragment, where i is the number of the

fragment for which the best match is calculated, j is the number of the fragment that is considered at the current iteration of the algorithm, and k is the number of the fragment with which the ith fragment had the best correlation in previous iterations. The accuracy of the algorithm was 99.8%.

In 2019, group of researchers [5] proposed a method for reconstructing objects consisting of rectangular fragments of arbitrary sizes, provided that the orientation of all parts is known, but their position remains uncertain. To solve this problem, an algorithm was developed based on the comparison of the angles of fragments using the MatchLift framework. MatchLift provides a convex optimization method for finding cyclically consistent matches from a set of noisy inputs. The accuracy of the algorithm was 87%. The system developed in 2021 [4] solves the problem where the fragments are convex polygons obtained by the intersection of an arbitrary number of lines. The solution is a greedy algorithm, simple but reliable. The algorithm starts the search with a randomly selected edge of the fragment and looks for a match with the edge of some other free fragment based on the Euclidean transformation of this fragment. The accuracy of the algorithm is 100%, but the computational complexity is $O(M^4)$, where M is the number of fragments.

In their 2018 study [6], research group proposed an algorithm for solving a problem in which fragments have erased boundaries and do not have a specific shape, and the range of permissible transformations between parts is continuous, that is, the number of ways to place fragments relative to each other is infinite. The algorithm proposed by the authors involves restoring the image at the edges of the fragments and arranging the fragments relative to each other in such a way so that the extrapolated slice edges intersect, but the original slice regions do not. The best placement is found through an iterative process. The area around one fragment that defines the set of possible placements for the second fragment is determined by the Minkovsky sum. Textural, statistical, and the length of the joined sides are used to assess the similarity between the tiles. The accuracy of the algorithm is 100% if the number of fragments $M < 10$, for $M \in [11; 20]$—92.6%, for $M \in [21; 30]$—98.5%.

In 2016, two authors [7] published a paper in which the problem of searching for neighboring fragments is solved using a genetic algorithm. They proposed their own algorithm that creates sets of combined fragments (partial reconstruction pool). The pool of partial reconstructions grows as a result of recombination and selection over generations. Which is resistant to noise and emissions, and a new genetic algorithm selection procedure has been proposed that balances fitness and diversity in the population.

A group of researchers proposed to solve the problem of assembling fragments of Turanian wall paintings on the basis of calculating the normals of 2D scans [9]. To determine the most likely connection of fragments into a single painting, the points of the surface obtained by transforming the fragments into a single surface are used. The same thickness.

The above list of studies allows us to conclude that depending on the features of the fragments (identical square, arbitrary fragments in the shape of a polygon,

the presence of damage at the edges of the fragments), as well as on the presence or absence of information about the orientation of the fragments, researchers use different approaches to solve this problem.

The most promising methods use algorithms based on heuristic approaches (greedy algorithms, minimal spanning tree), optimization methods (MatchLift, convex optimization), probabilistic models (genetic algorithms), and machine learning methods. Many modern works demonstrate high accuracy when assembling images from fragments, but the complexity of the algorithms can vary significantly. Algorithms that allow processing fragments that have a shape that is closest to real fragments of archaeological artifacts are usually based on the use of geometric characteristics.

It is worth noting that the neglect of some authors to use the color and texture of neighboring fragments is an ineffective tactic, since graphic information contains many features that can be useful in solving the reconstruction of archaeological artifacts. The graphic information contained in the images of fragments is very diverse and can include many features and components, and depending on the field of study, the informative value of certain features varies. In archaeology, the most informative features of images include the direction of marble veins, the pattern of carving on the surface, painting on the external and internal surface of the object, recesses and horizontal circles formed as a result of smoothing clay with fingers when rotating the pot on the wheel during the manufacturing process.

In archaeology, erosion and damage from impacts lead to the disappearance or destruction of fragments, changes in their boundaries. This makes it necessary to use all available information (including image features) to solve the problem of reconstructing frescoes, since geometric approaches based on the exact coincidence of the edges of fragments are not applicable for assembly if the boundaries of the fragments have not been preserved for a number of reasons.

