Automatic stitching of panoramas for geological images of polished sections

Gleb Nikolaev[1], Dmitry Korshunov[2], Alexander Khvostikov[1]

[1]Faculty of Computational Mathematics and Cybernetics, Lomonosov Moscow State University, Moscow, Russia
nickolaev.gleb03@gmail.com, khvostikov@cs.msu.ru
[2]Geological Institute, Russian Academy of Sciences, Moscow, Russia - d.korshunov@ginras.ru

Keywords: Image stitching, panoramas, geology, microscopic images, polished sections, seam removal.

Abstract

This paper addresses the challenge of constructing panoramic images from a set of overlapping microscopic images of geological polished sections.This task is valuable both for automatic mineral detection and for the creating of advanced virtual educational collections. We present LumenStone P1, a new dataset of polished section images for panorama stitching, along with an automatic method for constructing panoramas of these images. Comparisons with existing approaches, including commercial ones, demonstrate that our method is comparable in stitching quality. The numerical value of panorama stitching quality is 0.33 by RMSE for our dataset. The results demonstrate the effectiveness of the proposed method and its practical applicability for the preparation of series of images and further computer analysis of panoramic geological images of polished sections. The LumenStone P1 dataset intended for panorama creation, as well as the source code of the proposed method, are publicly available.

1. Introduction

In microscopy across various natural sciences, structural analysis of the studied sections often presents challenges. In ore microscopy, it is crucial to accurately assess the ratio of minerals and the forms of their intergrowths. This is essential for addressing both scientific problems (such as studying the mechanisms of crystal growth and determining the order of crystal formation) and applied problems (such as determining the size and cohesion of ore particles for processing). Traditionally, these issues were addressed through lengthy visual observation using a microscope and photographing the objects of study. However, with the advancement of computer vision, it is now possible to automatically obtain a vast amount of valuable information. Nevertheless, conventional photographs have a limited field of view and are unsuitable for subsequent computer processing. The solution to this problem is to obtain panoramic images of the sections.

Panoramic images are also utilized in another area: the creation of digital mineralogical collections. These collections serve two primary purposes: educational and industrial. For specialists in mineral diagnostics, acquired expertise is crucial, and digital museums significantly enhance opportunities for self-learning. For geologists involved in exploration and mineral processing, digital collections allow for more frequent use of analogies, enabling them to quickly find solutions to their problems. However, existing collections such as (Virtual Petrography, 2024) or (Virtual Microscope, 2024) share two common drawbacks. They are either extremely limited in scope or the published images are not panoramic, displaying a limited area that lacks important information.

There is currently no single, convenient solution for stitching panoramas. The most commonly used programs for stitching panoramas are graphic image processing software such as (Adobe Photoshop Software, 2024) or universal image analysis programs like Fiji with the stitching plugin (Preibisch et al., 2009) and (tshsoft PanoramaStudio, 2024). These programs allow users to stitch a series of photographs into one large panorama. Additionally, many researchers have developed their own specialized software solutions (Zukić et al., 2021), (Ro and Kim, 2021), (Zhang and Subasinghe, 2012). Major microscope manufacturers like Carl Zeiss and Leica also offer automatic photographing and image stitching systems (Research Microscopy Solutions, 2024), (Aperio Digital Pathology Slide Scanners, 2024), though these solutions are very expensive.

In this article, we have attempted to address the existing challenges in creating panoramic images of geological polished sections. We are making publicly available a dataset for creating panoramas, the resulting panoramas created with the proposed method, as well as the source code of our method as part of the Python package `petroscope` [1]. We hope that the provided materials and tools will be useful to researchers and professionals in the field of geology and will contribute to the development of methods for analyzing and visualizing geological samples.

2. Used data

The task of stitching panoramas from a set of geological images has received relatively little attention in the scientific community. Additionally, there are no open datasets available for this task, making it impossible to practically compare existing methods. To address this issue, we have collected our own dataset for stitching panoramas of microscopic geological images and made it publicly available and free to use.

This work continues our research in the field of developing algorithms for the automatic analysis of images of geological polished sections (Khvostikov et al., 2021), therefore, we are expanding the previously created LumenStone [2] dataset and including a subset of images called P1, specifically aimed at creating panoramic images from a series of polished section images.

[1] `https://github.com/xubiker/petroscope`
[2] `https://imaging.cs.msu.ru/en/research/geology/lumenstone`

Publisher's note: Copernicus Publications has not received any payments from Russian or Belarusian institutions for this paper.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-2-W1-2024-39-2024 | © Author(s) 2024. CC BY 4.0 License.

In this work, we consider samples from the same deposit and the same mineral association as in the LumenStone S1 dataset intended for solving the segmentation task. The material used was collected from 30 ore deposits of the CIS. Prepared samples of polished sections were analyzed using a Carl Zeiss AxioScope 40 microscope, photographing was carried out using a Canon Powershot G10. Each image was taken under x100 magnification, the resolution is 3396×2547.

