Measuring the Accuracy of Historical Maps

Emmet Morrin

A Dissertation

Presented to the University of Dublin, Trinity College

in partial fulfilment of the requirements for the degree of

M.Sc. Computer Science (Year 5)

Supervisor: John Dingliana

April 2024

Measuring the Accuracy of Historical Maps

Emmet Morrin, M.Sc. Computer Science (Year 5)

University of Dublin, Trinity College, 2024

Supervisor: John Dingliana

The goal of this project is to develop an automated system to score the differences between historical maps from different sources or periods. This involves the vectorisation and annotation of historical maps in order to more effectively compare the differences in the content of two maps. This dissertation focuses mostly on the vectorisation and annotation of coastlines of land masses and the following comparison of those results. The exploration of four different differences metrics are discussed in their effectiveness at comparing the differences in two map contours; the sum of absolute differences, Hausdorff distance, cosine similarity, and template matching.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Problem Overview

This dissertation will explore the annotation or vectorisation of historical maps, including the methods used in the past and their use cases. It will discuss the development of a tool used to vectorise maps for the purpose of comparing the coastlines of two similar maps. The use of various difference metrics will also explored, which can be used to numerically measure the difference between two vectorised maps.

Vectorising an image refers to the process of converting a traditional raster image into a vector image. Raster images are the majority of images we see in everyday use. These images are composed of a discrete grid of pixels. Vector images are composed of mathematically defined points, lines, and shapes. The ability to convert raster historical maps into vector maps allows for the further analysis and deeper visualisation of the data included on these maps.

This project holds significance for several reasons:

Knowledge about land over time: This project helps to show how the knowledge about the land we live on increases over time. The gain in knowledge can be caused by a number of contributing factors, such as advancements in mapping technology, the change in culture requiring more accurate maps, and the ability for mapmakers to work collaboratively. This can help us to show how mapping technology has increased over time, and influence the development of future technologies.

Error on Purpose: Error could have introduced to historical maps purposefully. This purposeful error could have been introduced to deceive others in history. This project stipulates that some powers may have wished to exaggerate land sizes for their own gain. A record of where this may have happened is likely lost to time, but using a tool such as this could help to highlight this tampering.

Data available only in Historical Documents: Because the land we use changes over time by our own influences such as urban sprawl or agriculture, data about the underlying land itself can be lost. This could be available in historical maps, but could have error which influence their usefulness. Being able to review how much error and account for that could provide a solution to this issue. Seeing the change in land use could also highlight the human impact on the planet.

Coastline Change: These maps act as a record of a coastline's position at a specific point in time. However, this project assumes that this change is minimal in comparison to the error in the maps themselves. This project could be extended in the future to consider the change in coastlines over time. Having a longer record of coastline change would help us to measure coastal erosion or sea levels and their acceleration. This could help us to predict future coastal changes and manage vulnerable areas.

1.2 Research Objectives and Contributions

The objective of this project is to develop an automated system to score the differences between historical maps from different sources or periods. This analysis is based on specific landmarks or contours that the user can select, or can be automatically chosen, on an image of a map. After the image map has been sufficiently annotated it can be used as a comparison against similarly annotated maps.

This project is split into two distinct pieces of work: vectorising map data, and comparing vectorised maps.

1.2.1 Vectorising Map Data

This project will use a number of computer vision techniques to transform regular raster images of maps into accurately annotated images, or vectorised images. This involves extracting key information from the map image, such as the coastline, borders, or points of interest, and saving them in a way in which can be used to compare against other similarly vectorised images.

1.2.2 Comparing Vectorised Maps

After at least two images have been vectorised as described in 1.2.1, they can be compared against each other. This project explores a number of different metrics for comparison and evaluates their effectiveness. The metrics evaluated are: Sum of Absolute Differences, Hausdorff Distance, Cosine Similarity, and Template Matching. Using a weighted sum of the appropriate metrics can provide an accurate measure of the similarity of two map images. Using the similarity of historical maps with modern-day maps provides their overall accuracy.

1.3 Document Structure and Contents

This dissertation will first explore the background of work in this area in Chapter 2, show the limitations present in current research and where this project aims to fit into those gaps in research. Then it will continue to describe the overview of the design of the tool presented in this project in Chapter 3. After the design, this dissertation will describe the actual implementation of the design in Chapter 4. It will then show the results and evaluation of the implementation of this project in Chapter 5. Lastly, the conclusions and future work will be discussed in Chapter 6.

Chapter 2

Background

This chapter discusses the existing tools and research in the area of this disseration. It will start with surveying the existing tools and their features when handling historical maps. Then, a review of the existing research in vectorising raster images, as well as vectorisation of images specific to historical maps. In the next section, research in measuring the differences between maps and their contours will be reviewed. Lastly, this chapter will end with a summary of the work reviewed and its applications to this project.

2.1 Existing Tools

This section aims to acknowledge and review a set of the existing tools which are applicable to the vectorisation of historical maps. In this section the following tools will be discussed:
ArcGIS (2.1.1), geographic information system (GIS) software.
GRASS (2.1.2), also GIS software.
GDAL (2.1.3), Geospacial Data Abstraction Library.
OpenCV (2.1.4), Open Computer Vision Library.
Potrace (2.1.5), software for converting raster images to vector images.
While this is not a comprehensive list of all of the tools available to anyone working with historical maps, the review of this list of tools will help to inform what transformations are possible and any possible limitations when attempting to vectorise historical maps.

2.1.1 ArcGIS

ArcGIS is a collection of geographic information system (GIS) software developed by ESRI. While originally released as a desktop application, ESRI has migrated ArcGIS to an online service model. They focus mostly on providing business-facing tools to analyse geospatial data. ArcGIS Online (the online service for ArcGIS) also supports a community collaboration model, which aids in the ability to share geographic data between different users.

The maps hosted on ArcGIS are vector maps first and foremost, but they mostly deal with current maps, not historical. While tools to analyse data are available through ArcGIS, these tools for analysis are concerned with current day information and their implications. Additionally, ArcGIS does not provide the necessary tools to either vectorise historical raster maps, or provide a means to measure differences between between two maps, visually or numerically. Tools do exist in ArcGIS to convert raster images to vector images, but these are not suitable for historical maps and require manual work to achieve this. Documentation of the tools in ArcGIS which can be used to vectorise maps is available at ESRI (2024).

2.1.2 GRASS GIS

In comparison to ArcGIS (2.1.1), GRASS GIS is an open source geographic information system software. While GRASS GIS is primarily for vector data, it also offers raster analysis tools. When using historical maps as scanned raster images, GRASS can convert them to vector format and use tools to compare those features. This conversion however does not intelligently convert the raster image to a vectorised version. It will convert any line and polygon information found within the raster image to a vector image without labelling any of the data. This means that the resulting vector image will require more processing in order to be able to extract specific data, such as coastlines, roads, borders, or rivers. Despite this however, some comparison of maps is possible without labeling the vectorised data. Tools provided such as the ability to overlay two maps on top of each other can be used to view the differences between two maps and highlight any changes present. This is limited however to maps which have a similar boundary, and the comparison is limited to internal structures within that boundary. This tool can also be used to georeference historical maps to real-world coordinates and view that historical map overlayed on that georeferenced point. These tools available however are entirely visual, and do not provide any functionality to numerically measure the differences between these maps.

2.1.3 GDAL

According to GDAL's web documentation (OSGeo, 2024), GDAL is a translator library for raster and vector geospatial data formats. As a library, it presents a single raster abstract data model and single vector abstract data model to the calling application for all supported formats.

It provides functionality to work with raster and vector images in the context of maps. For example, some of its uses include the ability to translate geographic data across different file formats and across different coordinate systems, georeferencing raster images like satellite imagery, warping geospatial data to adjust for any distortions present, and providing functionality to integrate multiple different programming languages to ensure interoperability. In the context of vectorising historical maps, which are likely have clarity errors, the use of GDAL on its own is often not accurate enough or sufficiently account for inaccuracies in the maps themselves to extract important information (Giraldo Arteaga, 2013).