3 Methods Used for Reassembly of Broken Frescoes

This section describes the methods used to restore broken frescoes. The recovery process consists of three main parts:

- extrapolation of graphic information from the edges of the fragments to the area around the fragment; The main task of this part of the software package is to predict graphic information on the erased edges of fragments. The implementation of this part is represented as an algorithm that determines the color value of pixels in an area outside the boundaries of the tiles with a confidence measure.
- finding a measure of affinity to determine the compatibility of fragments; The characteristics obtained from the predicted texture outside the fragment correlate with the graphical characteristics of the original region of possible neighboring fragments. To effectively compare the textures of fragments, a similarity measure has been introduced, which uses edge continuity, texture

patterns, and color similarity, and the task of assembling is set as its optimization.
– fragment placement using Fast Fourier Transform.

The proposed approach to automated mural restoration is based on a fast and reliable method that finds the best fragment placement that maximizes the overlap and continuity of the fragment textures while respecting geometric constraints (the original areas of the fragments cannot overlap).

After preprocessing the data, the first step is to predict the pixel values in the band around the fragment boundary; This step applies to all fragments separately. The prediction algorithm automatically fills in the extrapolation area using the information from the central part of the fragment. The basic idea behind extrapolating the image/texture of the original region of the fragment outward is that the correlation between the characteristics of the extrapolated region and its true adjacent region is significantly higher than between alternative pairs of fragments.

In extrapolation, the confidence of the extrapolated region is entered as a parameter. This parameter reflects the confidence level for each predicted pixel and is used by subsequent processes. The validity depends on the texture structure, such as edge continuity, texture roughness, and distance to the original fragment boundary.

The feature values are then calculated, both in the original fragment and in the extrapolated region. The proposed approach does not limit the number of features and does not limit the type of features of the image that can be used in the restoration of a fresco. Any textural feature that is thought to improve the success of an assembly can be incorporated into the process. The combination of feature values and image confidence is used to create a measure of the similarity of the corresponding fragments.

Finding the best placement of fragments during assembly is done using the Fast Fourier Transform. Initially, each fragment has a random position in space. FFT is used to iteratively find the fragment placements relative to each other that maximize the correlation between the fragments. To improve the assembly, the algorithm must be able to determine whether a particular arrangement of parts improves the assembly or not; This is done through a common measure of affinity, defined as the sum of the similarity measures of all points in space. During the assembly process, the similarity measure is optimized. In this way, the best placement of the fragments is found.

3.1 Extrapolation of Graphic Information of Fragments

The first step in the assembly process is to extrapolate each fragment to the area around its boundary by predicting graphical information [2]. Image extrapolation refers to the process of filling the area around a fragment in the form of a strip n pixels wide. The task is to fill the strip with a texture sample from the original region of the slice.

3.2 Cost Function to Determine Build Quality

When comparing or calculating the similarity of possible two adjacent fragments, a pixel-by-pixel comparison does not make sense. Therefore, a measure to determine the quality of the side is proposed, described below. Let us call it the function of cost F_{cost}. [1]

The total cost is the sum of similarity conditions and geometric constraints for all points on a predefined board or space. The only parameter of this indicator, which reflects the quality of the assembly of the parts based on textural features and geometric shape, is the transformation of T_i parts.

The F_{cost} value becomes negative if there is a good match between the images in the tiles. The correspondence between the parts increases, while the cost function is optimized.

