Beatrice Vaienti*, Isabella di Lenardo, Frédéric Kaplan

Cartographic Stemmatology: Segmenting Local Deformation Similarities in Historical Maps of Jerusalem

Keywords: Stemmatology, Deformations, Georeferencing, Similarity, Copy, Cartometry, Jerusalem

Summary: The advancement of computational tools for cartometric analysis has opened new avenues for the identification and understanding of stemmatic relationships between historical maps through the analysis of their planimetric deformations. The 19th-century western cartographic depiction of Jerusalem serves as an exemplary case study in this context, since the difficulties in conducting comprehensive onsite surveys and the fascination towards the depiction of this area led to a proliferation of maps characterized by frequent errors, distortions, and widespread copying. How can similarities and differences between maps be measured, and what insights can be derived from these comparisons? This paper introduces a methodology aimed at detecting and segmenting regions of local similarity across maps, corresponding to the portions that were either copied between them or derived from a common source. To detect these areas, the Ground Control Points from the georeferencing process are employed to deform a common lattice grid for each map. These grids, triangulated to maintain shape rigidity, can be partitioned under conditions of geometric similarity, allowing for the segmentation and clustering of locally similar regions that represent shared areas between the maps. By integrating this segmentation with a filter on the intensity of deformation, the areas of the grid that are almost non-deformed, and thus not relevant for the study, can be excluded. Our analysis of fifty 19th-century maps of Jerusalem, the most extensive dataset computationally explored to date, reveals the effectiveness of this method in selecting areas of local similarity, even when limited. The methodology is applied to the maps in the dataset representing the Russian Compound and used as a tool to aid in the creation of the stemmata of the area's representation and to uncover new historical insight. These findings underscore the potential of computational cartometry in revealing hidden layers of cartographic knowledge and contributing to the digital genealogy of map collections.

1. Introduction

During the 19th century, Jerusalem was the subject of a large number of western maps. Early in the century, the maps of the Holy Land were typically created by individuals or small groups, mostly Christians (Levy-Rubin and Rubin 1996: 1). However, in the case of the representation of Jerusalem, the difficulties of carrying out an extensive survey of the city resulted in inaccuracies, fabrications and on reliance on copying partially or completely previous maps (Rubin 1999: 55). As the century progressed, the mapping efforts became more organized, with a subsequent improvement of their quality (Goren et al. 2017: 2). As such, the abundant western cartographic production on Jerusalem, particularly in the first half of the century, serves as an ideal case for analyzing inaccuracies as sources of hidden historical insights and to derive information on copying practices by studying their propagation.

This contribution presents a novel methodology to segment areas where local planimetric distortion is similar between maps. This methodology is then applied to derive historical insights into the cartographic depiction of one of the first quarters built outside the Old City walls, the Russian Compound. By comparing the planimetric deformations of fifty selected maps of Jerusalem from 1810 to 1872 and examining the representation of the area, it is possible to propose a hypothesis for its

* PhD student, EPFL – DHLAB, Switzerland [beatrice.vaienti@epfl.ch]

stemmata and to clarify dating issues. Moreover, the results will also be used to investigate the predecessors of a map and trace the combination of sources used in its compilation.

The proposed novel methodology is based on the analysis and comparison of the deformed lattice grids obtained from the Ground Control Points (GCPs) created for georeferencing the maps. The objective is to detect local deformative similarities between maps, even when limited in space or when a map combines information from multiple sources. In these cases, a global alignment would not be suitable as it may be influenced by the dominant source, resulting in the misalignment of smaller similar regions or an overall misalignment if the sources are comparable in size. Consequently, it is necessary to avoid relying on the global alignment of the grids being compared, as their alignment may be skewed by outliers and incorrectly place distantly areas of limited local similarity.

Therefore, the methodology relies exclusively on local geometrical properties. Specifically, it leverages the property of geometrically similar rigid shapes to maintain a constant ratio between corresponding sides. After triangulating the deformed lattice grids to ensure the rigidity of the shapes, the geometrical configuration of every square subset of nine nodes, referred to as a *token*, is encoded. The encoded nine-point configurations are then used to perform a hierarchical clustering, based on the similarity of corresponding *tokens* across the compared grids. The result of this clustering can then be used to segment regions of shared similarities across maps.

These segmented regions represent areas that, across two or more maps, share the same deformation, indicating a copying relationship between a group of maps either directly or through intermediary maps. Through this discretization of copied elements it is possible to draw a similarity with the field of stemmatology: just as stemmatology traces the copying and transformation of texts through manuscripts, this methodology traces the copying and transformation of cartographic elements through maps. This opens up the possibility of adopting the analytical instruments established by stemmatology to study the propagation of map distortions.

In the following sections, the case study, the methodology devised to analyze it, and the results will be presented. Section 2 explores the peculiarity of the cartographic depiction of Jerusalem in the 19th century and the reasons for studying the representation of the Russian Compound. Section 3 provides background on studies that used deformation analysis as a tool to gain insight into cartographic production. Section 4 details the methodology for encoding local deformations and comparing them across maps. This methodology is then applied to derive the stemmata of the representation of the Russian Compound (Section 5). Finally, Sections 6 and 7 discuss and summarize the findings and implications of this study, respectively.

2. Case study: 19th-century maps of Jerusalem and the Russian Compound

In the 19th century, Jerusalem became a focal point of cartographic interest, particularly among western travelers driven by religious and historical motivations. However, limitations in accessibility, time for surveying, and local knowledge often led to misrepresentations and distortions in the maps. Moreover, maps resulting from individual endeavors, often created by travelers not specifically trained as surveyors or cartographers, likely lacked a specific interest in achieving geodetic and planimetric accuracy, which contributed to the high level of distortions in these maps.