2.1.4 OpenCV

OpenCV is an open source computer vision library specialising in real-time computer vision. It provides the necessary functionality to its users to conduct any image and video analysis (OpenCV, 2024a). Some of the functionality the algorithms OpenCV provides include the ability for object detection and recognition, feature extraction, image filtering and noise reduction, motion analysis, and a platform for machine learning capabilities.

In the context of vectorising historical maps, OpenCV can be used to interact with the base raster images and perform in-depth processing and analysis on the image. Using a suitable set of operations, a wealth of information can be extracted from raster images using OpenCV. It can also be used to perform image processing more effectively than some of the other tools discussed in this chapter.

2.1.5 Potrace

Potrace is a software tool designed specifically to convert bitmap images to vector graphics. This tool focuses primarily on black and white images, where it excels at converting shapes and outlines, while it is not ideal for complex colour images (Selinger, 2003). The process of vectorisation is slightly different than what has been discussed in this dissertation. Vectorisation using Potrace results in a vector graphics image, such as those used in image formats like SVG, PDF, and DXF. See Figure 2.1.5 for an example of this process. While the use of vector graphics can be used to analyse map images, this dissertation concentrates more on labeling specific polygons present within the image, rather than having only a vector representation of the raster map image. This is effectively a different vectorisation output than what Potrace can offer.

Figure 2.1: This Figure shows the result of using Potrace on an example raster image on the left, and producing a vector graphics image, shown on the right. Example images given by Potrace (2024).

2.1.6 Existing Tools Summary

While many powerful existing tools exist for any user wishing to analyse historical maps as discussed in this section, no adequate off-the-shelf tool exists to vectorise and compare the differences between historical maps. Integration between the tool presented in this dissertation and the tools discussed would aid in its usefulness and adoptability by professional cartographers already familiar with the tools discussed here, and in particular ArcGIS and GRASS GIS.

2.2 Vectorisation of Raster maps

The vectorisation of raster images and maps can be generally split into utilising two main techniques; the use of purely traditional image processing heuristics and techniques which mostly rely on the use of machine learning based approaches. Existing research in the vectorisation of maps also focused significantly more on details of smaller scale urban maps. This mostly included identifying zoned areas of the maps, such as for building information and land use.

As an overview of some of the traditional techniques used to convert raster images to vectorised ones, Wenyin and Dori (1999) surveys a series of techniques which can be used to extract information from line drawings. While most historical maps are not strictly line drawings, many of the techniques discussed in this article are useful when extended and applied to these maps. The use of Hough Transforms are discussed, which usually utilised in the detection of straight lines and ellipses in binary images. Different thinning algorithms are discussed, which are used to convert thick lines present in an image to a single pixel wide skeletal structure of the original lines. Particularly relevant to this dissertation, methods using contours to extract line data is also discussed as a less computationally expensive option to thinning algorithms. They find that contour-based methods can be used by approximating the middle of thicker lines by the in-between of two parallel lines, but that the junction of two or more lines presents significant problems for this alternative.

Chen et al. (2021) show the use of deep edge filtering and closed shape extraction. This research explored the use of multiple different learning technologies in order to extract the edge information from the raster map image, where it found using the Holistically edge detector (HED) the most appropriate learning technology for this application of edge filtering. Chen et al. state, "The second observation is that pre-trained HED is the best deep edge filter. The Holistically Edge Detector outperforms the other architectures both when used with a naive closed shape extraction or with a more advanced one such as our watershed module."

Figure 2.2 shows a set of this research's results from comparing the use of HED for edge filtering against another deep learning technology, U-Net. This figure shows on the left the intermediate stages of both of these two technologies, showing their produced Edge Probability Map (EPM). On the right it shows their final predicted shape, where green are correctly predicted shapes and red are missed.

After utilising HED for edge filtering, this research proposes the use of watershed segmentation to identify the connected components and finally using an off-the-shelf vectorisation tool, GDAL to obtain the final shapes. Watershed segmentation is a transformation applied to an image which mimics the process of rain pooling in valleys, as it flows down to lower points. Watershed lines would form at the ridges of these valleys. This process is similarly applied in image processing, creating watershed lines where different regions in the image meet. For a further in-depth explanation of watershed segmentation, see Serge and Meyer (1992).

Iosifescu et al. (2016) shows how without the use of deep learning, traditional computer vision techniques can also be used to vectorise urban maps. This research focused on extracting building and river information from urban maps, specifically that of Zurich. Iosifescu et al. propose one possible method to extract that building and river information. They extract each of the colour channels and perform separate image processing on each to target different information present in these maps. For example, in the maps explored in that project, water was usually indicated in a blue colour. One of the key steps they found that by subtracting the red channel from the blue channel, it was possible to extract only the river information. While this worked very well for the maps chosen in that project, it is likely that this is, along with other steps taken in that project, is specific to those

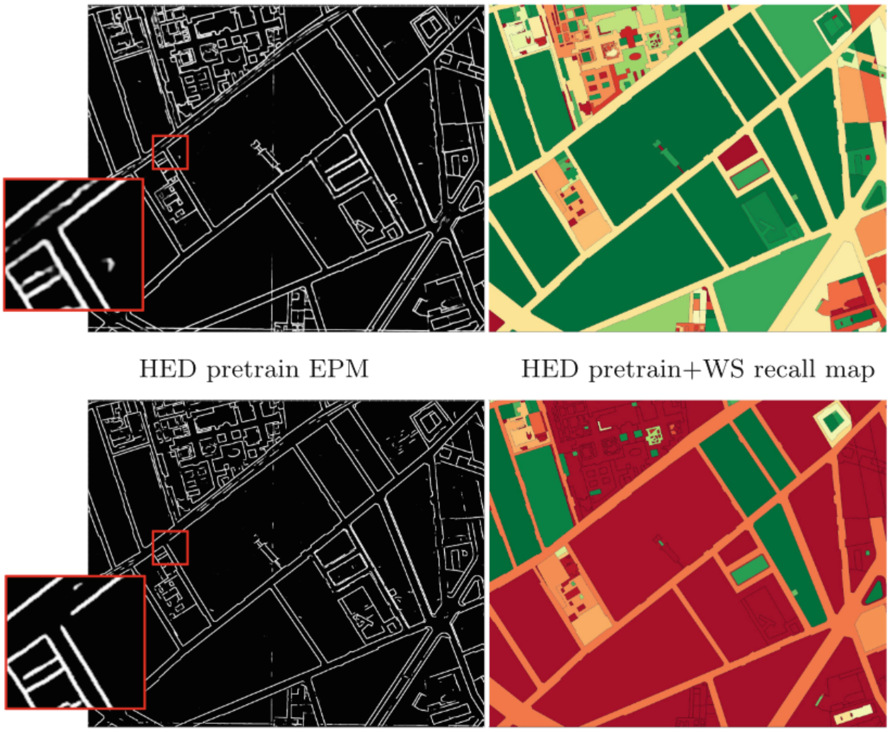

Figure 2.2: This figure shows on the left the intermediate stages of both of these two technologies, showing their produced Edge Probability Map (EPM). On the right it shows their final predicted shape, where green are correctly predicted shapes and red are missed.

maps chosen. This is acknowledged by Iosifescu et al., and despite these limitations, the methods presented in this project are useful in exploring different ways of vectorising urban maps.

Giraldo Arteaga (2013) explores the use of computer vision techniques to aid in the data extraction of historical maps from the database of maps from the New York Public Library. They aimed to help in the automatic data extraction of these maps which have, up until this project, been a manual task. In Giraldo Arteaga (2013), Giraldo Arteaga shows how useful the automatic vectorisation of maps can be when compared against that done manually. They state that, in the 3 years prior to this project, 170,000 polygons have been processed manually across New York City street atlases. In just under 24 hours, a whole atlas containing 55,000 polygons was processed using the methods outlined in this paper.

Giraldo Arteaga (2013) find that off-the-shelf tools such as GDAL "require either very clean and simple polygons or human intervention to assist the polygonization algorithm". As such they found that existing tools such as GDAL alone were not accurate enough for their use case. Instead, Giraldo Arteaga have proposed a set of three steps for the vectorisation process: Raster image thresholding, rough polygon extraction, and polygon analysis and simplification.

