

Martijn Meijers¹, Jules Schoonman²

Mapping the Edge: A Novel Approach to Georeferencing Historical Map Series

Keywords: IIIF (International Image Interoperability Framework), georeferencing, historic maps, map series, Allmaps

Summary: Libraries and archives across the globe have adopted the International Image Interoperability Framework (IIIF) to make high-resolution images of their digitized map collections available to the public. These images are often not (yet) georeferenced, which limits their findability, usability for spatial analysis and integration with other geographic data. IIIF maps can be manually georeferenced in Allmaps Editor by masking images and indicating Ground Control Points (GCPs). While effective, this process is also time-consuming and repetitive, especially when dealing with large map series, consisting of many sheets following a uniform layout. We therefore propose a novel and semi-automated method for georeferencing map sheets of such map series. The proposed approach employs computer vision techniques to detect the neat lines of map sheets. The crossing of these neat lines often corresponds to well-defined geographic coordinates in a sheet index, making them ideal candidates for georeferencing. By harnessing the inherent structure of map series, our approach reduces the manual labour involved in the georeferencing process. The method is adaptable and capable of handling slight variations in the layout of sheets. We demonstrate the effectiveness of our approach through a series of case studies involving different map series from various institutions worldwide. These results highlight the potential of our method to enhance the digital accessibility and re-usability of historical map collections.

Introduction

When making available digital collections of scanned maps, a key step is georeferencing. This means assigning real-world locations (latitude and longitude) to features on a scanned map image. Georeferencing is crucial because it allows these maps to be used with other geographic data and tools. Considerable attention has been devoted to the topic with various methods being applied (Fleet et al. 2012), (Hackeloeer et al. 2014), (Burt et al. 2019), (Luft and Schiewe 2021), (Basarić et al. 2022), (Gede et al. 2022) and different maps being used (Affek 2013), (Baiocchi et al. 2013), (Duménieu et al. 2018), (Heitzler et al. 2018). Traditionally, georeferencing is done manually, which can be very time-consuming for large collections. This research explores a novel approach that uses Georeference Annotations, which were developed to be used in conjunction with the International Image Interoperability Framework (IIIF), a standard for sharing images online. The new method is implemented in a special software tool to produce these annotations, thereby semi-automatically georeferencing large collections of map sheets. The proposed approach reduces the manual work required compared to traditional methods. Moreover, the method allows for an iterative refinement process, where users can inspect the results and adjust parameters to improve accuracy and check for errors, making adjustments as needed. We show that IIIF Georeference annotations and related tools have exciting potential for making georeferencing large map collections faster.

¹ Assistant professor, Delft University of Technology, GIS technology [b.m.meijers@tudelft.nl]

² Digital curator, Delft University of Technology, Library [j.a.schoonman@tudelft.nl]

International Image Interoperability (IIIF) and historical map collections

Libraries and archives across the globe have adopted the International Image Interoperability Framework (IIIF) to make high-resolution images of their digitized special collections available to the public, including large numbers of maps. IIIF is a set of open standards that allow high-quality digital assets to be delivered over the internet (Snydman et al. 2015).

The IIIF Image API is one of the established standards. It outlines a web service where the user can define specific parameters to retrieve a desired image. These parameters include the region, size, rotation, quality characteristics, and format of the image being requested. Moreover, the web service can be interrogated on which of these parameters (i.e. image operations) are supported. The web service can also specify a set of image sizes and image tiles, which have been pre-generated and stored (cached) and will thus yield a quick response.

Another standard is the IIIF Presentation API. This is a web service that returns documents structured in JSON-LD (a linked data format). These documents, known as Manifests, describe the structure and layout of digitized objects, or a collection of objects, and related metadata. The IIIF Presentation API is built on the Shared Canvas Data Model³, which is illustrated in Figure 1. This is a method for describing digital copies of physical objects using linked data.

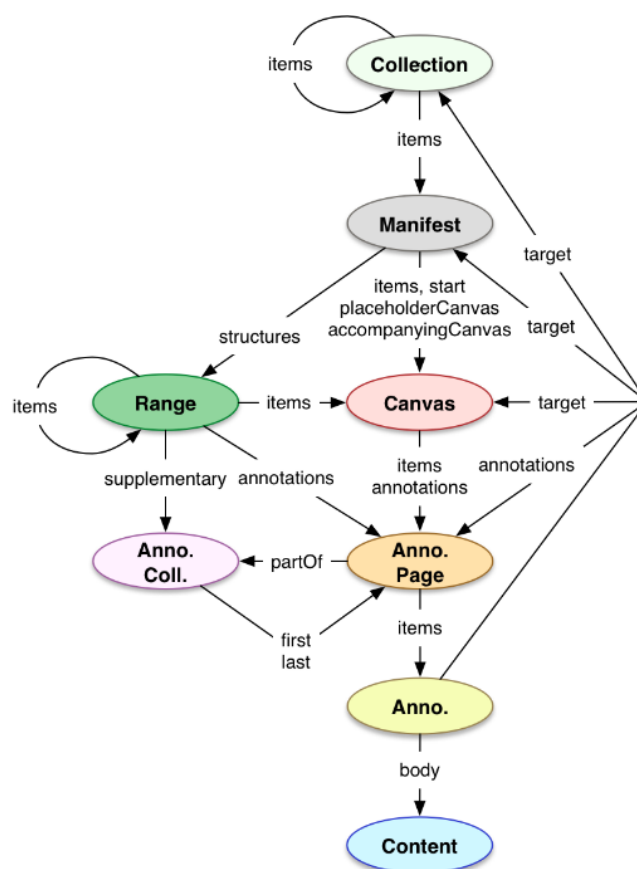


Figure 1: The concepts that the IIIF Presentation API defines and how they relate.