The tendency to copy existing sources completely or partially led to a proliferation of inaccuracies, making this dataset particularly valuable for study. Instead of dismissing these misrepresentations or correcting distortions, as one might when seeking accurate topographical information, these inaccuracies represent a worthy subject of inquiry for extracting additional historical insights. Rubin

(1990) introduced the idea of studying the maps of Jerusalem using a "cartogenealogic" method due to the large number of copies. He observed that "sometimes the similarity in marginal details, in distortions, in mistakes and errors, could prove the connection of imitation and copying more than the central details", which underlines the importance of even limited partial similarities. This "cartogenalogic" method was applied to maps of Jerusalem from the 15th to the 19th century (Rubin 1991), distinguishing between original maps (i.e., those based on a visit to Jerusalem, providing firsthand information) and copied maps (those created in Europe). Through genealogic research, the author identified families of maps, pinpointed the original map in each family, and reconstructed the chain of copies in the right order.

2.1 The Russian Compound

The Russian Compound, designed by the Russian architect Martin Ivanovich Eppinger in collaboration with his brother Fyodor Ivanovich (Vakh 2011:46), was built between 1860 and 1864 as part of a broader effort by the Russian Empire to establish its presence in the Holy Land (Gerd and Potin 2018; von Winning 2022). Intended to host Russian Orthodox pilgrims, it is among the first neighborhoods of Jerusalem built outside the city walls (Ben-Arieh 1975).
Among the selected maps, some depict the Russian Compound as early as 1859 and 1860, which is respectively before construction had even begun and when construction had just started. Moreover, some of these maps display a configuration of the compound that is inconsistent with the actual shape and number of its buildings (Figure 1).

Gelis, C. (1866) Plan Topographique de Jérusalem. Paris: Imp. Lemercier.

Pierotti, E. (1860) Plan de Jerusalem ancienne et moderne. Paris: E. F. Kaeppelin. 55 x 91cm

Wilson, C et al. (1865) Ordnance Survey of Jerusalem. Southampton: Ordnance Survey Office. 1:10'000

German Army, Aerial Photo (1917-1918) In Dalman G. Hundert deutsche Fliegerbilder aus Palästina (1925)

Figure 1: The first two maps by Pierotti and Gelis depict a configuration of the Russian Compound that is incompatible with the real one in terms of shape and number of buildings. The third representation, from the Ordnance Survey, is the most accurate, as evidenced by the 1917 aerial photo on the right (Dalman 1925: 19).

For this reason, the case of the Russian Compound is particularly interesting to study in the context of the cartographic representation of Jerusalem. These discrepancies raise questions about the dating of the sources and whether the incorrect representations might depict project plans for the area before construction. Using the proposed methodology to identify families and lineages in the reproduction of this area opens up the possibility of analyzing the reasons behind these inaccuracies and the involvement of mapmakers in the construction of the area.

3. Background: cartometric analysis of deformation and comparison

In the present contribution, the terms "deformation" and "accuracy" refer to *planimetric* deformation and accuracy, following Laxton's (1976: 40) definition as the assessment of *"the extent to which distances and bearings between identifiable objects coincide with their true values"*. In the

1970s, numerous studies on planimetric accuracy were conducted, mainly to determine a score of accuracy for old maps (Laxton 1976; Ravenhill and Gilg 1974; Stone and Gemmell 1977). The seminal works by Murphy (1978) and Blakemore and Harley (1980) then analyzed these efforts, laying the theoretical foundations of the discipline and highlighting the potential of accuracy analysis in understanding the origins of maps and the lineages of copies. Murphy (1978) conducted a comparative analysis of existing methodologies and highlighted the problem of methodological heterogeneity, a persisting issue even with the advancement of GIS (Livieratos 2006). Blakemore and Harley (1980) developed the theoretical discourse around the concept of accuracy and error, subdividing the types of accuracy in cartography into chronometric, geodetic and planimetric, and topographical.

Facilitated by the introduction of MapAnalyst, a dedicated tool to analyze and visualize distortions in old maps (Jenny et al. 2007; Jenny and Hurni 2011), various studies emerged, particularly focusing on accuracy assessment for land use and geomorphological changes detection (Tucci and Giordano 2011; Loran et al. 2018). Perthus and Faehndrich (2013) used MapAnalyst to assess the accuracy of two maps of the Holy Land, aiming to understand the production process.

The intuition that accuracy analysis could be more than a metric of reliability and the key to assessing stemmatic relationships (Murphy 1978; Blakemore and Harley 1980; Jenny and Hurni 2011; Jongepier et al. 2016) was explored in several studies that used accuracy analysis to understand the origins of maps (Bower 2011; Solchenbach 2021; Bartos-Elekes 2023). However, existing approaches either compared the positional accuracy of corresponding homologous points or, when employing the displacement vectors or the deformed lattice grids produced with MapAnalyst, relied on visual comparison. Both elements constitute a limit for scaling up the comparative analysis. On one hand, when working with many maps, visual comparison becomes impractical. On the other hand, finding the same homologous points across all maps may be impossible due to differences in scale and topographical content, making it challenging to identify a consistent set of homologous points.

4. Methodology: Encoding and Comparing Local Deformations Across Maps

The main objective of this section is to present a methodology for (1) processing local planimetric distortions of maps, represented by their deformed lattice grids, to encode them in a comparable format, and (2) using this transformed data for the local comparison of the maps.