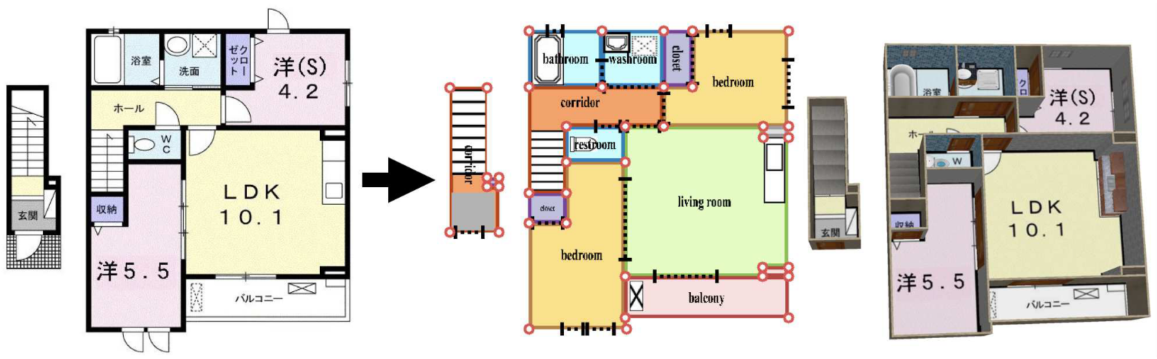

Figure 2.3

Stepping away from historical map specific research, Liu et al. (2017) presents a tool which can be used to convert a raster floor-plan image into a vector image which can be used to create 3d renders of that floor-plan (see Figure 2.3). Liu et al. outlines that their process is to be split into two intermediate layers before the final vector image is computed. The first of these layers includes utilising a convolutional neural network (CNN) to convert the raster input image into a set of "junctions". These junctions describe the essential information present in the raster image, such as wall corner data, door placements, furniture, etc. After this information has been extracted by the use of deep learning using a CNN, the resulting "junctions" are converted to a set of primitive shapes, i.e. points, lines, and simple polygons. After enforcing geometric and semantic restraints that helps to remove any inaccuracies from the previous step, the primitives found can be used to construct the desired vectorised image. The result of this process can be seen in Figure 2.3.

2.2.1 Vectorisation Summary

Overall, the existing research in vectorising historical maps has mostly focused on that of urban areas (Iosifescu et al. (2016), Giraldo Arteaga (2013), Chen et al. (2021)). The methods used to extract the data relevant to vectorise images will usually begin with a pre-processing step, converting the complex raster image of the map to a binary image where it is easier to perform image processing than on the original RGB image. The research explored in this review all chose different paths of vectorising this data, which was tailored to their use cases. There is a gap in the existing research for a solution which aims to be suitable for more general maps, rather than for a smaller subset of maps. This is likely due to the complexity and variance of different maps. These historical maps do not have a set of standards to which they are made, since they are produced across different cultures and time periods. The existing research explored in this review serve as a baseline for what different methods are useful when dealing with specific maps, and could be used for an extension to a larger set of maps.

2.3 Map Comparison

Research in the comparisons between historical maps in terms of their content has been mostly focused on land use over time and their implications. Pontius and Lippitt (2006) explores how land use has been changing alarmingly fast between 1971 and 1999 in certain areas of central Massachusetts. Pontius and Lippitt considers the case where errors present in the maps used for this detection in changes could explain some of those changes, rather than the land use itself being changed. This article explored how even how small percentage errors can reduce the significance of the findings from comparing two different maps. They found that in their example between these two time periods, "If the maps were 85 percent correct, then error could explain nearly all the observed differences between the maps", except for a few select cases where land use was converted to built land. This research shows the effects that inaccuracies present in maps have on the results of those analysing them. Being able to account for that error or even correct it can help us to make more accurate results from our analysis of historical maps.

Similarly, Loran et al. (2018) considers the land use and cover change of topographic maps. They argue that "there is a need for a standardized process to assess map comparability in a systematic way in order to avoid, or at least minimize, the detection of spurious landscape changes due to incompatible map series". In this article, they propose using a series of georeferenced points present on both maps, and use a root mean square error of those points to present an error metric for the maps chosen. Interestingly, using the difference between these two points, a visual distortion analysis of the maps was possible to be constructed against their modern day topographic data.

Smith and Cromley (2006) and the follow-on research article Smith and Cromley (2012) explore the coastline changes present in historical maps. In this case they are concerned with the coastline change of Connecticut, USA. Using the vectorisation process as outlined in Smith and Cromley (2006), two difference metrics were discussed in Smith and Cromley (2012). They compare the use of a transect-from-baseline technique, as was standard for the Federal Emergency Management Agency, and their proposed change polygon method.

A visualisation of the transect from baseline technique can be seen in Figure 2.4.

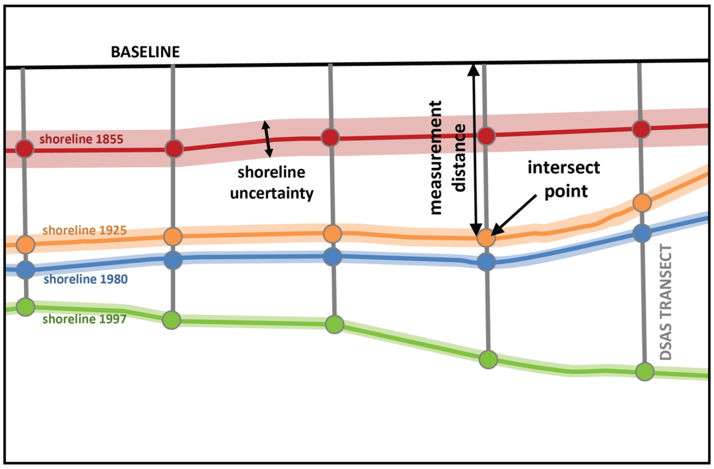

Figure 2.4: This figure visualises how the transect baseline technique is used. Using a reference baseline as labeled in this image, perpendicular lines from this baseline are used to calculate coastal change, from where they intersect with those perpendicular lines from the baseline. Baselines can be generated in multiple ways, and can be in front, between, or behind the coastlines being analysed. This figure was provided by Woods Hole Coastal and Marine Science Center (2018).

However, Smith and Cromley found two main problems with using this technique for measuring coastal change; the transect lines can miss the nearest coastline and transect instead with a coastal point further away, and portions of the coastal contour can be unintentionally excluded from the analysis. See Figure 2.5 for a visualisation of these two problems found in this research article.

Instead of using the transect from baseline method, they propose the use of a change polygon method. This method includes overlaying the two contours atop each other, and calculating the change between them as a positive or negative area change. This method was introduced in this research article to mitigate the two issues with the transect from baseline method as described above.

2.4 Summary

In this section, a list of tools available to anyone working with historical maps and their vectorisation was reviewed. This is not a comprehensive list of all of the tools available, but an overview of currently commonly used software, their uses, advantages, and limitations. Being able to integrate any new methods or research with these existing tools would heavily aid in their usefulness as professionals already familiar with these tools would be able to more easily integrate their ongoing work with that new methods or research.

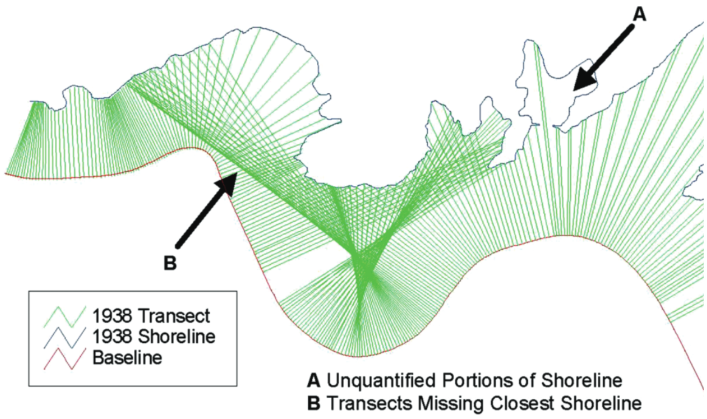

Figure 2.5: An example visualisation of two problems present with the transect from baseline method for measuring coastal change. Area A of the coastline shows portions of the coastline not represented in this methods analysis. Area B shows how the nearest coastline areas can be missed, and as such will overlap with other transect lines. This figure originated from Smith and Cromley (2012).