For example, a digitised version of a medieval manuscript can be described by a IIIF Manifest consisting of a series of canvases. Each canvas can be annotated with different types of information.

³ <https://iiif.io/api/model/shared-canvas/1.0/>

Most commonly, it references a single IIIF Image, depicting one digitised page of the manuscript. It may include more annotations, with different purposes. Annotations point to the entire canvas or a specific region of it and can provide additional information useful for visualizing the resource. For instance, text obtained from an Optical Character Recognition (OCR) process can be added as annotations to the respective canvases of the manuscript. This allows clients to overlay image and text or build a search index. There are several tools available that can interpret IIIF Manifests and display different types of annotations, such as Mirador⁴ or Universal Viewer⁵. Figure 2 illustrates the use of an IIIF Manifest of a historic map series in one of these tools.

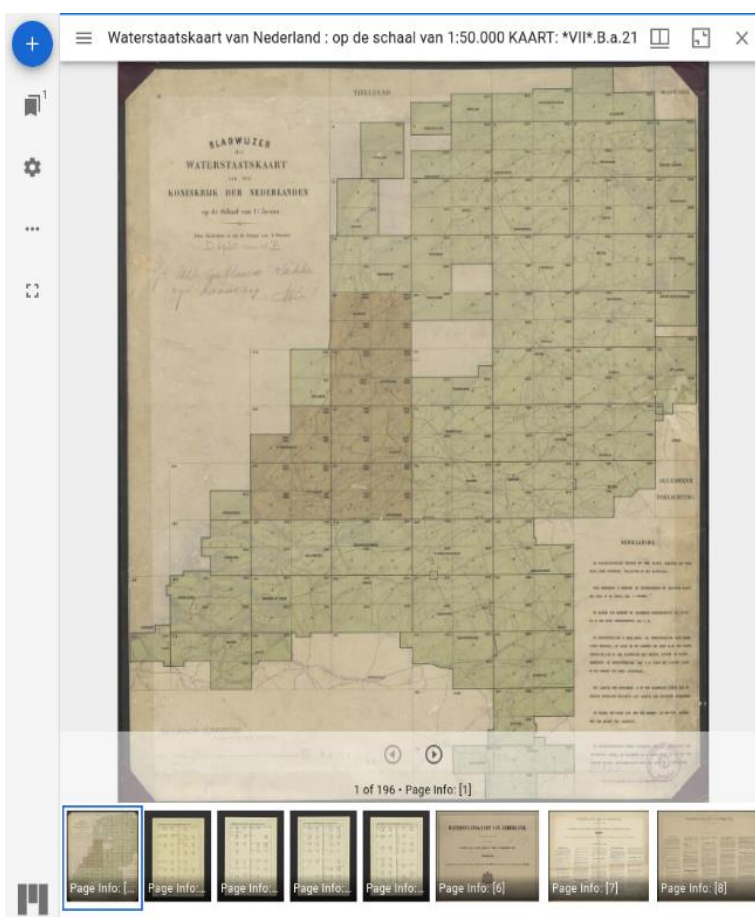


Figure 2: IIIF manifest from University Utrecht Library - Special Collections (describing the separate map sheets of the Waterstaatskaart) loaded in Mirador⁶.

Georeferencing IIIF maps

Among the digital collections of IIIF-supporting institutions worldwide, numerous maps and atlases can be found. Like for books, IIIF Manifests are available for cartographic resources that provide access to the available scans and metadata. Next to bibliographical metadata, additional geographic information might be available for such resources, but until recently there was no systematic way to add this to a IIIF Manifest. To fill this gap, the Georeference Extension⁷ was published by the IIIF Consortium in 2023.

⁴ <https://projectmirador.org/>

⁵ <http://universalviewer.io/>

⁶ <https://projectmirador.org/embed/?manifest=https://objects.library.uu.nl/manifest/iiif/v2/1874-389916>

⁷ <https://iiif.io/api/extension/georef/>

This extension to the IIIF Presentation API indicates how to georeference a IIIF Canvas, which means mapping the image coordinates of the resource (u,v) to geographic coordinates (ϕ,λ) in the WGS'84 datum. To be able to project from one coordinate system to another, two things are necessary: a set of corresponding points, with known coordinates in both systems, and a given mathematical transformation using these points. The corresponding points are termed Ground Control Points (GCPs), and for the transformation type different options are available. Such a mathematical transformation can be exact or inexact at these GCPs; An example of an exact transformation is Thin-Plate-Spline, an example of an inexact transformation is a polynomial transformation. When more control points are available than strictly necessary to specify the transformation, this redundancy –the additional points– can be used to make statements about the quality of the transformation (whether and how much points agree or deviate). In addition to the GCPs and the transformation type, the extension outlines how to specify a mask, which indicates the region of the image to interpret as a map, as there might be non-cartographic information around the map (such as explanatory notes) or multiple map fragments present in a single digital image (for example, in the case of an inset map). All this information can be encoded in a JSON-LD (JSON for Linking Data) Georeference Annotation.