3.3 Correlation and Fragment Placement with Fast Fourier Transform

The problem of assembling a broken fresco can be formulated as the optimization of the above cost function, the optimization task requires too much computational cost. Note that minimizing the distance function, D, is equivalent to maximizing the correlation between the fragments. [1]

Thus, the algorithm for assembling frescoes looks like this:

Algorithm 1. Fresco Assembly Algorithm

1: Extrapolate all fragments.
2: Restore the images in the extrapolated regions.
3: Place all fragments on a sufficiently large board B.
4: **while** the puzzle is not unambiguously assembled **do**
5: Find the best transformation for a randomly selected fragment t using FFT and move it accordingly.
6: **end while**

The main disadvantage of this algorithm is that there can be multiple solutions to the problem. These multiple decisions depend on the initial placement of the fragments on the (B) and the random selection of the tth fragment. In such situations, a developed measure of similarity is used. For the possible placement specified by the proposed methodology, F_{cost} is calculated. If the algorithm achieves a new solution that has a lower cost than the previous one, then the new solution becomes the best way to place the mural fragments relative to each other. Iterations continue until the last possible placement is the lowest cost.

4 Extract Features from Images

One of the most informative types of data that can be extracted from fragments is graphical and visual data. In this regard, the development of computer systems

to simplify the work of archaeologists using methods of extracting textural and structural features from an image and applying these features in algorithms for assembling archaeological artifacts is becoming widespread.

This approach allows you to more fully and accurately present information about the content of the image. Features allow you to highlight the most significant parts of the image to solve the problem of assembling murals, ignoring less important or noisy data. This makes the algorithms more efficient.

4.1 Histogram of Oriented Gradients

Histogram of Oriented Gradients (HOG) is a feature extraction technique widely used in computer vision tasks such as object recognition and human detection. HOG features are extracted from images to represent them as vectors that contain information about gradients and orientations. The calculated HOG features of the two images can be compared using various methods, such as cosine or Euclidean distance.

You can determine whether one image is relative to another based on the maximum correlation or the bias score.

4.2 Local Binary Patterns

Local binary patterns are a descriptor with low computational complexity proposed by T. Oyala in 1996 to extract texture features [11]. Due to its simplicity and resistance to changes in brightness and small image rotations, this method is a common and effective way to solve computer vision problems, in particular, face recognition.

4.3 Halftone Adjacency Matrix

A grayscale contiguity matrix (sometimes Gray Level Co-Occurence Matrix (GLCM) is a matrix that shows the distribution of gray levels between adjacent pixels in an image. GLCM describes the distribution of pairs of pixels based on their brightness or intensity. For each pair of pixel values, the number of times they occur in a particular relative position in the image is calculated.

Such matrices represent the texture properties of an image, but are sensitive to its scaling and rotation. They are not convenient for direct use in image analysis due to their large size; therefore, they are often only used to calculate various metrics for a more compact representation of textures. In this case, 14 characteristics proposed by R. Haralik in 1973 [10] and used for the task of texture classification are extracted for each image.

4.4 Gabor Filter

A Gabor filter is a linear filter used in computer vision and image processing to analyze texture, highlight edges, and detect structural features of an image. It is based on the Gabor function, which is a harmonic wave (sinusoidal function) with a Gaussian shell.

5 Description of Experiments

5.1 Preparation of Input Data

To simulate the separation of frescoes, the algorithm for extrapolating the graphic information of fragments described in 1.1 is implemented in the work. The algorithm is fed an image of the source fragment and a mask that indicates the target area that the algorithm will fill (Fig. 1). An example of the result is shown in Fig. 2.

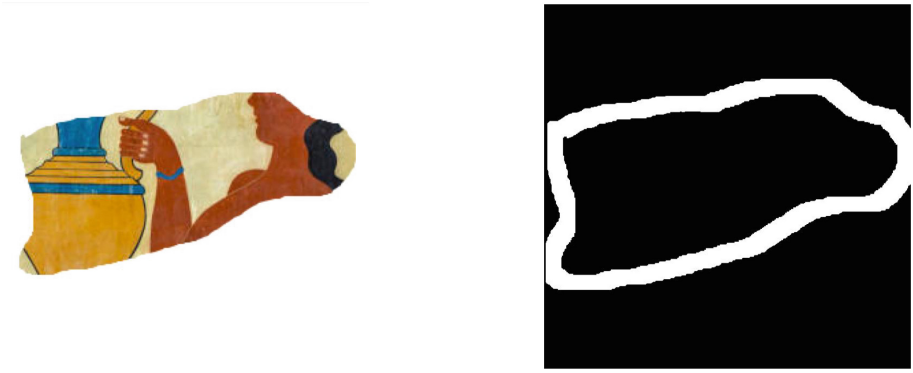

Fig. 1. Image of the source fragment and target area.