An overview of existing research in map vectorisation and comparison was also conducted in this section. Most research that exists in these areas is very specific to the maps they pertain to. There is a significant gap in providing a general automated method for either generating labelled vectorised maps from raster map images or comparing two different sets of map data against each other.

These reasons underline the significance of this dissertation. It acts as a starting point to explore different methods and techniques of vectorising general historical maps and comparing those vectorised maps. The methods taken are also planned to be as automated as possible to reduce the amount of labour intensive and error prone manual effort that is currently required to perform these image transformations and analysis.

Chapter 3

Design

In this chapter, the design of the tool produced as a part of this project to compare the accuracy of historical maps will be discussed. It will begin with the problem this tool aims to solve (in 3.1), and continues to explain an overview of the design (in 3.2).

3.1 Problem Formulation

The problem that this project aims to solve is to be able to provide a method or tool that can be used to measure the accuracy of historical maps. As discussed in Chapter 1, a tool such as this would be able to supplement existing and ongoing research in historical maps. Unfortunately, tools currently available to achieve this are unsuitable and require extensive manual input. As discussed in 2, while generic tools exist for vectorisation, simple visual comparisons, and warping, there is no off-the-shelf tool (at least not found in the research of this project) providing a set of features that are specific to historical maps and will produce a numeric error or difference value between two maps in an automated manner. The tool provided by this project is designed to solve this problem; to provide an automated method of measuring the accuracy of historical maps.

3.1.1 Identified Challenges

There are many issues which present themselves when working with this project. The first of which being extracting clean map data from the raster image itself. Although extensive scanned map images are available from various historical databases, these are not perfect. Additional artefacts within the map, while useful to any human reading the map, presents challenges to a computer trying to extract meaningful data from the image. Handwritten text data is prevalent in these maps, which presents a problem in identifying what is text and what is not. Misidentifying text means that text will be included in the

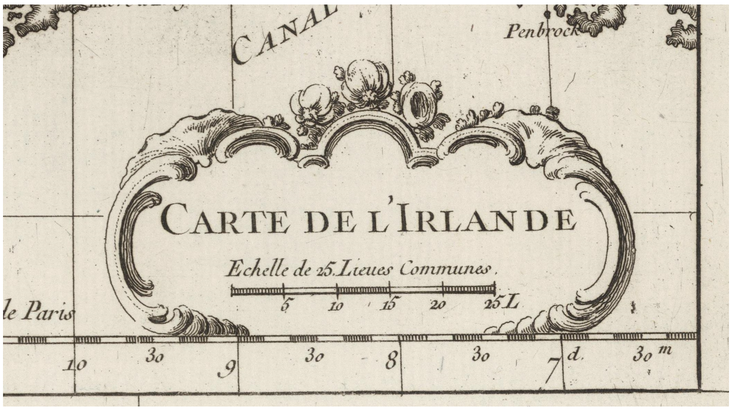

Figure 3.1: This figure highlights a decorative border present around text found in Figure 5.1; a historical map of Ireland from 1764. This border can often be confused as land coastline information.

coastline of a landmass, incorrectly reducing the accuracy rating as it includes that text boundary in the coastline.

Other artefacts in the image like decorative borders (See Figure 3.1) to text areas can also present a challenge, as it is very difficult to discern this from actual land. Particularly in single colour maps, longitude and latitude lines in the map can also be confused for land borders.

The projection used by historical maps can also change as mapping technologies change and for the specific maps use case. The most common map found is that of a Mercator map projection. For smaller land masses, the projection used to transform the curved surface onto a 2D plane has less of an impact than that of larger land masses. It is important to maintain the same map projection used, where that information is available to reduce any inaccuracies in measurement due to projection distortion. Brainerd and Pang (2001) presents a series of methods for visualising map projections and their distortions.

3.2 Overview of the Design

As mentioned in 1.2, the design of the tool produced as a result of this project is split into two distinct pieces of work; vectorising map images and comparing those vectorised maps. The following description of the design will show the methods used to mitigate the challenges presented to this project in 3.1.1. While these methods used are not perfect, this project contends that the result of using this tool significantly improves the time it takes to vectorise map images and provides a set of useful metrics to compare those vectorised maps. This dissertation will also discuss the limitations of the design proposed as a part of this project in 3.2.3.

3.2.1 Vectorising Map Images

The majority of the processing required to vectorise map images lies in cleaning the image. Historical maps will contain a lot of unneeded data that needs to be filtered out. These maps were multi-purpose and required inclusion of such extra information to fulfil this requirement. Having many points of information contained in the one map means it is more likely to be useful for any potential user of that map. However, a consequence of this is that there is a lot of visual noise in the map which is difficult to discern from the desired information (GanchimegGanbold, 2015). In this section, the steps taken to extract the relevant land boundary information (coastlines) will be discussed, and their implementation will be discussed in 4.1.

Obtaining text boundary data

Before performing any image editing operation, the text boundary information should be extracted from the image. It should be able to accurately identify, define, and interpret all text in the map image. As will be discussed in 4.1, this process is reliant on using a high quality copy of the scanned map image. In this regard, this step in identifying the boundary data of the text is done first to get the most accurate results. This data will then be used to subtract the text boundary area from the land area. This step is important in the cases for maps where text will "stick out" of land borders, so that that text will not be counted in the coastline contour.

Cropping the image to a manual setting

In many scans of historical maps, the original image of the scan will be the original unedited from the scanner. This will often contain additional unimportant space outside of the map area, such as blank white space, decorative page borders, or visible book spines and pages. The scans of historical maps can also found to be distorted because of practical limitations of scanning old paper maps. This optional step requires that the user of this tool select the four corners of the map. These four points would be selected in a custom GUI on the original map image. The tool will use these points to perform a perspective transform operation on the image and crop it to that selection. While keeping the aspect ratio of the image, the transformed image will also be normalised to an arbitrary set width. This helps to ensure that despite the resolution of any two maps being different, any operations done to those images will operate in a similar manner.

Threshold

In order to perform the following steps, it is required that the image given is a binary image. This is an image which only contains one colour channel and is only black or white. Thresholding can be done on an image in order to convert an colour image to a binary image. Generally, a threshold operation will take the greyscale of an image, and use a set value to compare each pixel against, resulting in a binary image. There are many methods of thresholding images each suitable for different use cases. The majority of map images are produced by high-quality scans of the original physical map. This results in these images having a uniform brightness across the image. As a result, this step does not require a method of thresholding that uses a variable threshold value such as adaptive thresholding. This project uses Otsu's Binarization (Otsu, 1979), which aims to select a threshold value which will apply to the whole image, such that the difference between the foreground and the background of the image is maximised.

Removing Vertical and Horizontal Lines

As mentioned in 3.1.1, longitude and latitude lines often found in historical maps can be confused with coastline data. The next step of this tool is to use a filter to remove any orthogonal lines from the image. The actual implementation of this filter will be discussed in 4.1.4, but this filter should produce a mask of the detected lines in the image. This mask can be used to isolate the longitude and latitude in the image, and subtract them from the binary image.

Morphological Cleaning

After the larger artifacts in the image have been removed (text data and orthogonal lines), the next step involves an "opening" and "closing" operation on the binary image. These two steps help to do the final cleaning steps of the binary image before it can be used to obtain the contour boundary.

Opening and closing morphological operations are fundamental operations in image processing. They both manipulate the image by iteratively applying erosion and dilation.

Erosion: Erosion works by essentially shrinking an image by an amount of pixels in all directions. It uses a structuring element which works similar to a convolution upon the image (although it cannot be expressed as strictly a convolution). This is usually used to remove small objects in an image, such as noise, separate marginally touching objects. However it does have the effect of reducing the size of any objects in the image.

Dilation: Dilation works as the opposite to erosion, it will essentially grow any objects in the image by an amount of pixels in all directions. Similarly, it uses a structuring element which describes the size and shape of the dilation. This is usually used to fill small holes in objects and connect closely located objects in the image. It will also thicken any objects in the image.

Opening Operation: An opening operation is performed by doing an erosion followed by a dilation using the same structuring element. Using these two operations in sequence will remove smaller objects while preserving the shape and size of any larger objects in the image. The remaining objects will also be smoothed as minor details are removed.