Allmaps

In conjunction with the publication of the Georeference Extension, new open-source tools became available to produce and consume the Georeference Annotations under the name of Allmaps⁸. In Allmaps Editor⁹, a IIIF Presentation Manifest of a cartographic resource can be loaded by pasting its URL. After adding the mask and the set of Ground Control Points for selected canvases in a visual editor, a Georeference Annotation can be copied or downloaded. The annotation is also stored on Allmaps's annotation server and can be directly loaded in Allmaps Viewer¹⁰, which overlays the IIIF Image on a web map. The application fetches the corresponding image tiles from the original IIIF image server, where the digital image is stored. Using the GCPs and the specified transformation type, Allmaps Viewer can then adjust the images to fit over a world map (in Web Mercator projection). This adjustment, or 'warping', is performed directly on the user's device (client-side) by transforming the fetched image tiles (using WebGL). This means that no server-side processes are used for the heavy lifting of image transformation.

Warping the image is one use case among many. Allmaps' modules also support the reverse transformation and converting geographical information 'back' to pixel coordinates. This allows clients to render geographical data on top of the original, unwarped image, e.g. drawing a tracked user's path of geographical positions on top of the original image. Georeference Annotations can also be used to build a spatial search index for exploring a collection of maps.

While the georeferencing process is effective for making the IIIF images also available as geographic information and increasing their usability, the manual annotation process can be time-consuming and repetitive. Especially when dealing with large map series consisting of many sheets following a uniform layout, where prior knowledge is present about the coordinates of the individual map sheets.

⁸ <https://allmaps.org>

⁹ <https://editor.allmaps.org/>

¹⁰ <https://viewer.allmaps.org/>

The role of the sheet index

Historically, map sheets that are part of a series often are laid out as non-overlapping quadrangles. These quadrangles not only defined the layout of the map, but also served as individual units for terrain surveying (Bakker et al. 2015: 13). In The Netherlands, for instance, a variety of map series have been created for various purposes, including topographic mapping, water management, and geological studies. Figure 3 provides three example sheet indexes for Dutch historic map series with the layout of the quadrangles.



Figure 3: Sheet indexes of: (left) Geological Map, (middle) Water Management Map, (right) Topographic Military Map

The sheet index, describing the location of each map sheet, and the corner coordinates of the sheets can be used as a starting point for georeferencing a map series. With the sheet index, the world coordinates of the individual map sheets can be derived and listed easily. Figure 4 gives an example of sheets, where the world coordinates are written near the corners. For georeferencing the sheets, this is already half of the information required for determining the GCPs. The corresponding image coordinates (u,v) for each corner is thus what we aim to find out automatically, instead of manually.

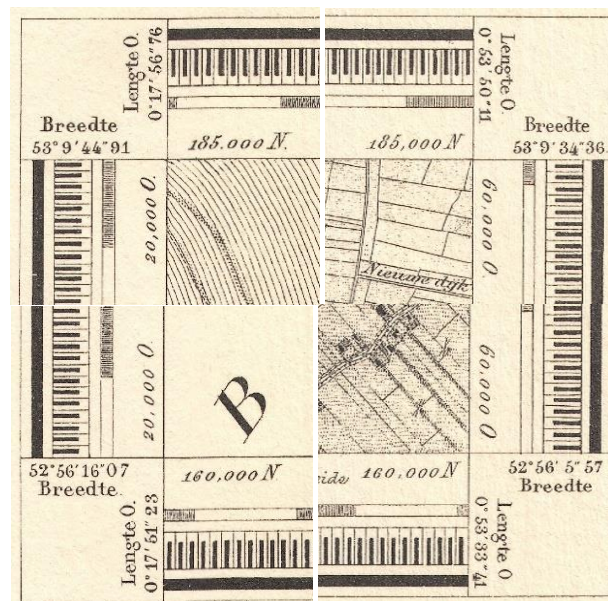


Figure 4: An example of a map sheet from the Dutch Topographic Military Map (Sheet 10, Sneek), where near the corners (top-left, top-right, bottom-left and bottom-right) the exact coordinates of the sheet are mentioned, which corresponds to the known sheet size in kilometer: 40×25 km

MapEdge

To georeference semi-automatically a series of map sheets with MapEdge the following steps are taken:

1. Digitize the sheet index and link the sheet geometries with known world coordinates (φ, λ) with the related IIF Image endpoints.
2. Set global parameters related to the layout of the series and the rims.
3. Run MapEdge, given this sheet index, on each map sheet individually. The IIF Image endpoint URL gives access to the digital image resource of the sheet. This step produces the image coordinates (u, v) of the Ground Control Points, together with the mask for the sheet (which part of the image to consider as map fragment).
4. Check the output produced by MapEdge. If we are satisfied, we keep the output (as Georeference Annotation). If we are dissatisfied, we change the used parameters of MapEdge and run MapEdge again on the sheet, until satisfied (or we decide to manually process the sheet).

Figure 5 illustrates the workflow steps that we have implemented in MapEdge in more detail.