To create each panorama, we use 20 or 25 photographs (4 rows of 5 photographs or 5 rows of 5 photographs). All images used in this work (10 sets for creating panoramas) are available in the LumenStone P1v1 dataset. Each image contains the information of the expected location in the panorama grid. The collected panoramas, resulting from the application of the presented method, are also available for download within LumenStone P1v1.

3. Proposed method

The proposed method for stitching panoramas from a grid of geological polished section images consists of two main parts. The first part involves image registration and includes image preprocessing, key point detection, key point matching, panorama stitching, and making adjustments. The second part improves the final visual representation of the panorama by removing (masking) the seams between the images. Each of these components is described in detail below.

3.1 Distortion correction

Geometric distortions, arising from the physical characteristics of cameras and often present in photographs, can have a strongly negative effect when creating panoramic images. The presence of distortion can lead to poor matching of key points, and the resulting panorama can start to noticeably bend in one direction or another. Moreover, as the size of the panorama increases, the final distortions will only intensify. Therefore, for higher quality and more stable stitching of panoramas, it is necessary to identify the geometric distortions produced by the camera and then correct them.

Since all images forming the panorama are taken under identical conditions using the same microscope and camera, it is sufficient to determine the distortion parameters once and then correct all available images.

Distortion can be calculated in various ways. In this work we used the Brown-Conrady distortion model, which defines the radial distortion as:

$$x' = x(1 + k_1 r^2 + k_2 r^4 + k_3 r^6),$$
$$y' = y(1 + k_1 r^2 + k_2 r^4 + k_3 r^6),$$
$$r = \sqrt{x^2 + y^2},$$

and tangential distortion as:

$$x' = x + 2p_1 xy + p_2(r^2 + 2x^2),$$
$$y' = y + p_1(r^2 + 2y^2) + 2p_2 xy,$$

where r is the distance from pixel with coordinates (x, y) to the center of the image. So, the distortion in the used model is determined with coefficients k_1, k_2, k_3, p_1, p_2. But in our case due to the nature of photography, tangential distortion can be omitted. In order to find all optical parameters we took a photograph of a ruler with laser markings placed every 10 microns (Fig.1). We calculated the intrinsic camera parameters focal length (f_x, f_y) and optical centers (c_x, c_y), which form a camera matrix:

$$\begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

and also found the radial distortion coefficients k_1, k_2, k_3. All the mentioned optical parameters were selected so that the lines in the photograph of the ruler became straight. The resulted parameters are: $f_x = 5.51 \cdot 10^4, f_y = 5.58 \cdot 10^4, c_x = 1.87 \cdot 10^3, c_y = 1.42 \cdot 10^3, k_1 = 21.78, k_2 = -134.8, k_3 = -7.36$. Coefficients of tangential distortion are omitted: $p_1 = 0, p_2 = 0$. Bicubic interpolation was used when resampling the images.

Figure 1. Microscopic photo of ruler used to evaluate the distortion.

It is also worth noting that such physical calibration can be avoided by stitching the panorama with different distortion coefficients and finding the optimal correction coefficient using binary search by analyzing the RMSE metric of key points matching.

3.2 Key points detection

One of the key stages in the method of creating panoramas is the step of registration of two adjacent images, which involves finding such a transformation of the coordinate system of one image into the coordinates of another so that the overlapping parts of the images coincide. There are numerous methods for solving the image registration problem, including modern methods based on deep learning (Sun et al., 2021). Despite this, methods based on key point detection, such as SIFT (Lowe, 1999) and SURF (Bay et al., 2006), remain relevant in this field.

For microscopic images of geological polished sections, which have high contrast and clear boundaries of minerals, key point-based methods work efficiently and reliably. In this work we used the SIFT method, one of the best classical methods for key point detection. In addition to detecting key points, this method also assigns each key point a descriptor — a vector of 128 floats. These descriptors are utilized for point identification and for finding correspondences in adjacent images. We used the original SIFT parameters as in (Lowe, 1999): number of octaves = 4, number of scale levels = 5, initial $\sigma = 1.6, k = \sqrt{2}$. An example of applying SIFT detector to the pair of adjacent images is shown in Fig. 2.

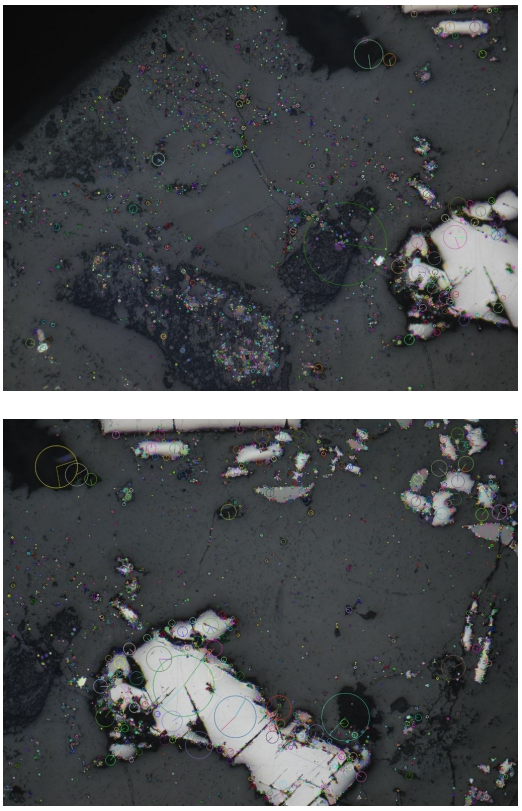

Figure 2. Example of two adjacent images with visualized SIFT key points.