Closing Operation: A closing operation is performed by doing a dilation followed by an erosion using the same structuring element. Using these two operations in sequence will fill any remaining holes and can connect any closely located objects while retaining the overall shape and size of any objects in the image.

In this step, an opening operation will be done first to reduce noise in the image. Next, a closing operation will be done to connect any slightly broken objects in the scene. These could have been caused by a number of reasons in the previous steps. Where a longitude or latitude line will intersect a land boundary and result in a "cut" of pixels into the land. See Figure 5.4 for example of this "cut" effect from longitude and latitude lines. Also where text intersects with the land, the removed text boundaries can disconnect two sections of a coastline. The closing operation will restore those pixels at the land boundary, but will not regain the whole longitude or latitude line, stopping only at the land boundary.

Obtaining the contour boundary

After the previous cleaning steps, obtaining an accurate boundary of objects in the scene becomes possible. Any connected pixels in the image can be described by their contour. Ideally this would be only the coastlines of any land in the map image, but often there will be additional objects left. Only the largest contour boundary was used as the main object in any image. While other contours will exist in an image, the largest one should represent the main object in the scene.

Manual Inspection and Corrections

After obtaining a list of contours in the map image, this will be shown to the user of the tool. At this point, the user can review the resulting coastline contour and fix any remaining errors. The user can optionally add or subtract any contour shapes through the use of a custom GUI. While the goal of this tool is to have minimal manual intervention, this step is necessary to ensure that the tool is working properly and to fix any errors that may have occurred throughout this process. Finally, the resulting contour is saved and can be used to compare against other maps which have undergone this process.

3.2.2 Comparing Vectorised Maps

Once at least two map images have had their coastline data extracted, they can be compared against each other for their similarity. In this dissertation, the effectiveness of four difference metrics will be assessed; the sum of absolute differences, Hausdorff distance, cosine similarity, and template matching.

Sum of Absolute Differences

The sum of absolute differences is the simplest of the difference metrics used in this project. This process is similar to the change polygon method presented by Smith and Cromley (2012). A filled polygon of both contours is constructed, and they are subtracted from each other. To compute this, the two contour polygon areas are converted to a fixed size raster image, where the filled polygon is centered in that image, and the count of the pixel differences between the two is taken as the value returned by this metric. A visualisation of this is shown in Figure 5.8.

Hausdorff Distance (Huttenlocher et al., 1993)

Hausdorff distance describes the maximum distance any point on one contour is from the closest point on the other contour. Since this metric only relies on the maximum value, it is more susceptible to outlier points on the coastline than with the other three metrics discussed here.

Cosine Similarity

Traditionally, cosine similarity is used to measure the difference in any two vectors. It is calculated by computing the cosine of the angle between those two vectors. In the context of image processing, an n-dimensional vector describing features in the image is generated and used to compare against the vector from another image.

Template Matching (OpenCV (2024b), Brunelli (2009))

Template matching is a process defined in OpenCV, designed to find any locations in a larger image that closely resembles a provided smaller template image. This is done by using the template image as a convolution over the larger image and obtaining the similarity between the two at each pixel. The similarity metric between the section of the larger image and the smaller image can be chosen from a few various options provided by OpenCV, such as the sum of squared differences, correlation coefficient, or normalized cross correlation. In this project, this process is used in an alternative way to make use of the difference metric provided. Instead of providing a larger image and a smaller image, two images of the same size are provided, and a single difference value is returned. These two images are obtained by filling a polygon described by the contour of the two images being compared.

3.2.3 Limitations

This dissertation proposes these steps to produce an accurate method of both vectorising map coastline data and comparing that map data against other similarly vectorised maps, however it does have a number of limitations which could be expanded upon in future work.

Errors in text boundary data: False negatives in the text boundary data will result in a less accurate representation of the coastline. This is difficult to work around, since no optical character recognition (OCR) software is perfect. Even in cases where the OCR does work, since text can often cross the coastline boundary, this can result in having a bay in the land where one does not exist.

Vertical and Horizontal land borders: Land whose borders are significantly vertical or horizontal can be filtered out by the step which removes those lines. This can present issues for internal borders which were not chosen based on geographical features, but does not often occur for coastlines, as is the focus of this project.

Multiple contour boundaries: By ignoring any contours but the largest, important information could be lost such as important islands which a user would wish to include in the comparison. This will also be an issue for a single country, where it is split across multiple islands, even if they are very closely positioned, such is the case for New Zealand or Japan.

User Input: As mentioned before in 1.2, "The objective of this project is to develop an automated system to...", while most of the steps in this design are done automatically, this still requires some manual user input. This takes place in the form of selecting the corners of the map from the overall scan, and potentially cleaning the resulting contour data. Both of these points of user input are optional for the operation of this tool, but they can be essential in order to obtain accurate and useful results.

Rotation: While the steps presented here are designed to be invariant of scale and translation, it is not invariant of rotation. For this design, it is assumed that the maps provided are always in the same rotation, that is that vertically upwards in the map images is geographic north. A solution which is invariant of rotation could be discussed in future work.

3.3 Summary

This chapter described the approach to the design of the tool proposed in this dissertation to vectorise the coastline of historical map images, and four different difference metrics to compare that coastline data. It began with identifying the problem this project aims to solve along with any challenges expected to be faced. Then an overview of the design of the tool was given, including all of the steps to be taken to vectorise map images and compare vectorised image data. Lastly, this chapter ended with the limitations of this design, which could be explored further in future work.

Chapter 4

Implementation

In this chapter, the implementation of the tool used to vectorise and compare historical maps is discussed. It intends to describe how the design defined in Chapter 3 was implemented and how the research objectives described in 1.2 were achieved. The structure of this chapter is similar to Chapter 3, explaining in more detail how each step of the design was implemented.

The main technologies used in this project were Python, OpenCV, and Optical Character Recognition (OCR) software. Python is the programming language used in this project to provide the platform to produce the tool described in this dissertation. OpenCV (Open Computer Vision Library) is a commonly used image processing library which contains much of the functionality needed to perform the steps defined in Chapter 3. The specific OCR library that was used in this project was EasyOCR (Jaided AI, 2024).

4.1 Vectorising Map Images

4.1.1 Obtaining text boundary data

Optical Character Recognition is still a difficult problem in the field of computer vision. As such, the development of an OCR specifically for the use of historical maps is outside of the scope of this project. Despite this, the benefit of using an OCR is still significant in producing accurate coastline boundary data. As mentioned at the beginning of this chapter, the EasyOCR library was used. While some other OCR software was explored, EasyOCR was found to be more effective at text detection in the context of historical maps. Since most historical maps include irregular fonts, odd angles, noise, and can often be handwritten, identifying and interpreting text within a map image is not always perfect. Although it did not very often correctly interpret the text correctly, EasyOCR was able to identify text locations accurately, which for the purposes of this project is what was required from this OCR software. EasyOCR returns a list of polygon corner data along with their interpreted text. Using the polygon data, the text in the image can be subtracted from the image after the thresholding operation. As mentioned in 3.2.1, this step is done first to the original image, before any other step can distort the text data and the result is saved for when the threshold of the image is obtained.

4.1.2 Cropping the image to a manual setting

This is done by allowing the user to select four points on the image selected. Because the image provided can be a high resolution scan and this often cannot be all shown on a monitor, care is required to maintain the position of the selected points in relation to the actual pixels on the original image. After the four points are selected, the image will undergo a perspective transformation to crop the image to only include the actual map data. The aspect ratio of the image is preserved from the original and the specific width of the cropped image was chosen arbitrarily at 1024. If a map should not be the same aspect ratio as the image it originated from, this would require additional input from the user to specify that desired aspect ratio.

4.1.3 Threshold

As described in Chapter 3, this tool uses Otsu's Binarization in order to obtain a binary image from the cropped map from the previous step. This is done in two stages, where first the image is converted to a greyscale image. While the use of only the blue channel of the image to further separate the value difference between ocean and land could be effective for some maps, using an approach which uses all three colour channels to produce a greyscale image proved more effective at producing a greyscale image that supports Otsu's Binarization.