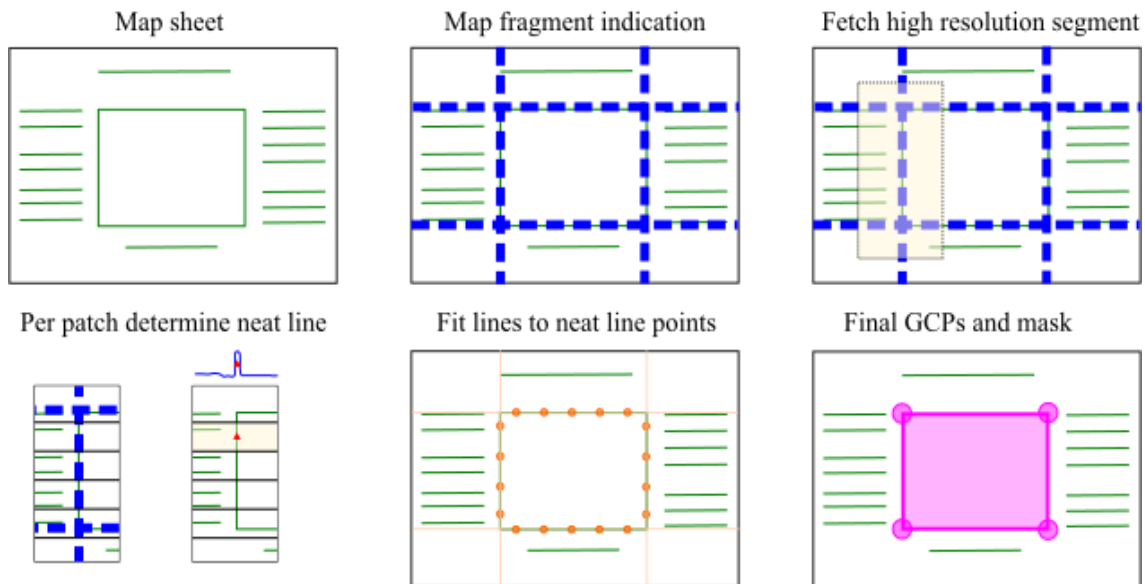


Figure 5: MapEdge workflow steps

Per map sheet the following steps are carried out: MapEdge reads the user settings specified to use for processing a map sheet. An overview image, at low resolution, is retrieved from the IIF Image endpoint, and a rough indication is obtained of where the map fragment is placed on the map sheet. Note, the IIF Image API protocol permits us to retrieve the image at multiple resolutions. The rough indication of the map fragment is obtained by running the image through a series of image conversion steps, which leads to a black and white image. We then use a 1D histogram approach to obtain the indication of the rims' position. We count how many of the pixels along each axis are black. We threshold the obtained histogram, which then gives a list of peaks that potentially form the rims of the map fragment. For each peak, the centre and the width are obtained. In the user settings, the expected characteristics of these peaks are described a-priori (these depend on the layout of the map frame). Subsequently, all identified peaks are examined in pairs. These pairs are then ranked based on a fuzzy metric (Zadeh 1965) that considers the distance between each pair

and their proximity to the location specified in the user settings. Figure 6 shows an example of a detected map fragment.

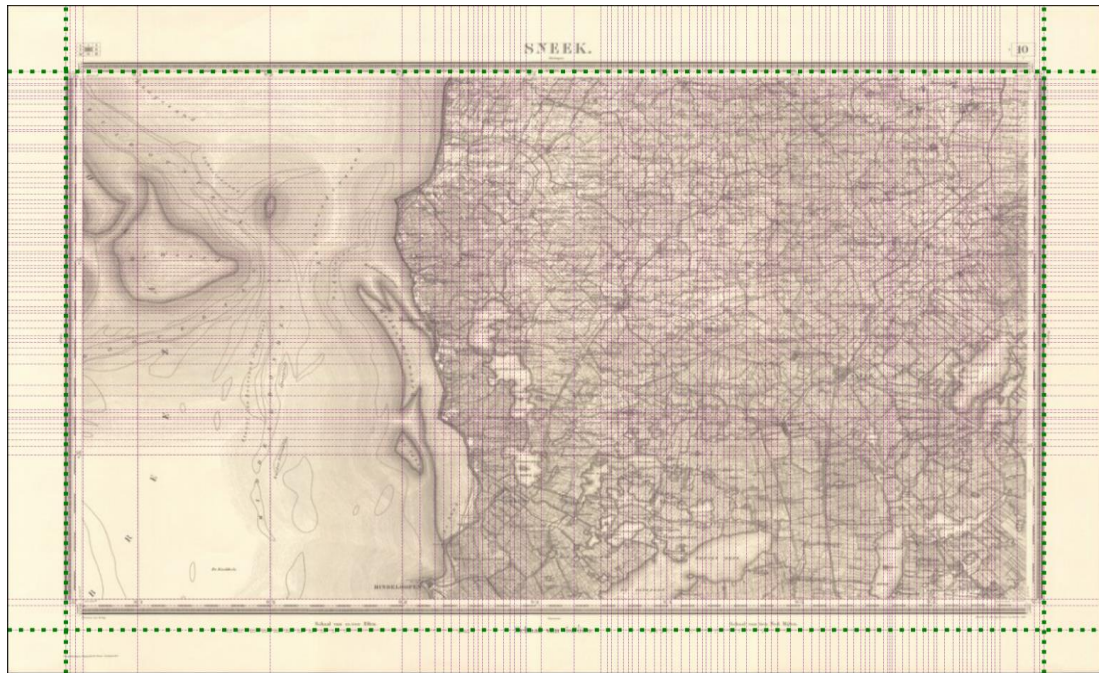


Figure 6: Approximate rectangle detected; the green stipple lines indicate the final pairs used