3.3 Key points matching

The next step of the method involves finding pairs of key points in two adjacent images.

Considering the characteristics of the available data, specifically that there cannot be a strong rotation or stretching between two adjacent images of the panorama, we chose one of the simplest methods: K-Nearest Neighbors Matching (Jakubović and Velagić, 2018). This method selects $K = 2$ descriptors from the second image for each descriptor in the first image that are most similar to it based on MSE. After that a filtering procedure, which is inspired by (Lowe, 2004) is performed. For each key point we calculate the ratio of

$$\eta = dist(p, B1)/dist(p, B2),$$

where p is the considered key point, $B1$ and $B2$ are the first and second nearest key points on the second image. After that, we sort the points in ascending order of η and choose the first 75 points. This is justified by the fact that the smaller the value of the ratio, the more confident we are in the correct correspondence of this point.

The visualization of selected key points and their matches for a pare of adjacent images is shown in Fig. 3.

3.4 Rough panorama stitching

The next step after finding correspondences between all pairs of adjacent images is determining the homography matrix, which is a matrix of perspective transformation from the homogeneous coordinates of one image to the coordinates of another. Generally, 4 pairs of points are sufficient to uniquely determine the homography between two images, but typically, the number of found point pairs is significantly greater than 4, and an overdetermined system of equations is solved to find the target homography. The homography computed in this way is more accurate due to the reduced influence of noise. To find the best possible homography from the previously selected key points, the RANSAC algorithm (Fischler and Bolles, 1981) is applied. Homography hypotheses are constructed based on $n = 50$ randomly selected pairs of points, and then we check how well the hypothesis transforms the points from one image to the points in the other image. To do this, we compute the number of key points for which the distance from the transformed point to the corresponding point in the second image is less than the threshold T:

$$dist(LM, M') < T,$$

where $M = (x, y)$ is a key point in the first image, $M' = (x', y')$ is the correspondent point in the second image, L is the operator performing the homography transformation, and T is the threshold value. The key points that meet the threshold conditions are called inliers. In this work we use $T = 0.5$.

The described computations are repeated for 1000 random samples. The homography hypothesis with the largest number of inliers is retained. The final homography matrix is built using all found inliers.

The described method of calculating homography matrix allows to stitch a pair of adjacent images by applying the transformation to one of the images. An example of stitching two images in such a way is shown in Fig.4.

The described calculations are applied to all pairs of adjacent images, thus forming homography matrices between all neighboring images. If matrix H_1 defines the transformation between the first and second images, and H_2 between the second and third, then $H_2 \cdot H_1$ will transform the first image into the coordinates of the third. By sequentially multiplying the homography matrices in this manner, we can obtain transformations that link any two images in the panorama.

When stitching the panorama, we choose a reference image (the central one), and then compute the transformations that align all other images to the reference. Sequences of homography matrix multiplications of minimal length are selected. If there are multiple sequences of multiplications of the same length, the final homography matrix is computed as the element-wise median of the multiplication results of all such sequences.

As a result of this stage, we obtain a panorama in which all images are aligned.

3.5 Panorama adjustment

The approach described above allows to find reasonably good correspondences between pairs of adjacent images for alignment. However, when assembling the entire panorama and multiplying the homography matrices, errors accumulate, leading to significant mistakes and discrepancies. To improve the results, it is necessary to adjust the values of the homography matrices across the entire panorama minimizing the overall loss functional. Typically, in 3D computer vision, this is achieved using the bundle adjustment approach (Triggs et al., 2000). However, our task is simpler because, due to the nature of the imaging, the angles of the photographs are consistent, and perspective distortions are minimal. In this work, we use a simplified version of this method, described below.