4.1.4 Removing Vertical and Horizontal Lines

In order to create a filter which isolates horizontal and vertical lines in the binary image, this project makes use of a custom structuring element used alongside a set of erosion and dilation operations. Through experimentation, a value for the structuring element for both the vertical and horizontal lines was found to be optimal at the width or height of the image divided by 25. Taking the horizontal lines as an example, a structuring element of size $(width/25, 1)$ is constructed and used to do a single erosion followed by a dilation. This will isolate any long, thin, and horizontal lines in the binary image.

This results in any lines longer than a 25th of the image being isolated, another dilation step is used to increase the width of any detected lines in the mask to better remove them from the binary image. Finally, the mask produced is used to subtract the lines from the binary image.

4.1.5 Morphological Cleaning

In this step, an opening operation followed by a closing operation was used to first remove any noise in the image. Generally the amount of noise left in the image after the previous steps is quite small. As a result, only a 3x3 structuring element over 1 iteration of erosion and dilation is used to remove any "salt and pepper" noise in the image. Connecting remaining contour data requires a larger amount of work. A 5x5 structuring element used over 5 iterations of dilation followed by erosion is used to connect any slightly disconnected land in the image. While this step is much less involved than the previous, it is essential as a final correction of the image so it can be used to obtain the desired contours present in the map image.

4.1.6 Obtaining the contour boundary

Using a binary image, obtaining the contour boundary in OpenCV is very simple.

But this will return all contours present in the image, not just the main contour interested. For this project, the contour with the largest area is used as the coastline data used for comparison. Any other contour information is still saved in case the user wishes to fix any errors present, such as connecting any two contour areas. The mean area and standard deviation of the set of contours found is also saved and used. Any area which is larger than the mean area plus one times the standard deviation is detected as potential land, and shown to the user in a different colour from areas which are below this threshold. This selection criteria is not perfect, and is prone to false positives and false negatives. Despite this, this project proposes that this is a useful visualisation to the user when viewing and editing the contours present in the image. The contour with the largest area is also shown in a different colour from the other contours present to distinguish it further.

4.1.7 Manual Inspection and Corrections

The last step in the vectorisation process is the manual inspection of the contours detected. This allows the user to view any contours found and edit them if there are any inaccuracies present. The editing process is done by allowing the user to select an area to add or subtract from the contour. Adding can result in conjoining two separate contour areas, and subtracting can result in splitting one contour area into multiple separate contours. In this tool, the user is able to swap between the two edit modes, and drag over the area they wish to edit. With each point of the cursor on the image, a polygon of the area can be constructed. The implementation of additive editing involves drawing a polygon of the contour on a fresh binary image, and adding the user's polygon on top of this. The contour of the binary image is computed again with the users edits, and the process can be repeated again until the user is done making any edits. Similarly to additive editing, subtractive editing removes the users selected polygon from the contour polygon area.

Once the user is finished making any optional edits to the contour, the points are saved, and can be used to compare against any maps which have also undergone this process.

4.2 Comparing Vectorised Maps

The implementation of the metrics used is less involved than that required to vectorise the map image. Before the contour data can be fed into each of the metrics explored in this project, they need to be appropriately imposed on top of each other. Since any comparison between map data should be invariant of scale and position within the map, any contour data used should be transformed to remove any differences in scale and position. This is done by scaling the contour to fit within the bounds of a square, and translating one of the contour sets so that its centre is in line with the centre of the second contour set. The centre of the contour area is found by its centroid. Although fitting the contour to the bounds of a square will distort the image and could introduce inaccuracies in its results, this dissertation asserts that it is necessary in order to have a fair comparison between different historical map images. Once the contours are fitted to each other, they can be given as an input to the metrics explained here.

4.2.1 Sum of Absolute Differences

For the implementation of the sum of absolute differences, the two contours given are used to create two binary images containing only the polygons defined by those contours. The difference between these two binary images is obtained and the count of the pixels remaining provides the metric for the sum of absolute differences. This metric provides a direct area difference between the two contours.

4.2.2 Hausdorff Distance

The Hausdorff distance between two contours can be calculated using OpenCV. This method will return a number which represents the maximum distance from one contour to the other (see 3.2.2 for a more detailed explanation). For every point on one contour, the distance to every point on the other contour is calculated. As a result, the Hausdorff distance algorithm has a computational complexity of $O(N^2)$, where N is the amount of points in a contour. Because of the computational complexity of this algorithm, the method provided by OpenCV takes significant processing time over the other metrics in this dissertation.

4.2.3 Cosine Similarity

In order to use the cosine similarity algorithm as a difference metric, an n-dimensional vector is required as an input to the function. For this project, the binary image produced from the contour is flattened into a one dimensional array and used as the input to this function. While this is an easy-to-implement method of obtaining a vector from the image for the use in this algorithm, creating a more suitable vector of the image would produce more accurate results than what was used in this project. One such example of a method to obtain this vector would use a Convolutional Neural Network (CNN) to measure specific features in the image. Developing such a CNN is outside of the scope of this dissertation, but should be explored in future work in order to make better use of this difference metric.

4.2.4 Template Matching

As mentioned before in 3.2.2, OpenCV's template matching usually takes a larger image to search through and a smaller image to be used as the template. In this project however, two images of the same size containing only the area defined by the two contours being compared is given as an input to template matching. This results in a single metric of how similar these two images are to each other in the bounds of 0 to 1. There are multiple calculation methods available for template matching. In this project, the normalised square difference was used to obtain the difference metric produced by template matching.

4.3 Summary

In this chapter, the implementation of the tool described in Chapter 3 was discussed. This tool aims to contribute towards solving the problem described in 1.2, where a tool would be produced which can aid in the automatic scoring of historical maps from different sources or periods.

Chapter 5

Evaluation

This Chapter presents the results of the two systems developed as a part of this project. In the experiments overview (5.1), it will begin with a detailed inspection of each step in the vectorising and comparison process when applied to a single historical map. In the results section (5.2), it will show the results of multiple different historical maps when compared against modern day maps using the difference metrics explored in this dissertation. At the end of each section for a difference metric, its effectiveness will be assessed. Lastly, an experiment will be conducted on a set of constructed images in 5.3.

5.1 Experiments Overview

The following experiment using the tool to vectorise a historical map will be conducted on a map of Ireland. The map chosen for this experiment can be seen on Figure 5.1. This is only chosen as a baseline for comparison, the tool is not limited to only Ireland. Examples of results of vectorisation of other images can be seen in Figure 5.16 from Figure 5.15, and four more examples conducted on historical maps are included in the appendix of this dissertation (See Figures 1, 2, 3, and 4). This map of Ireland was chosen as an example for this experiment to show the results of each of the different steps taken in the vectorisation process. In some other maps explored, such as those found in the appendix, some of the steps in the vectorisation process are not visually apparent. After the map chosen has been vectorised it will be used in a comparison against a modern day map. The modern day map chosen is that taken from Google Maps (Figure 5.7, left).

5.1.1 Vectorisation

The map chosen is a high resolution scan of the map found in Figure 5.1. The following sections will show the steps taken with this figure, and at each step show the processing

Figure 5.1: A high quality scan of a map of Ireland, originating 1764 (Stanford Libraries, 2024).

done to the image in that step.

Text Boundary Data

Figure 5.2 shows the result of EasyOCR's identified text boundary and the interpreted text. The figure shows a zoomed area of the text boundary for clarity. As can be seen in this figure, not all of the text present in the image was identified. It also does not often correctly interpret the text when it is found. While it did correctly interpret "wne" of "Downe", it did not identify the whole word as text, missing "Do", since it is overlapping with the coast. Even when it does identify the whole text boundary, such is the case for "Carikfergus", it does not interpret the text correctly, returning "Etatikfergu". Despite the errors present, being able to identify so much of the text on the outside of the coastline helps immensely in distinguishing that text from the coastline.

After the text locations have been identified, this information will be saved for after the image has been transformed into a binary image, see 5.1.1. The text locations are used to subtract them from the binary image. The result of this process can be seen in 5.2, on the right. Because of the errors present from the OCR, not all of the text in the image is removed, but despite this, almost all of the text on the coastline is removed. This helps the result to have an accurate outline of the coastline present in the map with minimal manual edits required.