This novel methodology is designed to detect local deformative similarities between two or more georeferenced maps, represented through their deformed lattice grids, even when limited to a small area. This is achieved by avoiding the step of global alignment between the deformed lattice grids and instead relying solely on the local geometric configuration of the nodes. While global alignment is effective for detecting complete copies (i.e., maps whose planimetric information completely matches that of another), it can fail when the information of a derivative map comes from multiple sources or when the copied portion is limited. In these cases, global alignment, such as with an affine transformation, would find the best overall match across the grids, potentially placing small corresponding areas at a high distance from each other.

To rely uniquely on the local geometric distortion of the lattice grid, the geometric configuration of each squared portion of 9 points of the grid, referred to as a *token*, is encoded. In particular, a *token* represents the geometric disposition of a point in relation to its eight surrounding points, encoded as the mutual distances between these points. This method leverages the property of geometrically

similar rigid shapes, which maintain a constant ratio between corresponding edges. By also incorporating the lengths of the diagonals, the rigidity of the shape is ensured, guaranteeing that only actually geometrically similar shapes will be classified as such (Figure 2).

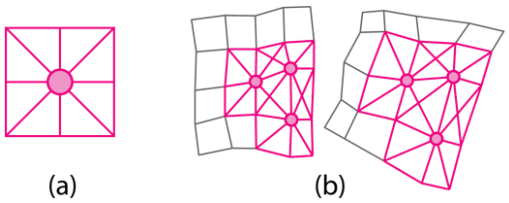

Figure 2: (a) The 9-point portion of the grid (called "*token"*) with the considered edges, and (b) how similar tokens look when comparing two deformed grids.

Section 4.1 shows how to encode the deformative *tokens* from the deformed lattice grid, creating the basis for the comparison, Section 4.2 presents how to exclude from further analysis the parts of the grid that are almost non-deformed and the ones that are outside of the alpha shape defined by the Ground Control Points (GCPs), and Section 4.3 illustrates the use of the prepared and filtered *tokens* to perform a comparison and clustering of the local deformations.

4.1 Encoding local deformation

This section outlines the process of using Ground Control Points (GCPs), i.e. the pairs of corresponding points between each map and a reference map[1] used for georeferencing, to create a deformed lattice grid for each map. It then explains how to use this grid to create a *token* representing the geometric configuration of each node. The result of this step is a matrix with the same shape as the initial grid, where each node contains a vector with the relevant edge lengths (i.e. the *token*). The steps needed to perform this part of the methodology are the following:

1. **Creation of the regular grid:** the process begins by identifying Ground Control Points (GCPs) on the map through georeferencing. These points are essential as they establish a correlation between specific locations on the map's pixel space and their real-world coordinates in the local reference system (CRS). Next, these correlated points are used to construct a rectangular bounding box in the local CRS of the area depicted in the maps. This bounding box serves as a reference framework for creating a regular grid within the local CRS. To enable future comparisons, the grid is created with a constant distance between neighboring nodes and from a shared starting point, to ensure alignment between the grids (Figure 3).

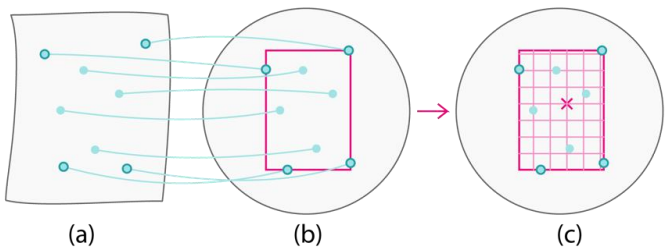

Figure 3: The Ground Control Points and the creation of the aligned, regular grid inside the bounding box. (a) Represents the map's image and the points that were found to perform the georeferencing. (b) Shows the location of these points in real-world coordinates. (c) Displays the creation of a regular grid in the local CRS starting with a common fixed point (indicated with an x) and using the extent of the GCPs as a bounding box for the grid.

[1] In our case we used OpenStreetMap as a reference, in the local EPSG (28193).

2. **Deformation of the grid with the Ground Control Points (GCPs):** the second step involves applying a thin plate spline transformation to the regular grid using the GCPs. These points provide vectors that map real-world coordinates to pixel points. By using these vectors to define a thin plate spline transformation, the regular grid in the local CRS can be transformed into a deformed lattice grid in the pixel space of the map. The choice of Thin Plate Spline interpolation is deliberate, as it allows the grid to elastically adapt to local deformations (Figure 4).

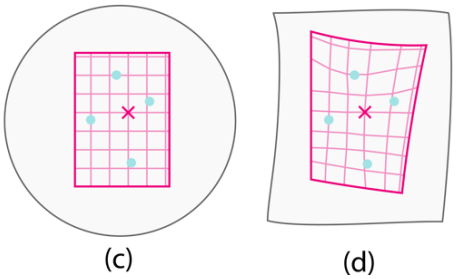

Figure 4: By using the vectors defined by the transformation from the real-world coordinates to the pixel space the regular grid can be deformed using a Thin Plate Spline interpolation.

3. **Encoding the *tokens*:** to detect areas of nearly perfect geometric similarity, the property of similar geometries to maintain a constant ratio between corresponding sides is used. To effectively employ this method, it is essential to first define the specific geometric shapes to be compared and to ensure their rigidity. In quadrilateral geometries (the edges of the grid), a constant ratio does not necessarily imply geometric similarity, as a shape with varying internal angles (e.g., a rhomboid) can lead to the same ratios between corresponding edges. To address this, we collect the distances between each node and its surrounding eight points (including diagonals) as well as the mutual distances between these eight points and their direct neighbors among the eight points (Figure 5). The collection of these edges, thanks to the diagonals, forms a rigid shape, as needed. These distances are gathered in a vector for each node, always following the same consistent order of 16 elements. This vector represents our *token*, which can be used in subsequent analyses and comparisons.