For the majority of map sheets, the user configurations tend to be consistent. However, variations may arise due to the positioning on the sheet or the dimensions of a specific sheet. To accommodate these variations, we have implemented an ‘overwrite mode’ for the settings. In this mode, a generic configuration is employed as the default. However, if a more specific configuration is provided, it supersedes the default setting for that particular sheet. This approach ensures flexibility and customization while maintaining a consistent base configuration.

With the rough indication (rectangle), a high-resolution segment of the map sheet, specifically around the edges of the indicated rectangle, can be fetched from the IIIF Image endpoint to minimize the amount of image data transferred. We again apply a conversion from colour to a black and white image (and, if needed, we improve the contrast of the image). Subsequently this rectangular fragment is subdivided in small, non-overlapping rectangular patches. For each patch, we search for a good indication of the location of the frame lines of the map sheet.

To do this, we obtain a 1D histogram (for the number of black pixels) along an axis of every patch and rank the relevant peak(s) that corresponds to the black line(s) of the map frame. A difficulty we have to solve is that such peak characteristics can slightly vary from patch to patch. We let the user specify in a fuzzy way the extents of the peak (e.g. the width of the black rim we are looking for is a best match if it is between 6 and 8 pixels, but we also accept peaks between 4 and 10 pixels wide). For the case that a map frame around the map sheet consists of multiple neat lines, we added the option to find a pair of lines at the rim of the map fragment in tandem. Next to the individual sizes (width) of both lines in pixels, also the distance between the pair of lines is then added in the user settings and considered in the ranking process. Figure 7 gives an example.

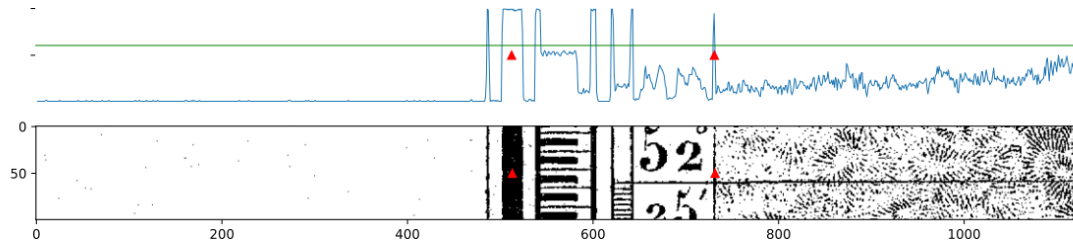


Figure 7: Thresholding (green line) 1D histogram of the number of black pixels (blue line) gives peak candidates, that are ranked according to their likelihood of being the rim of the map sheet we are looking for. The red triangles indicate the chosen peaks.

Once the centre of the peak is determined, a 2D point can be placed back on the original image (by applying a series of linear transforms), leading to a set of sampled 2D points along each side of the map sheet.

The points thus found are processed with a random sample consensus (RANSAC) algorithm (Fischler and Bolles 1981). We fit a straight line through the points belonging to one of the four neat lines (O'Leary et al. 2004). We select randomly a subset of the points from all rim points, that create the line. For all points associated with a rim, we then calculate the perpendicular distance to the fitted line, and we determine the number of points that are closer to the line than a specified threshold (obtained from the settings). These points are called the inliers. This process is repeated several times, while remembering which fitted line so far yields the highest number of inliers. This way, we end up with 4 fitted lines (per map fragment side one line).

The intersection of the fitted lines results in four corners. In our case, the intersected points serve two purposes. They form the (u,v) image coordinates of the GCPs and they also make up the corners of the polygonal mask. This mask determines which part of the image will be transformed and displayed in the map viewer.

When the corners are found, we perform some consistency checks on the rim points harvested by MapEdge. Firstly, we verify if a sufficient number of points have been sampled along each rim and if the RANSAC algorithm has detected an adequate number of inliers. Secondly, following the determination of the corner points, we conduct an internal consistency check on the sides of the map frame. The expectation is that the opposite sides should be of equal length, as should both diagonals. For each pair of sides, we examine the ratio of the smaller distance to the larger distance. If the sides are of equal length, this ratio should be 1. Thirdly, we perform an external consistency check against the known dimensions of the map sheets. Given that the size of the map fragments (e.g. in centimetres) is documented, we can combine this information with the scan resolution, expressed in dots per inch, to calculate the expected size of the sheets in pixels. This allows us to evaluate the deviation of the side distances from this expected value. Fourthly, we check the distribution of the perpendicular distances of the inlier points against the fitted line. This helps us see if a straight line is a good way to describe the edge of the rim. If the paper of the sheet is bent out of shape and the points are on a curved line (either bulging out or caving in), the distances of these points from the straight line will be different than if they were on a straight line. This can give us a clue about whether and how the paper of the scanned sheet is deformed. Lastly, we visualize all corners detected, together with the respective IIF image parts around the corners. This makes it easy to browse through all detected map corners. All checks together ensure a systematic assessment of the obtained results.