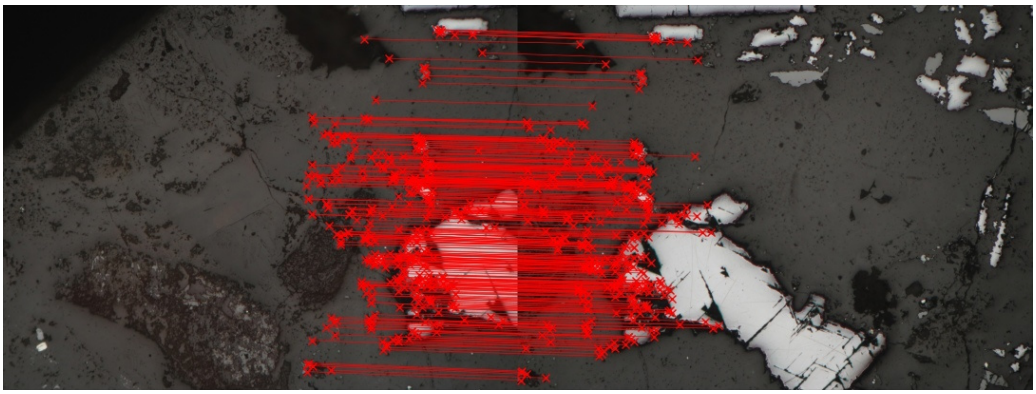

Figure 3. Visualization of filtered key points and their matches for the two adjacent images.

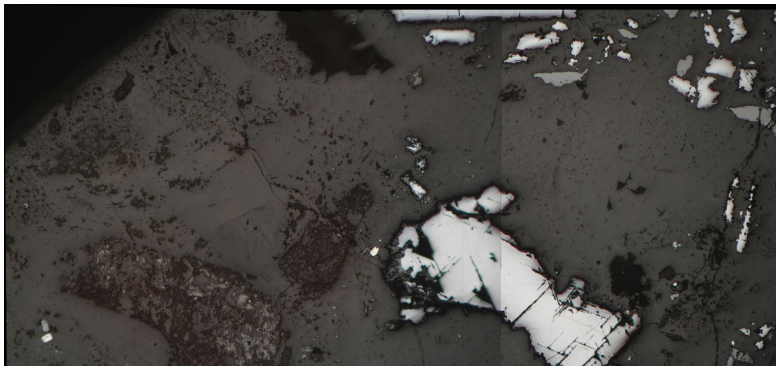

Figure 4. Result of stitching of two adjacent images. The intersection region is overlapped with the left image.

The inliers (pairs of key points) obtained while finding the best homographies can be used to assess the quality of the resulting panorama. Let image 1 be transformed into the panorama by homography H_1 and its adjacent image 2 by homography H_2. Let the selected key points of the first image be $p = (p_1, ...p_n)$ and the corresponding points of the second image be $p' = (p'_1, ...p'_n)$, where each $p_i = (x_i, y_i)$, $p'_i = (x'_i, y'_i)$. We would like the transformed points p by matrix H_1 and the transformed points p' by matrix H_2 to result in the same points, as they reflect the same object in different photographs, and their images should coincide in the final panorama.

Using the least squares method, we will vary the coefficients of all transformation matrices to minimize the sum of the squared distances between the images of each pair of selected key points across the entire panorama:

$$\sum_{(j,j')} \sum_{i} ||H_j p_i - H_{j'} p'_i||^2 \to min,$$

where i iterates over all selected key points for the pair of adjacent images and (j, j') represent all possible pairs of adjacent images with corresponding homography matrices H_j and $H_{j'}$.

To solve the least squares problem, we use the Levenberg-Marquardt algorithm (Moré, 2006) and consider the homographies obtained in the previous stage as the initial approximation.

3.6 Seam removal

After completing the previous steps, that perform image registration, we obtain the correct positioning of all the source images in the panorama. The final important step is to determine the pixel values of the panorama in the overlapping areas where two or more source images overlap. This step directly affects the visual perception of the panorama.

In this work we implemented a rather simple method that does not change the colors and the brightness of the images and only finds the least significant continuous seam line in the region of image intersection. The implemented method was inspired by (Agarwala et al., 2004) designed for photo montage problems.

The task comes down to labeling the overlapping region of two images into two classes: 0 and 1 (pixels from the first image are labeled 0, and pixels from the second image are labeled 1). To implement this approach, we used the Graph Cuts (Boykov and Kolmogorov, 2004). Let’s represent the overlapping area of the two images (specifically, the left and right images) as a 4-connected graph, where the pixels of the panorama in the overlapping region are terminal vertices, and two non-terminal vertices correspond to labels 0 and 1. The edge weights are defined as follows:

$$E(s,t) = ||I_1(s) - I_2(s)|| + ||I_1(t) - I_2(t)|| + \\ + ||\nabla I_1(s) - \nabla I_2(s)|| + ||\nabla I_1(t) - \nabla I_2(t)||,$$

where s, t are neighboring pixels, I_1, I_2 are the left and right images, and ∇I_i is the gradient magnitude of the image.

Additionally, edges connecting the graph to non-terminal vertices need to be defined. Pixels on the left side of the overlapping region are connected by edges with infinite weights to one non-terminal vertex, while pixels on the right side are connected to the other non-terminal vertex. The Graph Cuts algorithm

is then applied to the constructed graph. It finds the minimum energy cut and removes some edges to leave two independent subgraphs connected to their respective non-terminal vertices. Pixels remaining in the left subgraph are labeled 0, and pixels in the right subgraph are labeled 1. The considered images can then be merged by filling the overlapping region in the way that pixels labeled with 0 are taken from the left image and pixels labeled with 1 are taken from the right image.