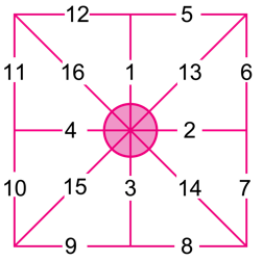

Figure 5: The order of the sides of the token used to collect the lengths of its sides in a vector.

The matrix containing the *tokens* is now suitable for assessing similarity directly through local geometric properties, without requiring any alignment. The next two parts of the methodology will detail how to filter out portions of the grid that should be excluded from the comparison and how to identify clusters of similar deformation across the filtered grids.

4.2 Filtering out the areas with low deformation

For the comparison of the deformed grids, we are interested in using only points of the grids that are relevant for studying the local distortions. This involves excluding points that are part of the grid uniquely due to its rectangular shape, specifically those at the corners defined by the bounding box of the Ground Control Points (GCPs). This preliminary filtering can be achieved by drawing the alpha shape (a concave hull with a desired level of complexity) around the GCPs and using it to filter out these points.

The second filtering step addresses the hypothesis that **a common lack of deformation does not necessarily indicate correlation**. Therefore, the analysis should include only tokens with a sufficient level of deformation. This filtering is achieved by comparing the tokens to a reference undeformed token, which is the edge lengths vector obtained from an undeformed square of 9 points. Tokens that closely resemble the undeformed configuration in terms of geometric similarity can be filtered out as non-deformed[2]. The combination of these two filters results in the set of relevant tokens that will be used for further analysis (Figure 6).

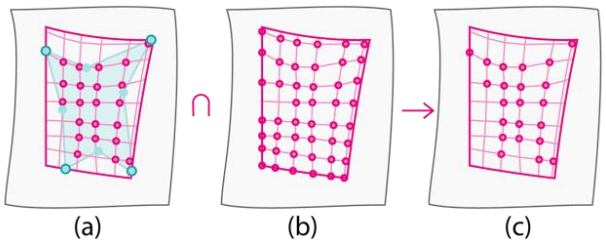

Figure 6: Using an alpha shape (a) it is possible to exclude the points of the grid that were outside the actual bounds of the Ground Control Points, and that were just included because of the rectangular shape of the bounding box. By intersecting this filtering with the one on the intensity of deformation (b) it is possible to select only the relevant points (c).

4.3 Clustering similar deformation

The third and crucial step consists of comparing the filtered *tokens* from the deformed grids of the maps under analysis. The goal in this step is to cluster the 9-point square subgrids that exhibit almost perfect similarity across maps. In particular, this analysis is carried out across the entire dataset of fifty maps. Therefore, a preliminary step of padding is necessary: while the grids are aligned with one another, they still have different extents. To address this, the maximum extent across all fifty maps is used to pad them to the same size with NaN values.

Now that all matrices of tokens have the same extent and are perfectly aligned, the value of non-NaN tokens across all maps can be analyzed. To do so, each point of the padded matrix with at least two comparable *tokens* is considered, and the pairwise ratio of the vectors is computed[3]. For each pair of compared *tokens*, the ratio of corresponding elements is calculated, resulting in a vector of the same length as the *tokens*, containing the ratios. Since a constant ratio corresponds to perfect geometric similarity, the difference between the maximum and minimum values in this ratio vector

[2] One limitation of this approach is that it relies only on geometric similarity as a constraint for filtering. Consequently, a token that is scaled up or down but remains geometrically undeformed will still be classified as undeformed.

[3] The ratio is computed after normalizing the token vectors, using the vector with the highest maximum as the denominator.

is used as a distance measure between the tokens. The closer this value is to 0, the closer the tokens are to perfect geometric similarity (Figure 7).

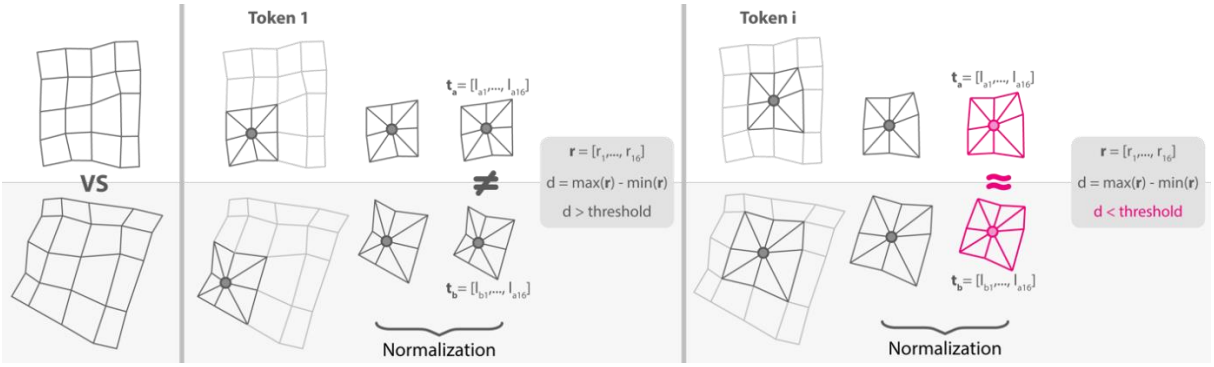

Figure 7: The figure shows the comparison of two grids at two different *tokens.* The first *token* is not detected as similar, as the distance is higher than the threshold, while the second (*token* i) shows a case of similarity.

Pairwise distances are then fed to a complete-linkage hierarchical clustering, where a threshold distance of 0.2 is used to determine the maximum distance between any two points in different clusters for merging them into the same cluster. By collecting the results at each point we can create a new matrix that contains the clustering of maps at the token level (Figure 8).