Fig. 2. Result of the algorithm for image extrapolation.

In addition, a script has been implemented that automates the division of the fresco into fragments. An image of the fresco and a mask are supplied to the input, which indicates where the cuts will be (Fig. 3). The output is several images - one for each fragment (Fig. 4).



Fig. 3. Image of the original image and the fragment mask.



Fig. 4. Results of the algorithm.

5.2 Selection of Image Features

Experiments were carried out on two fragments to determine which features of the images give the highest quality of assembly. The results are presented in Table 1.

The algorithm is based on the textural features of the images, so the following features gave the maximum accuracy during assembly: Local Binary Pattern (LBP), Histogram of Oriented Gradients (HOG), Gabor filter. They were chosen as features of the images that will be used in further experiments.

An example of how the algorithm works with the selected features: the initial placement of fragments is shown in Fig. 5, and the result of the work using HOG, LBP and Gabour filter in the listed order is shown in Fig. 6.

Table 1. Results of the experiment.

Sign	Proposed offset, (x, y)
HOG	(304, 96)
LBP	(300, 100)
Gabour filter	(300, 100)
GLCM	(0, 0)



Fig. 5. Initial arrangement of fragments.



Fig. 6. Results of the algorithm for two fragments using HOG, LBP, Gabour filter, respectively.

5.3 Build Quality Metrics

Texture Affinity Index. The Textural Similarity Index (TSI) is a metric that is used to quantify the degree of similarity between the textures of two images. Unlike traditional metrics such as MSE (root mean square error) or SSIM (structural similarity index), TSI takes into account the textural characteristics of the image, such as gradient histograms or statistical properties of local textures. TSI measures the difference between the textures of two images using texture descriptors such as:

- GLCM (Gray-Level Co-occurrence Matrix)
- LBP (Local Binary Patterns)

These descriptors are described in Chap. 2.

In general, the TSI index can be defined as follows:

$$TSI(I_r, I_t) = 1 - D(T_r, T_t)$$

where:

- I_r—reference image,
- I_t—image under test,
- T_r, T_t—texture image descriptors,
- $D(T_r, T_t)$—function of distance between texture descriptors.

To compute TSI based on GLCM, you can use the cosine distance between feature vectors:

$$D(T_r, T_t) = \frac{\sum_i T_r^i T_t^i}{\sqrt{\sum_i (T_r^i)^2}\sqrt{\sum_i (T_t^i)^2}}$$

The LBP-based method uses the Hellinger distance between the histograms of the reference image and the image under test:

$$D_H(T_r, T_t) = \sqrt{1 - \sum_i \sqrt{T_r^i T_t^i}}$$

Interpretation of the results:

$TSI \approx 1$ means that the textures are almost identical.

$TSI \approx 0$ indicates a complete difference in textures.

Euclidean Matching Accuracy. Euclidean Matching Accuracy (MA) is a metric used to measure the accuracy of predicted coordinates (or points) relative to their true values

For the coordinate prediction problem, the accuracy of the match (MA) is usually defined as the Euclidean mean distance between the predicted and true coordinates.

Interpretation of the results:

$MA \approx 0 \rightarrow$ predicted points perfectly match the reference points (excellent accuracy).

$MA < 1$ pixel $\rightarrow$ a good match.

$MA > 1$ pixel $\rightarrow$ degradation of the prediction quality.

The lower the MA, the higher the accuracy of the prediction.

5.4 Results of Experiments

Let's consider the first experiment. 6 fragments were served at the entrance (Fig. 7).



(a) (b) (c)

Fig. 7. (a) Original image, (b) Sections imitating fragments of a broken fresco (c) Image prepared for the input of the algorithm.