An example of calculated seam for two adjacent images in a row is shown in Fig.5.

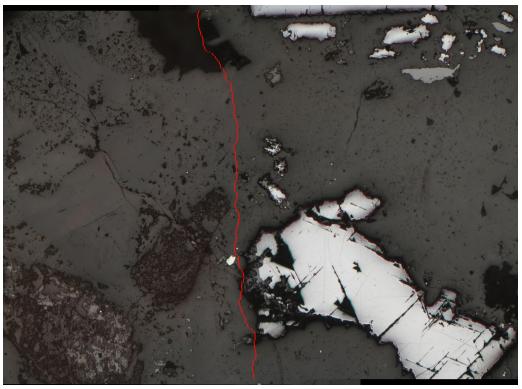

(a) Constructed seam

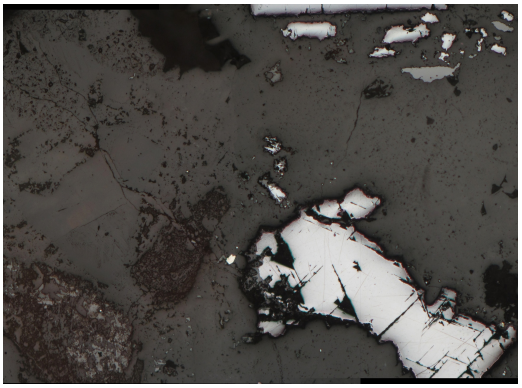

(b) Resulting combination of images

Figure 5. Example of constructing a seam for a pair of adjacent images.

The described approach is applied sequentially to all images in each row of the panorama, removing all vertical seams. Then, the same approach is applied to the resulting rows, removing all horizontal seams.

4. Results and comparison

During the development of the method, it is useful to numerically evaluate the quality of the panorama stitching. One of the few available methods, which we used in this work, is calculating the root mean square error (RMSE) of the distance of the used key points from their corresponding matched points in the finished panorama:

$$RMSE = \sqrt{\frac{\sum_{(j,j')} \sum_i ||H_j p_i - H_{j'} p'_i||^2}{n}},$$

where i iterates over all selected key points for the pair of adjacent images and (j, j') represent all possible pairs of adjacent images with corresponding homography matrices H_j and $H_{j'}$ and n is the number of such pairs. The smaller is the resulting value, the higher is the quality of the assembled panorama. This can be seen with the naked eye and is additionally confirmed by an expert geologist. This approach is valid in the considered task since the characteristics of the images of polished sections lead to the accurate and stable key points detection.

It is worth noting that in this work, the quality of the panoramas, as expressed in RMSE, was significantly influenced by the distortion correction step (which reduced RMSE from 25 to 3.561), the refinement step (which reduced RMSE from 3.561 to 1.619) and distortion correction step (which further reduced RMSE from 1.619 to 0.317). The resulting RMSE value of 0.317 indicates that the images are aligned with subpixel accuracy enabling the practical use of the developed method.

The final results of the proposed method are shown in Fig. 6, 7. We also compared the results of panorama stitching of the proposed method with Fiji’s stitching plugin (Preibisch et al., 2009) and two paid software solutions: (Adobe Photoshop Software, 2024) and (tshsoft PanoramaStudio, 2024) for two panoramas from LumenStone P1v1. For Photoshop we tested two different settings of panorama stitching (with and without automatic distortion correction), PanoramaStudio was applied in fully automatic way.

As it can be seen from Fig.8 all solutions produce artifacts, some of which are more significant, some less. Photoshop without distortion correction makes rough errors in image registration, Photoshop with enabled distortion correction does much better but nevertheless can make small image registration errors. Fiji performs good in image registration, but is not able to mask the seam between images and averages pixels intensity in intersection region, which leads to significant blur. PanoramaStudio performs well in registration and seam removal, but changes the overall brightness of the image. The proposed method is not always good enough with seam removal since the hue value of the source images varies from image to image. This is a common issue for microscopic geological images of polished sections (e.g. it can be easily seen in the panoramic images of sections an (Virtual Microscope, 2024). The reason for the color shift is the uneven distribution of differently reflective minerals in the rock and, consequently, in each photographic area. This leads to incorrect color perception by the camera matrix every time the content and composition of ore minerals changes. The best way to deal with this problem is apply recolorization as it is done by Photoshop and PanoramaStudio, however this should be done with great care, since colors play very significant role in mineral identification, especially for automatic methods.

It is worth noting that the artifacts shown in the Fig.8 are quite rare and, in general, all methods involved in the comparison performed well.