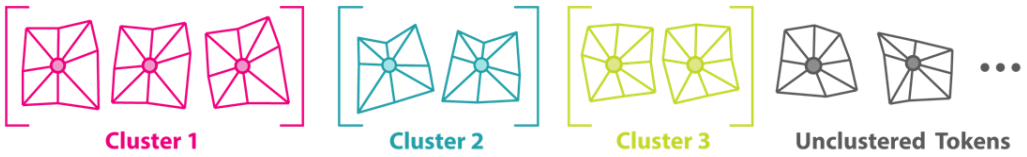

Figure 8: For every node of the deformed grid we cluster together the tokens that are similar to one another by using complete-linkage. In the figure, a possible situation is shown, where among all compared tokens three clusters were found. The tokens that are too distant to be clustered with another one are recorded as outliers, meaning that their information is not repeated in any other map.

This matrix can now be employed to study local similarities between maps. One useful derivative is represented by the regions of adjacent nodes shared by the same common maps, which can be derived using the *tokens'* adjacency. In the next section, the clustering results will be used to gain insight into a specific case study.

5. Results: stemmata of the representation of the Russian Compound

The regions of local similarity obtained with the proposed methodology can be used to visualize the clusters of maps sharing the deformation at a specific point or the regions of similarity shared by a selected subset of maps. These two filtered visualizations are useful for understanding the results in cases where a relatively high number of maps leads to a complex overall interpretation, such as in the present study.

This section presents the application of the methodology to a specific area of interest: the Russian Compound. By focusing on the maps representing this site and selecting a point within it, historical and stemmatic insights into the cartographic representation of the area can be deduced.

As anticipated in Section 2.1, this area represents a particularly interesting case study. On one hand, some of the maps depict it in a way that is not compatible with the actual shape and number of

buildings. On the other hand, some of these maps represent the area before the date reported in historical sources for the completion of its construction.

By setting the query point within the Russian Compound and selecting the maps depicting it from the dataset, our methodology identified and clustered together maps with similar local distortions. This process revealed two distinct clusters at the queried point (Figure 9), which can now be used to identify the predecessor within each cluster, i.e., the earliest map from which that particular representation was copied. The clustering results will be analyzed to identify the predecessor, also considering the maps that were clustered with no other map.

Figure 9: The two regions of similar maps obtained for the queried point (represented as a black dot) and the corresponding maps on the right.

The cluster on the bottom, composed of Wilson's Ordnance Survey maps (in both scales) and Whitney's map, consists of maps in which the area is represented correctly in terms of building number and shape. The predecessor in this case is represented by the Ordnance Survey map, as also declared in the title of Whitney's map, which represents its reduction.

The cluster on top is composed of earlier maps, some predating the area's completion, and needs a more thorough analysis to understand the origins of its content. This cluster differs from the second in the outer shape of the compound, which is slightly deformed, and more noticeably in the number and shape of the buildings, which are not compatible with the actual configuration of the Compound. A possible explanation is that the initial author depicted either a project or the ongoing construction, likely having direct access or connection to the site.

Given this information, the most likely predecessor within the cluster is Pierotti's 1860 map, as he was directly involved in the creation of the Russian Compound[4], unlike John Bartholomew. Therefore, Bartholomew's map must be a copy, indicating a potential dating error in one or both maps, as neither provides a date. The existence of a manuscript map from Pierotti dated June 1, 1860, held

[4] Recent research indicates that Ermete Pierotti, an Italian architect and engineer, acquired several key lots in Jerusalem between 1857 and 1859 for the Russian Consulate (Vakh 2014). This is supported by his mention in Boris Mansurov's letters; Mansurov was appointed by Grand Duke Konstantin Nikolaevich, with the support of the Tsar, to look for building sites for the Russian Empire in the fall of 1858 (Winning 2022). Moreover, Pierotti referred to himself as "Ispettore delle fabbriche russe" in an 1864 letter reported by Legouas (2013), implying his official role. However, his involvement in the project is not explicitly mentioned in official documents but can be inferred from "mere scraps of information" (Vakh 2014: 201), such as the ones mentioned or his appearance in a photograph labeled "the builders of the Russian Compound in Jerusalem," dated 1858 (Vakh 2011, Vakh 2014: 195).

in the National Library of Spain (Pierotti 1860) closely resembling this map suggests that Pierotti's map could be indeed from around 1860, while Bartholomew's map should be a later copy. Notably, Pierotti's misrepresentation continued to be reproduced in maps for several years after the site's completion. This indicates that later authors copied from Pierotti or a derivative without taking ground measurements or even accessing the area, as the error is particularly evident. The origins of the content of John Bartholomew's map can be further investigated by plotting all its detected predecessors among the fifty maps employed as a dataset (Figure 10).