Results

We have implemented MapEdge using the Python programming language. The source code of MapEdge is available online¹¹. The main libraries we used are OpenCV (for image processing operations), Numpy for the line fitting approach (linear algebra operations) and Requests, for retrieving the data from the IIIF services.

Table 1 gives details on the map series we have processed with MapEdge: the Dutch ‘Topografisch Militaire Kaart’, the first edition of the Dutch ‘Waterstaatskaart’ and a Japanese map for North Korea. The table also specifies the repository where the scans are stored as IIIF images, the intended map scale of the map series, and the estimated production date of the maps. For the Dutch series, digital sheet indices were not yet available and we had to reconstruct these based on the available historical documentation¹².

Table 1 Processed map series.

Map series	Owner IIIF server	Sheets	Date	Scale
Topografisch Militaire Kaart ¹³	Delft University of Technology	65	±1850	1:50,000
North Korea ¹⁴	Stanford	308	±1916	1:50,000
Waterstaatskaart (first edition) ¹⁵	Utrecht University	184	±1860	1:50,000

The processing of the sheets took on average 15 seconds per map sheet, including the transmission of the image data (note, this is much dependent on the performance of the IIIF Image server and the network bandwidth available).

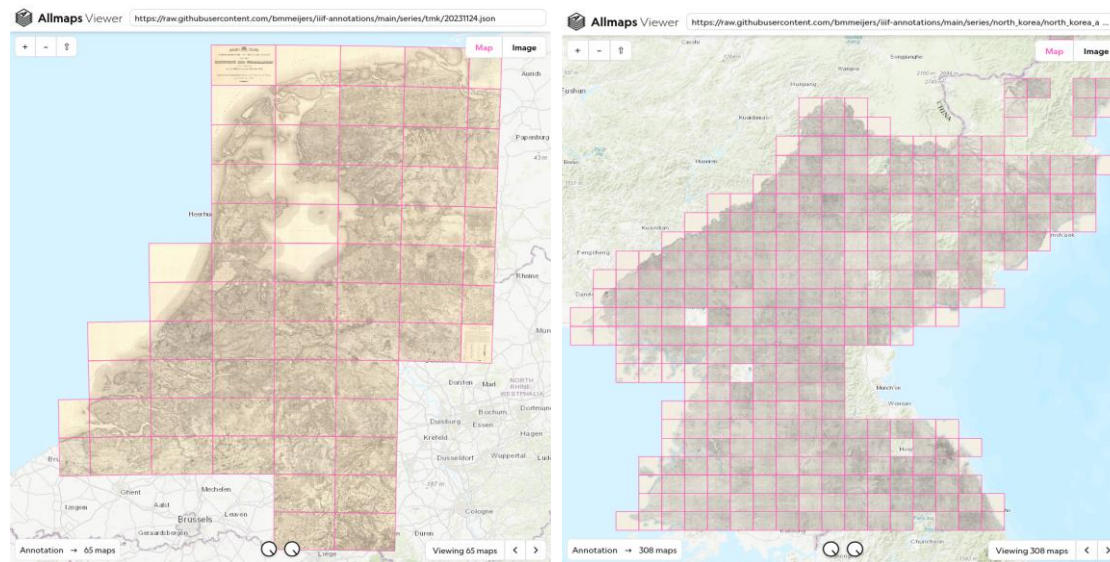


Figure 8: Complete map series visualized in the Allmaps viewer, after processing with MapEdge: (left) Water management map (right) North Korea

¹¹ <https://github.com/bmmeijers/mapedge/>

¹² <https://observablehq.com/@tudelft/sheet-indices>

¹³ <https://viewer.allmaps.org/?url=https://raw.githubusercontent.com/bmmeijers/iiif-annotations/main/series/tmk/20231124.json>

¹⁴ https://viewer.allmaps.org/?url=https://raw.githubusercontent.com/bmmeijers/iiif-annotations/main/series/north_korea/north_korea_annotation_page.json

¹⁵ https://viewer.allmaps.org/?url=https://raw.githubusercontent.com/bmmeijers/iiif-annotations/main/series/waterstaatskaart/uu/editie_1/latest.json

Figure 8 shows two of the map series we processed displayed in the viewer of the Allmaps platform. The Allmaps viewer, with its default features, allows one to examine the composite of the map series. Additionally, it provides a separate, linked view where one can inspect the original, high-resolution map sheet images in their original, unwarped form, with the polygonal mask from the Georeference Annotation superimposed.

After producing the Georeference Annotations, we performed our proposed checks. For the Topografisch Militaire Kaart (topographic, military map), these checks revealed that for this series MapEdge was able to perform extensive sampling, with numerous points identified along all edges of the map fragments (attributed to the clean scanned images). Size differences between left/right and top/bottom sides respectively were in the order of magnitude of 10 pixels horizontally and 6 pixels vertically and were assumed to be acceptable (given the size of the fragments scanned at 600 dpi, 12500 pixels wide, 14800 pixels height).

Upon visually inspecting the detected points in combination with the image, we observed that all detected corners aligned with the corners of the neat lines. In certain instances, this alignment was exact (pixel perfect), while in others, there was a minor offset of a few pixels. We attributed this deviation to distortions in the paper and inherent imperfections arising from the scanning process.