Among the existing methods discussed, PanoramaStudio and Photoshop are the most suitable for working with polished images. Fiji, due to excessive blurring of images during stitching, cannot be used for further analysis of polished sections images. However, Photoshop and PanoramaStudio have their drawbacks in ore petrography, related to their versatility. They can also be too liberal in modifying image colors, which is not good for subsequent automatic analysis (Indychko et al., 2022). Therefore, developing our own specialized software is necessary for the integration of visual geological methods when analyzing large amounts of information about ore deposits. Larger images

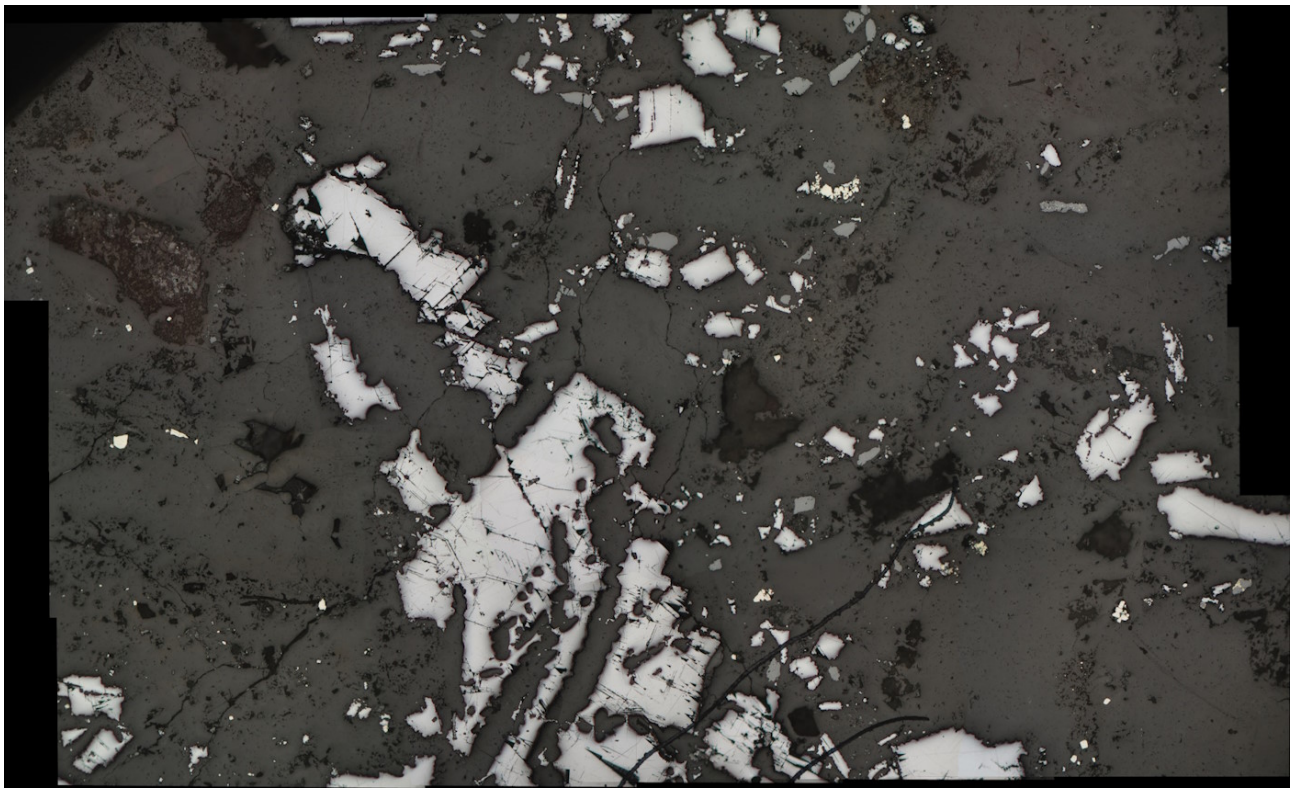

Figure 6. Example of resulting panorama stitched with the proposed method. LumenStone P1v1 dataset, panorama 1 (4×5).