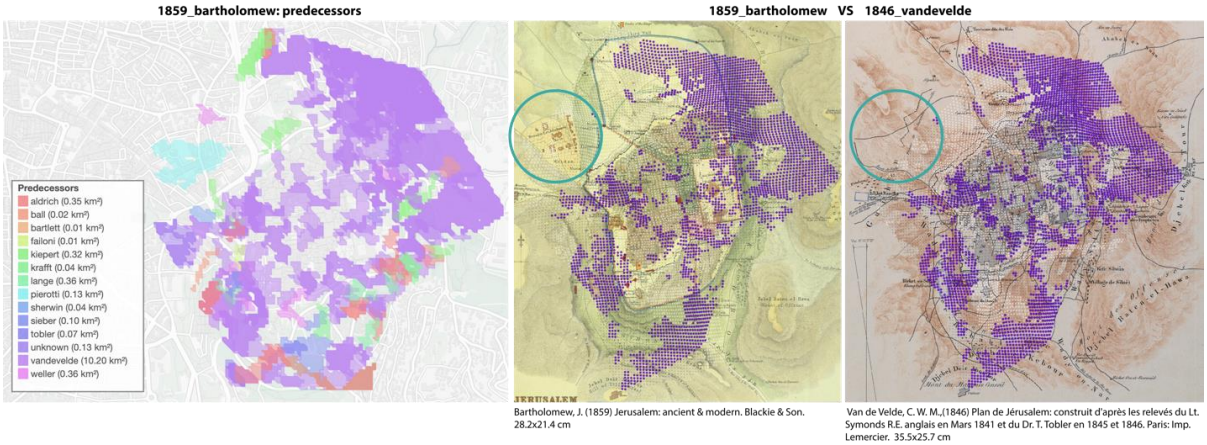

Figure 10: On the left, the authors of the possible predecessors of Bartholomew's map are listed, along with their extent in square kilometers. On the right, the tokens detected as highly similar are plotted in purple, while a blue circle highlights the area of the Russian Compound that was added to Bartholomew's map.

The results suggest that Bartholomew combined two main sources of information in its map: an 1846 map from C.W.M. Van de Velde, which didn't show the Russian Compound, for the main part, with the wrong representation of the site done by Pierotti around 1860.

Now that Pierotti's map has been identified as the predecessor, how can the difference in the number and shape of the buildings in the area be explained? Vakh (2014) suggested that Pierotti may have played a role in convincing Surayya-Pasha, his patron and friend (and consequently Sultan Abdulmecid) to gift the neighboring lot for the Russian Compound with the hope of participating in its construction as an architect. However, Vakh concluded that Pierotti no longer received a salary from the Russian Consulate by the end of 1859, which indicates that he was likely not successful in his endeavor. In another article (Vakh 2011) three projects for the area were reported. Despite an initial overall rotation of the area, the buildings in these projects resemble those that were eventually constructed and not the ones drawn by Pierotti. It is therefore possible that in his map Pierotti tried to propose a variation to the initial plans by slightly modifying the number of buildings, perhaps hoping to secure a role as an architect in the project.

Four maps among the ones representing the area were not clustered with any other at the queried point: Weller (1860), Pierotti (1864), Gelis (1866), and Pierotti (1869). It is possible to explore their connections to other maps by plotting all their predecessors (Figure 11).

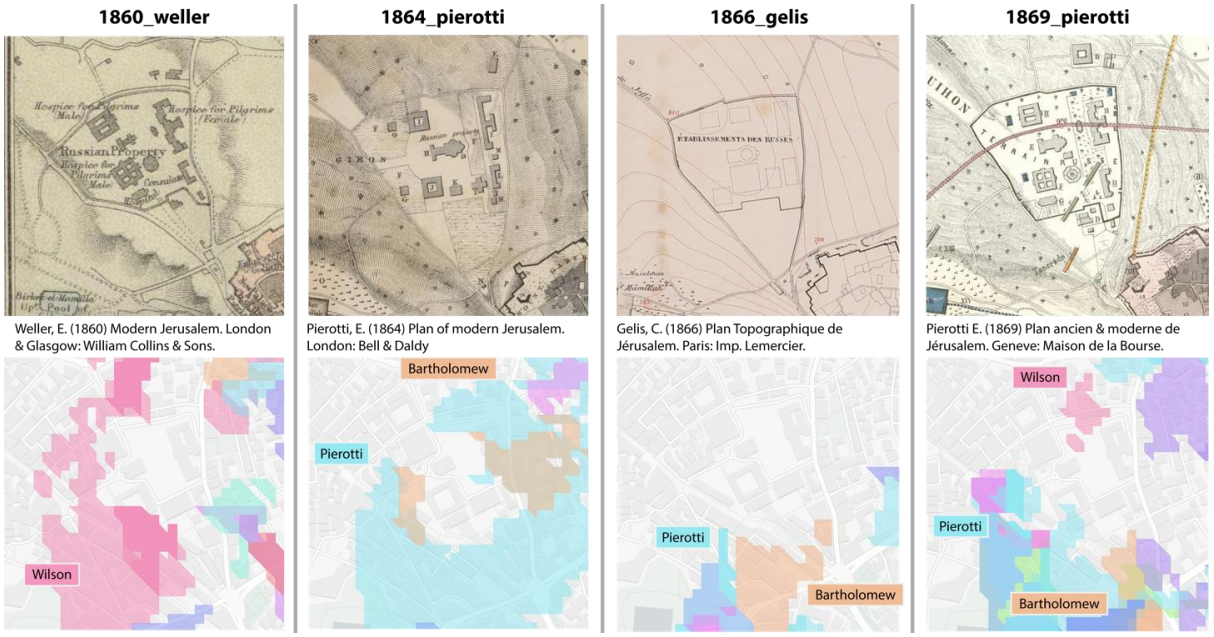

Figure 11: The regions of similarity shared with any previous map for the four unclustered maps.

Weller's map, dated 1860 according to the online archive of the National Library of Israel, shows a region of similarity with Wilson's map. However, its reduced size, simplified information, and intended publication (Collins' Pocket Atlas of Scripture Geography) make it unlikely to have been the predecessor of a highly influential map like the Ordnance Survey maps by Wilson. The latter are most likely the predecessors of Weller's map, which must necessarily be dated after them. The reason why the region of similarity does not cover the entire area can be explained by deformations resulting from the simplification of the map.