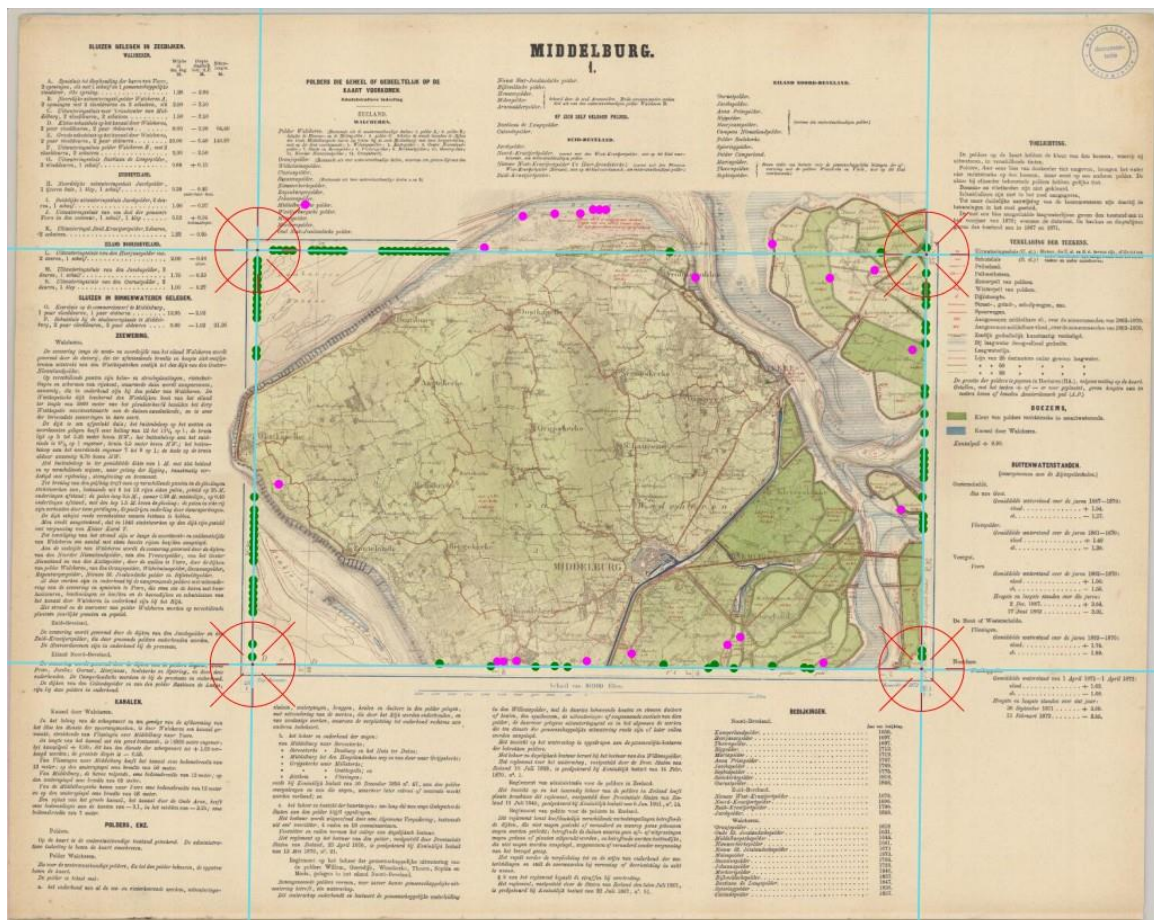


Figure 9: A processed sheet (Middelburg 1) of the Dutch historic water management map. In dark green the points found along the rim, which were accepted as inliers. The pink points are considered outliers. In blue the fitted lines with the RANSAC algorithm. The red crop marks indicate the intersection points of the lines, forming the GCPs in image coordinates and the corners of the polygonal mask for the sheet. Note, a neat line is partially absent around the top right corner.

For the Waterstaatskaart (water management map), we found another interesting case. Figure 9 shows that with our approach we can detect the corners and can produce a correct Georeference

Annotation, even though the map fragment sticks out partially from the map frame. Sampled points were missing around the top-right corner, but for the corner the image coordinates could be identified by fitting the two lines based on the other (far away) points found along the top and right neat lines. Furthermore, performing the quality checks on the Waterstaatskaart revealed three map sheets where exceptions in sheet size were not digitized correctly in our sheet index and had to be corrected afterwards. Another issue revealed was that for 3 of the 184 sheets neat lines were missing around (a part of) the map frame, and corners could not be deduced. For these particular sheets we resorted to manually doing the work in Allmaps Editor. Moreover, as for some sheets map fragments were sticking out of the frame, we had to manually correct the masks in the Georeference Annotations.

In order to evaluate the reproducibility of our method, we extended its application to a series of Japanese maps of North Korea. These IIF images, published by Stanford University, were used together with an obtained digital sheet index. Given the similarity in the layout across the entire series and the availability of high-quality scanned images, the MapEdge tool was successful in fitting a substantial number of points along the rims to determine the corners. This demonstrates the applicability of our method when used for different map series.