The result of the program using HOG (a), LBP(b) and Gabour filter(c) in the listed order in Fig. 8.



Fig. 8. The result of the work when using various features.

The effectiveness of the assembly using various features of the image can be evaluated in Table 2:

Table 2. Metrics of the quality of the algorithm

Metrics	HOG	LBP	Gabour filter
Euclidean precision (MA)	0.9px	1.5px	1.2px
TSI Ratio (GLCM)	0.945	0.894	0.928
TSI Ratio (LBP)	0.9762	0.9756	0.9761

At the next stage, an image with a lost fragment was fed to the input of the algorithm (Fig. 9). Such features of the input data are less informative, since graphical information is the most important aspect for the most accurate assembly, which, in turn, complicates the task for the algorithm.



(a) (b) (c)

Fig. 9. (a) Original image, (b) Sections imitating fragments of a broken fresco (c) Image prepared for the input of the algorithm.

The result of the program using HOG (a), LBP(b) and Gabour filter(c) in the listed order in Fig. 10.

The effectiveness of the assembly using various features of the image can be evaluated in Table 3:

Table 3. Metrics of the quality of the algorithm

Metrics	HOG	LBP	Gabour filter
Euclidean precision (MA)	1.9px	3.1px	2.7px
TSI Ratio (GLCM)	0.923	0.885	0.91
TSI Ratio (LBP)	0.9786	0.9682	0.9787



Fig. 10. The result of the work when using various features.

6 Conclusion

Summing up the work done, it can be noted that an algorithm has been developed that automatically restores the frescoes. The proposed approach allows you to find the best placement of fragments, which maximizes the coincidence and continuity of the fragment textures while respecting the geometric constraints that consist in ensuring that the fragments do not intersect with each other.

References

1. Ercil, A.: Optimization for automated assembly of puzzles. In: Proceedings of of EUROPT (2008). https://doi.org/10.1007/s11750-010-0156-6
2. Criminisi, A., Perez, P., Toyama, K.: Object removal by exemplar-based inpainting. In: CVPR, pp. 721–728 (2003)
3. Son, K., Moreno, D., Hays, J., Cooper, D.B.: Solving small-piece jigsaw puzzles by growing consensus. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
4. Kreis, B., Dengler, N., de Heuvel, J., Menon, R., Perur, H., Bennewitz, M.: Compact multi-object placement using adjacency-aware reinforcement learning. In: 2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids), pp. 698–705 (2024)
5. Wang, T., et al.: Robust and flexible puzzle solving with corner-based cycle consistent correspondences. In: Proceedings of EG UK Computer Graphics and Visual Computing (2019)
6. Derech, N., Tal, A., Shimshoni, I.: Solving archaeological puzzles. Pattern Recogn. **119**, 108065 (2021). https://doi.org/10.1016/j.patcog.2021.108065
7. Sizikova, E., Funkhouser, T.: Wall painting reconstruction using a genetic algorithm. In: Proceedings of EUROGRAPHICS Workshop on Graphics and Cultural Heritage (2016)
8. Gur, S., Ben-Shahar, O.: From square pieces to brick walls: the next challenge in solving jigsaw puzzles. In Proceedings of IEEE International Conference on Computer Vision (ICCV) (2017)
9. Brown, B.J., et al.: A system for high-volume acquisition and matching of fresco fragments: reassembling Theran wall paintings. In: ACM Transactions on Graphics (Proc. SIGGRAPH), vol. 27, no. 3 (2008)

10. Haralick R.M., Shanmugam K., Dinstein I.: Textural features for image classification. In: Proc. IEEE Trans. Syst. Man Cybern. **SMC-3**(6), 610–621 (1973)
11. Ojala, T., Pietikainen, M., Harwood, D.: A comparative study of texture measures with classification based on featured distributions. Pattern Recogn. **26**(1), 51–59 (1996). https://www.sciencedirect.com/science/article/abs/pii/0031320395000674. Accessed 20 May 2022