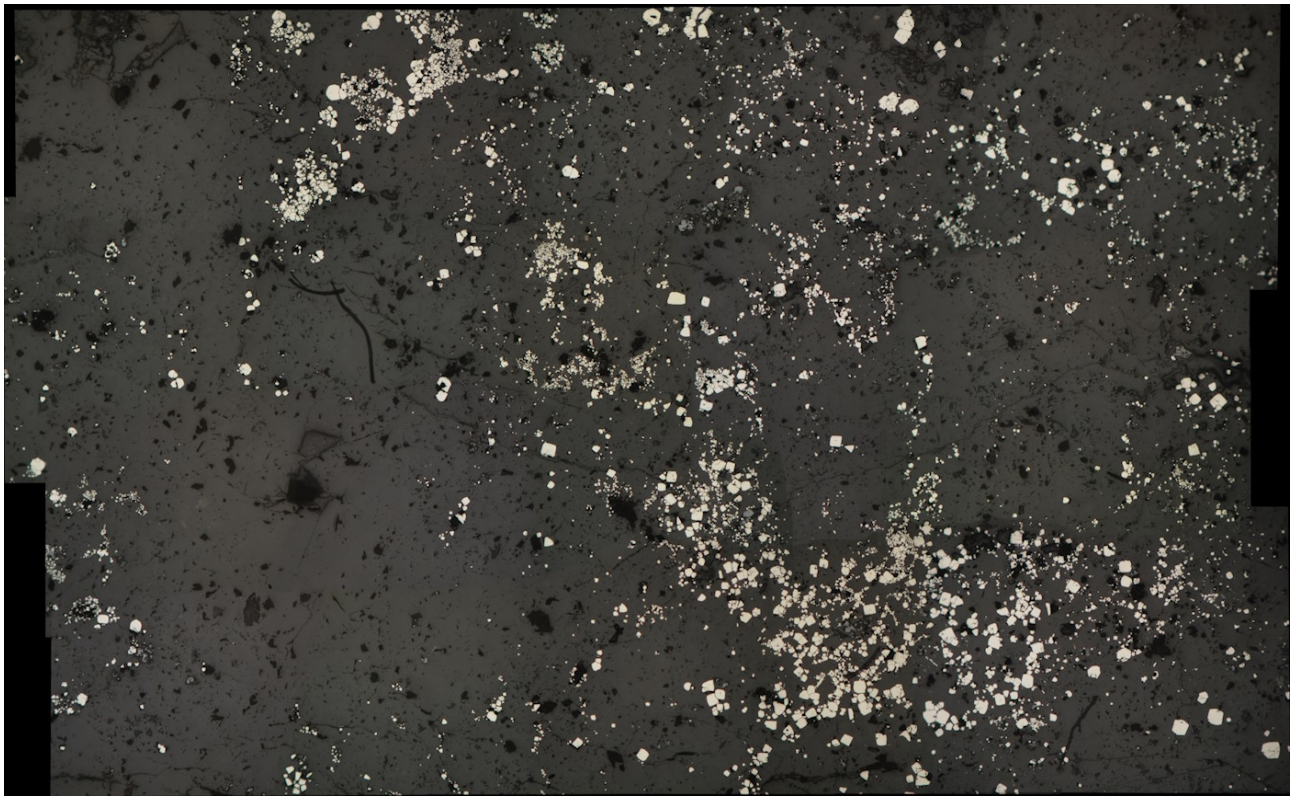

Figure 7. Example of resulting panorama stitched with the proposed method. LumenStone P1v1 dataset, panorama 2 (4×5).

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-2-W1-2024-39-2024 | © Author(s) 2024. CC BY 4.0 License.

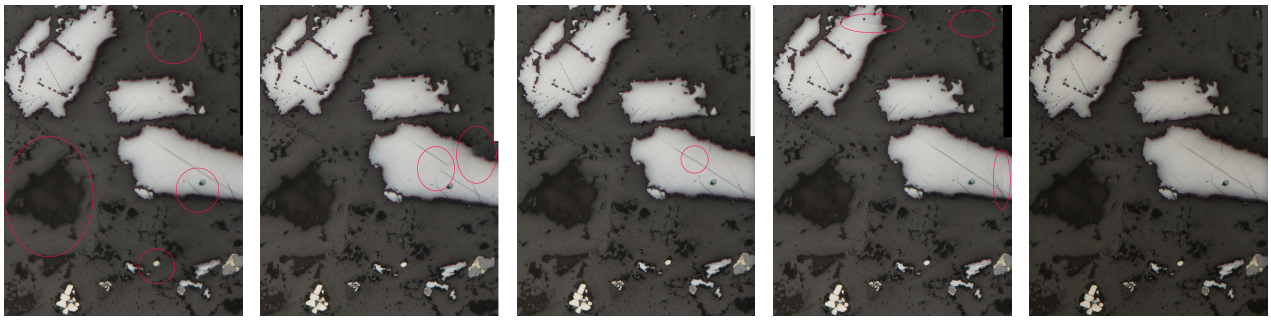

(a) Fiji (b) Adobe Photoshop without distortion correction (c) Adobe Photoshop with distortion correction (d) Proposed method (e) PanoramaStudio

Figure 8. Enlarged panorama fragment assembled with different methods (LumenStone P1v1 dataset, panorama 1). Common artifacts: (a) blurry image intersection region, (b) serious errors in image stitching, (c) minor errors in image stitching, (d) discoloration of the seam line, (e) image brightness shift.

(panoramas) are more statistically representative for automatic analysis of mineral composition (Khvostikov et al., 2021) and for the creation of virtual collections.

5. Implementation details

The proposed method of panorama stitching was implemented in Python 3 and is available as a separate module inside python `petroscope` package [3]. The package is free to use, the included docs provide instructions on setting up the package and running the proposed method for a set of images.