Discussion of obtained results

Although the results we obtained for the map series we processed are promising, the following is a list of items that we like to investigate and further improve:

- The algorithm's detected corners may deviate slightly from the exact corners due to distortions in sheets. To achieve pixel-perfect accuracy, we could rerun the main algorithm of MapEdge (sample neat line points, fit lines, obtain intersection points) around the detected corners to obtain an even better estimate of the corners. We can also explore other techniques, such as pattern matching. The IIF Image API is helpful to obtain high resolution image data around the detected corners.
- Considering each neat line as two separate halves and estimating the location of the straight line for each half separately could potentially reduce the influence of distant points (that reside near the opposite corners).
- The method of finding points along the neat lines is robust against some noise in the peak points we derived. However, confusion can still occur, for instance, when map features inside the map fragment close to the rim share similar characteristics with the rim side. For example, linear features, such as graticule lines, can occupy the same number of black pixels as the neat line.
- Although MapEdge obtains a relatively low number of control points (albeit from a large number of detected neat line points giving insights in how accurate the neat line was sampled) georeferencing could be improved further by adding also in the interior of the map fragment more control points. Our approach then serves as an initial step in this process, followed by a reverse transformation of world coordinates to image coordinates. Where map features are supposed to be located on the image (for example, survey network points present as icons for church towers or intersections of graticule lines with known world coordinates) can then be inferred from this reverse transformation.
- Our standalone Python program could be integrated into Allmaps Editor to provide interactive, browser-based georeferencing assistance.
- Questions about governance and linking Georeference Annotation information arise. While results can be stored in the Allmaps platform's annotation server, it would be beneficial to

develop mechanisms to contribute back the Georeference Annotations to the original IIF image repositories.

Concluding remarks

We have introduced MapEdge, an open-source software designed to semi-automatically generate Georeference Annotations for a series of map sheets. Our established process enhances the repeatability of image coordinate derivation for map sheets available as IIF images that display similar characteristics. The IIF Image API enables efficient image data retrieval by transferring only the relevant parts of high-resolution images, keeping the image data at the source institution. This process also facilitates the inspection of the derived corner points, enabling quality checks to ensure the results meet the required standards.

The Georeference Annotations obtained serve as a pivotal link, enabling the transition between image and world coordinates. In the context of linked data, the combination of the IIF approach and available Georeference Annotations positions the geographic location as a key connector. This allows for the integration of diverse types of historic information. Our approach therefore contributes to findability, usability for spatial analysis and integration with other geographic data of IIF maps.

References

- Affek, A. (2013). Georeferencing of historical maps using GIS, as exemplified by the Austrian military surveys of Galicia. *Geographia Polonica*, 86(4):375–390.
- Baiocchi, V., Lelo, K., Milone, M. V., and Mormile, M. (2013). Accuracy of different georeferencing strategies on historical maps of Rome. *Geographia Technica*, 1:10–16.
- Bakker, N., Altena, V., Verberck, M., and Berkhout, T. (2015). *200 Jaar Kaarten Maken in Beeld: 1815-2015*. Uitgeverij 12 Provinciën.
- Basarić, M., Jakovljević, G., and Sekulović, D. (2022). Georeferencer X – tool for automatic georeferencing of the topographic maps. *Abstracts of the ICA*, 5:1–2.
- Burt, J. E., White, J., Allord, G., Then, K. M., and Zhu, A.-X. (2019). Automated and semi-automated map georeferencing. *Cartography and Geographic Information Science*, 47(1):46–66.
- Duménieu, B., Abadie, N., and Perret, J. (2018). Assessing the planimetric accuracy of Paris atlases from the late 18th and 19th centuries. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, SAC 2018. ACM.
- Fischler, M. A. and 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.
- Fleet, C., Kowal, K. C., and Pridal, P. (2012). Georeferencer: Crowdsourced georeferencing for map library collections. *D-Lib Magazine*, 18(11/12).
- Gede, M., Kerkovits, K., and Varga, L. (2022). Automatic Georeferencing of the 1951–53 Topographic Map Series of Hungary. *e-Perimtron*, 17(4):192–202.
- Hackeloeer, A., Klasing, K., Krisp, J. M., and Meng, L. (2014). Georeferencing: A review of methods and applications. *Annals of GIS*, 20(1):61–69.
- Heitzler, M., Gkonos, C., Tsorlini, A., and Hurni, L. (2018). A modular process to improve the georeferencing of the Siegfried map. *e-Perimtron*, (13):85–100.

- Luft, J. and Schiewe, J. (2021). Automatic content-based georeferencing of historical topographic maps. *Transactions in GIS*, 25(6):2888–2906.
- O’Leary, P., Harker, M., and Zsombor-Murray, P. J. (2004). Fitting coupled geometric objects for metric vision. In Chanda, B., Chandran, S., and Davis, L. S., editors, *ICVGIP 2004, Proceedings of the Fourth Indian Conference on Computer Vision, Graphics & Image Processing, Kolkata, India, December 16-18, 2004*, pages 171–176. Allied Publishers Private Limited.
- Snydman, S., Sanderson, R., and Cramer, T. (2015). The International Image Interoperability Framework (IIIF): A community & technology approach for web-based images. *Archiving Conference*, 12(1):16–21.
- Zadeh, L. (1965). Fuzzy sets. *Information and Control*, 8(3):338–353.