In his 1864 map, Pierotti seems to have corrected the outer shape of two buildings, which is now correctly squared and not rectangular. However, the number of buildings and the overall shape of the lot remains the same and when looking for its predecessors, his previous maps emerge as sources.

In the case of Gelis' 1866 map, there is no detected similarity even in the outer shape of the area. Visually, it is noticeable that the representation and disposition of the buildings differ from both previous clusters and are not compatible with the actual configuration, raising questions about the reasons behind this misrepresentation, especially considering that the area was completed by this point.

Pierotti's 1869 map, on the other hand, maintains similarity in the outer shape with his earlier maps (as indicated by the colored regions on the sides) but depicts the shape and number of the buildings correctly. The fact that the similarities with Wilson's map are extremely limited in extent could indicate either a new independent ground measurement by the author or a deliberate slight modification of the building proportions.

From the results of this analysis, we can derive a stemmatic schema illustrating how this area's representation was likely copied and modified in time (Figure 12).

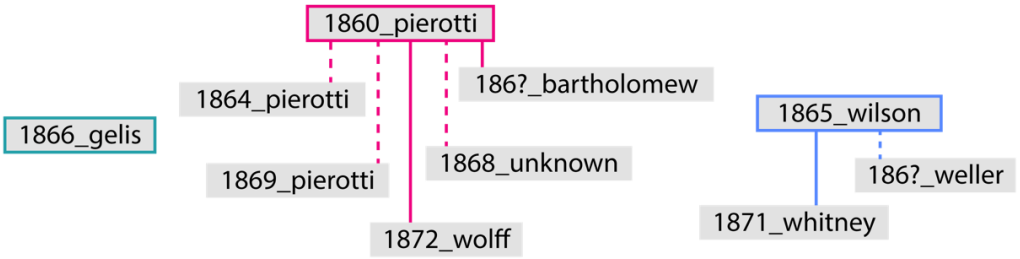

Figure 12: Stemmatic schema representing how the area's representation was introduced, copied, or modified based on the deformative analysis paired with information from secondary sources. The dashed line indicates a map that shares similarities with the original map but presents some modifications from it.

Gelis' incorrect representation of the area could not be related to any other source among the considered ones and therefore stands isolated. The first family is represented by Pierotti's 1860 map, from which several other maps stem. His map from 1864 has a first partial correction of two buildings' shape, and the one from 1869 finally corrects their number, while maintaining the old shape of the area. The 1868 map appears as a reduction of Pierotti's 1860 map, but with a removal of some of the non-realized buildings (while keeping the wrong shape of the others). The second family stems instead from Wilson's Ordnance Survey maps.

6. Discussion and limitations

This analysis helped shed light on dating errors in both clusters, provided clues for assumptions about Pierotti's representation, and offered further proof of his involvement in the project, consistent with previous studies based on literary sources. Ultimately, the methodology's results enabled the formulation of a hypothesis for the stemmata of the representation.

It is worth noting that the insights provided by the presented methodology act as a magnifying glass, offering valuable support for forming and motivating hypotheses in combination with traditional historiography tools.

However, in the application of our methodology for detecting and segmenting regions of local similarity across historical maps, several factors introduce inherent variability and potential limitations and should be considered when interpreting the results.

First, the choice of thresholds for filtering out non-deformed areas and for defining similarity in the clustering process determine the sensitivity of the methodology in detecting relevant distortions and similarities. Then, the scale of the maps and the level of detail in the depicted information also play a critical role. Maps at different scales might exhibit different levels of deformation, and simplifying information to fit a particular scale can alter the representation of local features. Such simplifications might smooth out deformations or exaggerate them, affecting the detectability of local similarities. In maps where detailed features are generalized, small but significant distortions might be lost, impacting the methodology's ability to detect copying relationships accurately. Moreover, the distortions introduced by the map support or by the digitization process can introduce another layer of variability where some similarities may go undetected.

7. Conclusion

The digital analysis of planimetric deformations in maps constitutes a promising methodology for identifying copying relationships, especially when used across many maps. The novel methodology introduced in this study has proven effective in identifying regions of local deformation that are

similar across maps, clustering together the sources that represent an area with similar deformations. The deformed lattice grids obtained from georeferencing, padded and aligned across maps, constitute the object used for the comparison. By leveraging the properties of geometric similarity of 9-point squares of the grid, the methodology successfully worked also in cases of limited, partial similarity.

A selection of fifty 19th-century maps of Jerusalem between 1810 and 1872 was used to analyze the history of the cartographic representation of the Russian Compound, one of the first quarters built outside the walls of the Old City. The clusters of similarity obtained from their comparison revealed the presence of two clusters of similarity among the maps representing the area inside the dataset. The clustering results, combined with historical insights, made it possible to advance new hypotheses for the dating correction of some of the maps and the lineages of copies between them. Moreover, the propagation of an evident error in the representation, stemming from Pierotti's 1860 map, can also provide information on the authors who most likely did not take ground measurements or access the site.

By providing clusters of local similarities, this methodology supports the study of map distortion propagation, especially in cases where cross-influences are largely present. Moreover, the possibility to query a specific point or filter the regions based on a selection of maps allows for an easier interpretation of the results in large collections.

In conclusion, integrating computational cartometry with traditional historiography tools can uncover hidden historical insights and support the creation of genealogies of map collections even when the collections are large and when differences may be slightly visible.

References

Bartos-Elekes, Z. (2023). Behind the first Habsburg map of Transylvania – comparative analysis of contemporary manuscript maps. *International Journal of Cartography* 9 (3): 507-524. doi:10.1080/23729333.2023.2240861.

Ben-Arieh, Y. (1975). The growth of Jerusalem in the nineteenth century. *Annals of the Association of American Geographers* 65 (2): 252-269.