All the obtained panoramas from the LumenStone P1 dataset are also publicly available in PathScribe software (Khvostikov et al., 2023) as a part of larger collection used for educational purposes at the Faculty of Geology, Lomonosov Moscow State University.

6. Conclusion

In this paper we presented a new method for stitching the panoramas from the microscopic images of geological polished sections and the corresponding publicly available dataset LumenStone P1v1.

The obtained results are comparable to the paid industrial solutions (e.g. PanoramaStudio, Adobe Photoshop) and can be used by geologists in order to better obtain an representative overview of the entire surface of the sample. The accuracy of panorama stitching (calculated for the registration part of the method) is 0.33 by RMSE for the LumenStone P1v1 dataset.

The main problem of the proposed method is noticeable seams when stitching images that differ in color hue. This problem can be only solved with recolorization of the source images, but has to be performed very carefully to keep the opportunity of automatic segmentation of minerals. This will be the main direction for our future research. Besides that the future work will be focused on automatic distortion correction, automatic ordering of images for panorama stitching and general improvements on the performance of the method.

7. Acknowledgements

This work was supported by the Russian Science Foundation, project no. 24-21-00061.

[3] `https://github.com/xubiker/petroscope`

References

Adobe Photoshop Software, 2024. https://www.adobe.com/products/photoshop.html (13 June 2024).

Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D., Cohen, M., 2004. Interactive digital photomontage. *ACM SIGGRAPH 2004 Papers*, 294–302.

Aperio Digital Pathology Slide Scanners, 2024. https://www.leicabiosystems.com/digital-pathology/scan/ (13 June 2024).

Bay, H., Tuytelaars, T., Van Gool, L., 2006. SURF: Speeded up robust features. *Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006. Proceedings, Part I 9*, Springer, 404–417.

Boykov, Y., Kolmogorov, V., 2004. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE transactions on pattern analysis and machine intelligence*, 26(9), 1124–1137.

Fischler, M. A., Bolles, R. C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 381–395.

Indychko, O., Khvostikov, A., Korshunov, D., Boguslavskii, M., Krylov, A., 2022. Color Adaptation in Images of Polished Sections of Geological Specimens. *Computational Mathematics and Modeling*, 33(4), 487–500.

Jakubović, A., Velagić, J., 2018. Image feature matching and object detection using brute-force matchers. *2018 International Symposium ELMAR*, IEEE, 83–86.

Khvostikov, A., Ippolitov, V., Krylov, A., Mikhailov, I., Malkov, P., 2023. Pathscribe: new software to work with whole slide histological images for education and research. *Proceedings of the 2023 8th International Conference on Biomedical Imaging, Signal Processing*, 63–70.

Khvostikov, A., Korshunov, D., Krylov, A., Boguslavskiy, M., 2021. Automatic identification of minerals in images of polished sections. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 44, 113–118.

Lowe, D. G., 1999. Object recognition from local scale-invariant features. *Proceedings of the seventh IEEE international conference on computer vision*, 2, Ieee, 1150–1157.

Lowe, D. G., 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60, 91–110.

Moré, J. J., 2006. The levenberg-marquardt algorithm: implementation and theory. *Numerical analysis: proceedings of the biennial Conference held at Dundee, June 28–July 1, 1977*, Springer, 105–116.

Preibisch, S., Saalfeld, S., Tomancak, P., 2009. Globally optimal stitching of tiled 3D microscopic image acquisitions. *Bioinformatics*, 25(11), 1463–1465.

Research Microscopy Solutions, 2024. https://www.zeiss.com/microscopy/en/products/imaging-systems/axioscan-7-for-geology.html (13 June 2024).

Ro, S.-H., Kim, S.-H., 2021. An image stitching algorithm for the mineralogical analysis. *Minerals Engineering*, 169, 106968.

Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X., 2021. LoFTR: Detector-Free Local Feature Matching with Transformers. *CVPR*.

Triggs, B., McLauchlan, P. F., Hartley, R. I., Fitzgibbon, A. W., 2000. Bundle adjustment—a modern synthesis. *Vision Algorithms: Theory and Practice: International Workshop on Vision Algorithms Corfu, Greece, September 21–22, 1999 Proceedings*, Springer, 298–372.

tshsoft PanoramaStudio, 2024. https://www.tshsoft.de/en/index (13 June 2024).

Virtual Microscope, 2024. https://www.virtualmicroscope.org (13 June 2024).

Virtual Petrography, 2024. https://planetearth.utsc.utoronto.ca/VirtualMic/ (13 June 2024).

Zhang, J., Subasinghe, N., 2012. Extracting ore texture information using image analysis. *Mineral Processing and Extractive Metallurgy*, 121(3), 123–130.

Zukić, D., Jackson, M., Dimiduk, D., Donegan, S., Groeber, M., McCormick, M., 2021. ITKMontage: a software module for image stitching. *Integrating Materials and Manufacturing Innovation*, 10, 115–124.