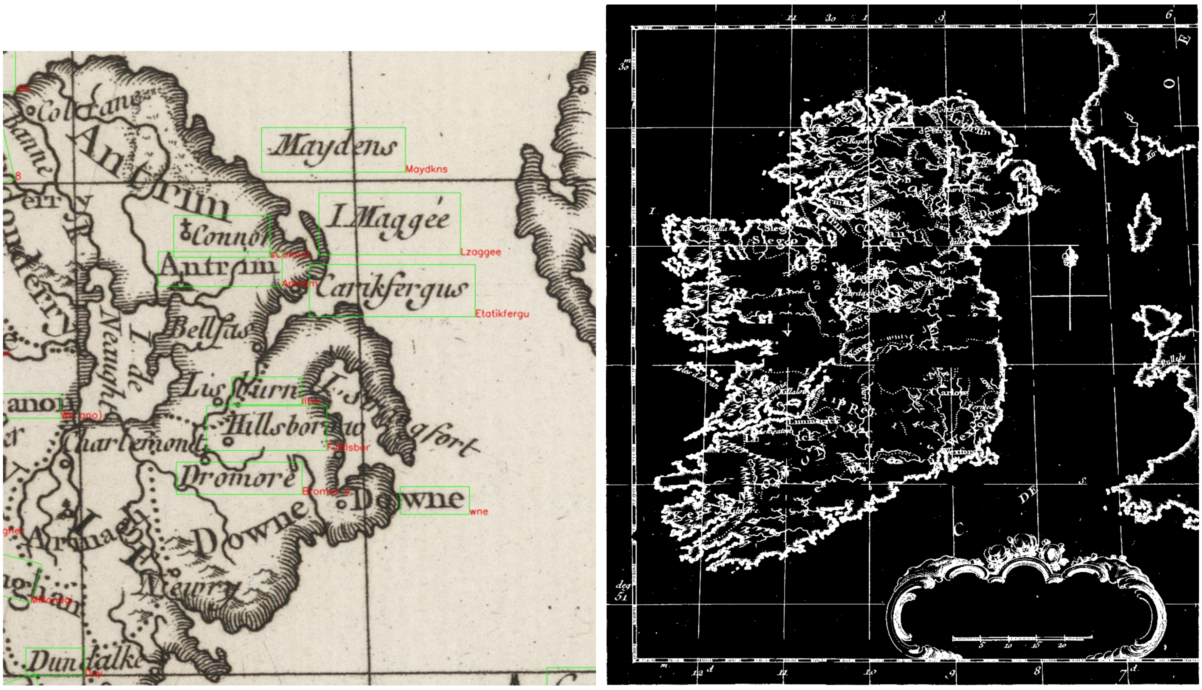

Figure 5.2: Left: A Zoomed image showing the detected text boundary data and interpreted text information.
Right: The result of removing the text boundaries from the threshold image of Ireland

Image Cropping

After the image is chosen, the user selects the four corners of the desired map information. Additional parts of the image can be seen which are not relevant to the map itself, and as such should be cropped out of the image. For example, in Figure 5.1 the borders of the book that this map appeared in can be seen. Removing this additional detail helps to normalise the process across different maps which may have similar additional parts in their scanned image. The result of this can be seen in Figure 5.3 on the left.

Threshold

Figure 5.3 on the right shows the result of a threshold operation executed on the cropped image seen left. While this process used Otsu's Binarization to select a threshold value,

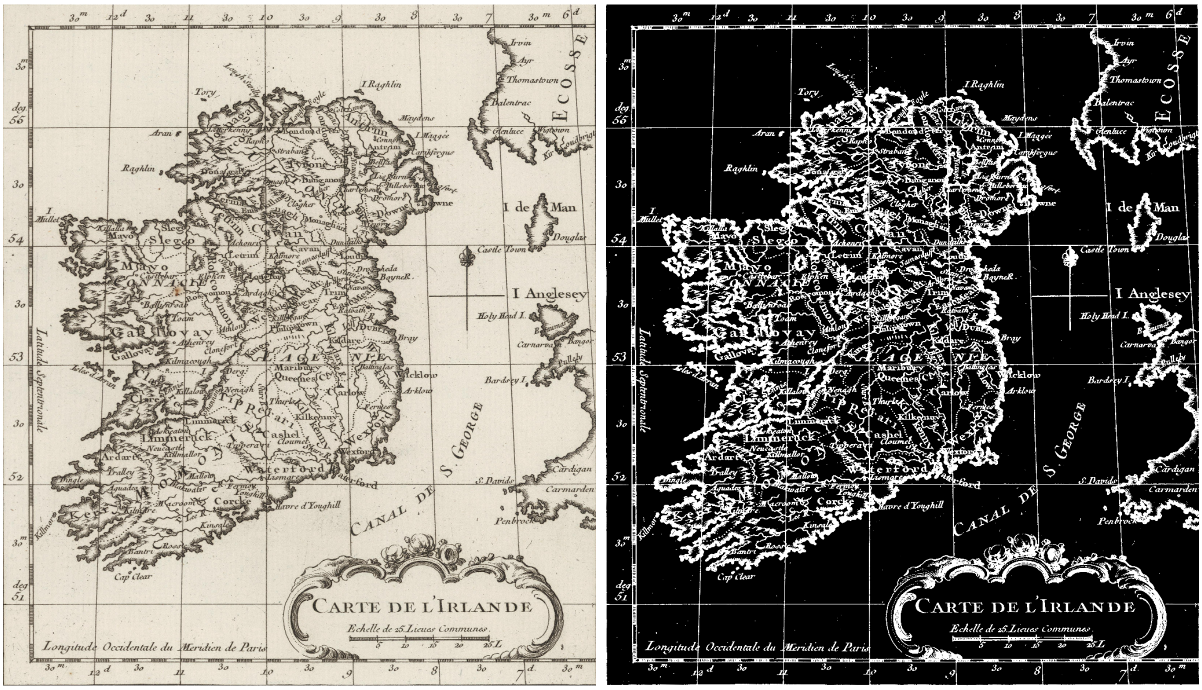

Figure 5.3: Left: Four corners of the map content have been selected, and used to perform a perspective transform to the image seen here. Right: Thresholded binary image of Ireland using Otsu's Binarization for the threshold value.

a Guassian Blur step was executed before this. The blur step helped to filter out any noise present within the image and provide a cleaner histogram of image values for Otsu's Binarization to operate on. A consequence of this algorithm, while not present in this example, is that the foreground (white pixels) may not be the boundary data. If the background in the original image is darker than the foreground, this process may result in the ocean being selected as the foreground. This can be resolved by simply inverting the image after this step in the cases where this occurs.

Vertical and Horizontal Lines

Figure 5.4 shows the result of removing any vertical and horizontal lines present in the image. While successful as removing many of the lines present, the process as described in Chapters 3 and 4 is not as effective as desired. This process does not completely remove all of the lines, since they are often at slight angles. It does however make many of the lines thinner, which means they will be removed in the next step, when an opening operation is done to this image. This step will also disconnect some of the land boundary where those lines exist. Other designs which do not have this effect should be explored in future work, since connecting those contours should not be reliant on the morphology operations in the next step.

Figure 5.4: Horizontal and Vertical lines have been filtered out of the image.

Morphological Operations

Figure 5.5 shows the result of performing an opening operation followed by a closing operation on Figure 5.4.

Contour Boundary and Manual Corrections

After the image has been sufficiently processed in order to isolate the coastline data, any contours present in the image are extracted. The contours found in this experiment can be found in Figure 5.6 on the left. This figure shows the boundaries found overlayed on the cropped image of Ireland, seen in Figure 5.3. The main contour of Ireland is seen in green, selected by being the contour with the largest area. Other smaller contours are identified in red and blue, such as the decorative border around the text at the bottom

Figure 5.5: Remaining noise and disconnected areas have been removed and connected together by the use of Opening and Closing Morphology operations.

and the UK at the right. At this point, if there are any errors present, the user can edit them at this stage. This could include closing any disconnected sections of the contour, or unwanted connected sections, such as disconnecting smaller islands or removing text areas identified as land. In Figure 5.6, no such manual corrections were made.