Blakemore, M. J. and Harley, J. B. (1980). The search for accuracy. *Cartographica* 17 (4).

Bower, D. I. (2011). Saxton's maps of England and Wales: The accuracy of Anglia and Britannia and their relationship to each other and to the county maps. *Imago Mundi* 63 (2): 180-200. doi:10.1080/03085694.2011.568704.d

Dalman, G. (1925). *Hundert deutsche Fliegerbilder aus Palästina*. Gütersloh: Bertelsmann. Universitätsbibliothek Tübingen.

Gerd, L. and Potin, Y. (2018). Foreign affairs through private papers: Bishop Porfirii Uspenskii and his Jerusalem archives, 1842–1860. In A. Dalachanis and V. Lemire (Eds.), *Ordinary Jerusalem, 1840-1940: Opening New Archives, Revisiting a Global City.* Leiden: Brill, 100–117. doi:10.1163/9789004375741_008

Goren, H., Faehndrich, J., Schelhaas, B., Weigel, P. (2017). *Mapping the Holy Land: The foundation of a scientific cartography of Palestine.* Tauris Historical Geography Series, 11. London: I.B. Tauris.

Jenny, B., Weber, A., Hurni, L. (2007). Visualizing the planimetric accuracy of historical maps with MapAnalyst. *Cartographica* 42 (1): 89-94. doi:10.3138/carto-v42-1-089.

Jenny, B. and Hurni, L. (2011). Studying cartographic heritage: Analysis and visualization of geometric distortions. *Computers and Graphics* 35 (2): 402-411.

Jongepier, I., Soens, T., Temmerman, S., Missiaen, T. (2016). Assessing the planimetric accuracy of historical maps (sixteenth to nineteenth centuries): New methods and potential for coastal landscape reconstruction. *The Cartographic Journal* 53 (2): 114-132. doi:10.1179/1743277414Y.0000000095.

Laxton, P. (1976). The geodetic and topographical evaluation of English county maps, 1740-1840. *The Cartographic Journal* 13 (1): 37-54.

Legouas, J.-Y. (2013). Saving Captain Pierotti? Palestine Exploration Quarterly 145 (3): 231-250. doi:10.1179/0031032813Z.00000000047.

Levy-Rubin, M. and Rubin, R. (1996). The image of the Holy City in maps and mapping. In N. Rosovsky (Ed.), *City of the Great King: Jerusalem from David to the Present.* Cambridge, MA and London, England: Harvard University Press, 352-379. doi:10.4159/harvard.9780674367098.c20

Livieratos, E. (2006). On the study of the geometric properties of historical cartographic representations. *Cartographica* 41: 165-176. doi:10.3138/RM86-3872-8942-61P4.

Loran, C., Haegi, S., Ginzler, C. (2018). Comparing historical and contemporary maps - a methodological framework for a cartographic map comparison applied to Swiss maps. *International Journal of Geographical Information Science* 32 (11): 2123-2139. doi:10.1080/13658816.2018.1482553.

Murphy, J. (1978). Measures of map accuracy assessment and some early Ulster maps. *Irish Geography* 11 (1): 88-101. doi:10.1080/00750777809555719.

Perthus, S. and Faehndrich, J. (2013). Visualizing the map-making process: Studying 19th century Holy Land cartography with MapAnalyst. e-Perimetron 8: 60-84.

Pierotti, E. (1860). Plan de Jérusalem ancienne et moderne: Primo studio di Ermete Pierotti, fatto su la località [Manuscript map]. Scale [ca. 1:12,333]. 26.5 x 29.5 cm. Biblioteca Nacional de España, MR/42/494. In digital form, http://bdh-rd.bne.es/viewer.vm?id=0000069851&page=1.

Ravenhill, W. and Gilg, A. (1974). The accuracy of early maps? Towards a computer aided method. *Cartographic Journal* 11 (1): 48-52.

Rubin, R. (1990). The map of Jerusalem (1538) by Hermanus Borculus and its copies - a carto-genealogical study. *The Cartographic Journal* 27 (June): 31-39.

Rubin, R. (1991). Original maps and their copies: Carto-genealogy of the early printed maps of Jerusalem. *Eretz Israel* 22 (David Amiran Volume): 166-183. (In Hebrew).

Rubin, R. (1999). *Image and reality: Jerusalem in maps and views.* Jerusalem: Magnes Press.

Solchenbach, K. (2021). Comparing old maps with cartometric methods. *e-Perimetron* 16 (2): 55-62.

Stone, J. C. and Gemmell, A. M. D. (1977). An experiment in the comparative analysis of distortion on historical maps. The Cartographic Journal 14 (1): 7-11.

Tucci, M. and Giordano, A. (2011). Positional accuracy, positional uncertainty, and feature change detection in historical maps: Results of an experiment. *Computers, Environment and Urban Systems* 35 (6): 452-463. doi:10.1016/j.compenvurbsys.2011.05.004.

Vakh, K. A. (2011). Velikii kniaz' Konstantin Nikolaevich i russkoe palomnichestvo v Sviatuiu Zemliu. K 150-letiiu osnovaniia Russkoi Palestiny. 1860–1864. Moscow: Indrik. (In Russian).

Vakh, K. A. (2014). Ermete Pierotti in the Russian service: New biographical discoveries. *Zeitschrift Des Deutschen Palästina-Vereins* 130 (2): 194-204. In digital form, *http://www.jstor.org/stable/43664933.*

Von Winning, A. (2022). Jerusalem: The Russian Compound. In *Intimate Empire: The Mansurov Family in Russia and the Orthodox East, 1855–193*6. New York: Oxford University Press, 53-85.