5.1.2 Comparison

Once the contours have been extracted from at least two images, those contours found can be compared against each other. For the following comparison, the image from 1764, seen in Figure 5.1, with the resulting contour boundary, seen in Figure 5.6, was compared against a modern day map of Ireland from Google Maps. The modern day map and its contour can be found in Figure 5.7. The results of this comparison will be discussed in 5.2, along with the results of other maps, but a visualisation of the differences of these two contours can be seen in Figure 5.8. In this comparison, while both contours are

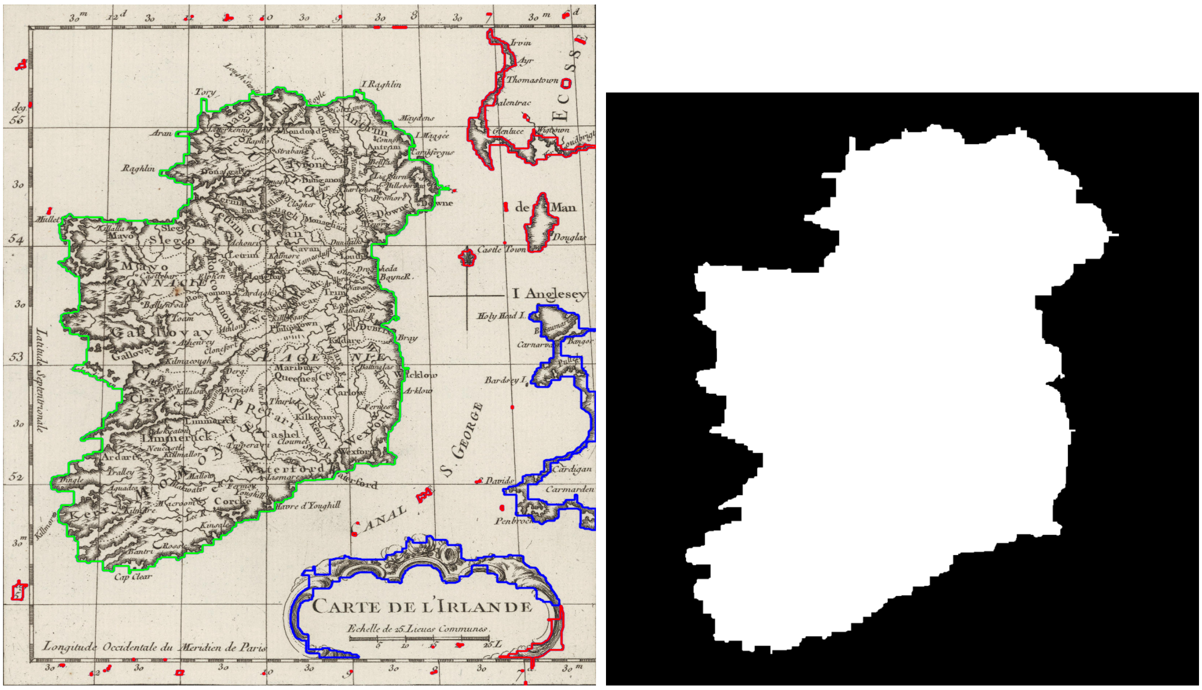

Figure 5.6: Left: The contour information found within the image is used to label the original image (Figure 5.1) with that contour information.
Right: The resulting filled contour of Ireland from the process outlined in this dissertation.

centralised on top of each other and normalised in size, it can be seen that in general, the map from 1764 is wider in most areas of the map. The calculation for the sum of absolute differences for this comparison involves the sum of the areas seen in this figure, of both red and green areas. On the right of Figure 5.8, the Hausdorff distance found for these set of contours is visualised. This value for the Hausdorff distance could be improved however, since the map from 1764 includes the Aran Islands, while the modern day map does not. In practical applications of this tool, ensuring that both maps used have the same content, in terms of the country and the included islands or additional land would be the responsibility of the user. However, for the demonstration of this tool this manual correction was not done. A visualisation of cosine similarity or template matching in this context is not feasible, see 3.2.2 and 3.2.2 for an explanation of these difference metrics.

5.2 Difference Metrics Results

For the results of each of the difference metrics, a number of historical maps from different sources and different periods of time were used.

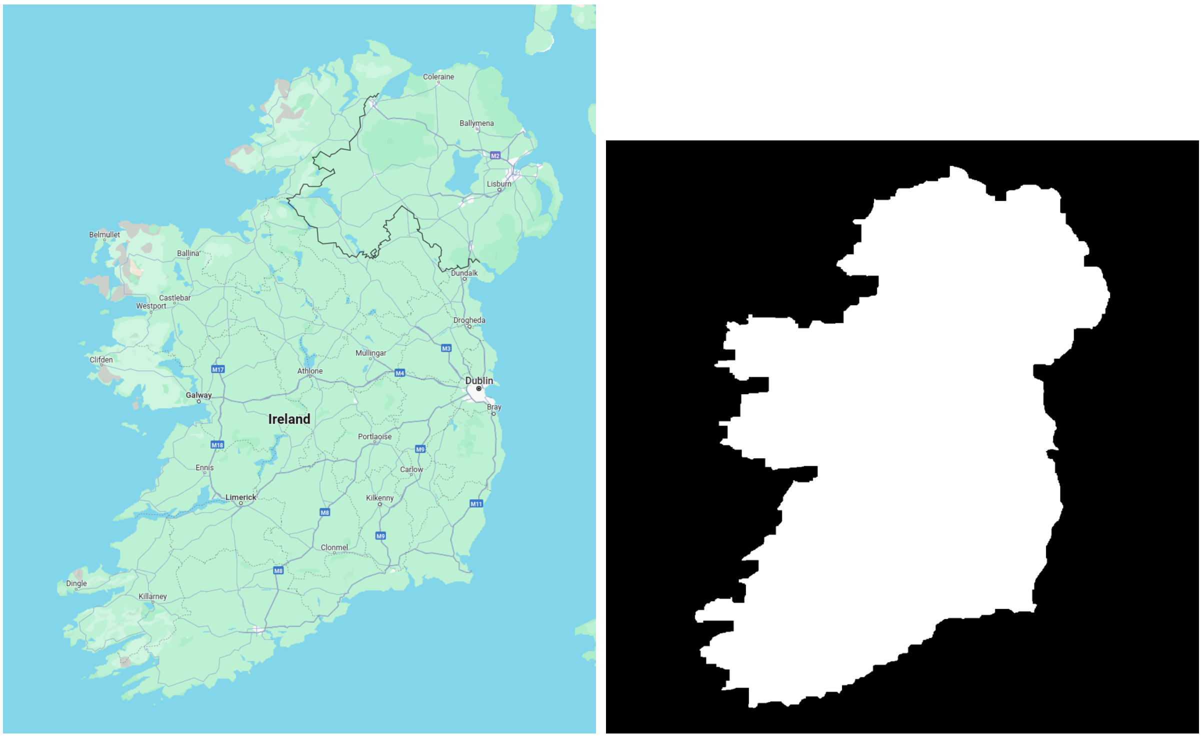

Figure 5.7: Left: Modern day map of Ireland, from Google Maps.
Right: Contour boundary of Ireland from the modern map seen left.

1692: This 3206x3819 scan of a map of Ireland was used as the oldest historical map in this evaluation. It was produced by Giovanni de Rossi, Augustin Lubin and Antonio Barbey around the years of 1692-1699. It shows a approximation of the land of Ireland as was known at the time.

1764: This was the map used as the example for vectorisation in 5.1.1. It has a resolution of 5930x6816 and was produced by Jacques Nicolas Bellin in 1764. Even without impirative analysis, it visually appears to be significantly more accurate than that from 1692. It also contains areas of Britain on the right hand side of the map image.

1825: This map image has a resolution of 6022x7517 and was produced by Aaron Arrowsmith Jr. in 1825. Compared to the previous maps, it has text completely surrounding the entire map, that without any text removal would impact the accuracy of the detected coastline greatly.

1888: This map image has a resolution of 4726x5647 and was produced by Karl Sohr, Heinrich Berghaus, Friedrich Handtke, and H.J. Klein in 1888. It has a coloured background, which is darker than the map content itself, has text surrounding the entire coastline, and also has shallow water around the coastline in a darker colour. The combination of these complications in this image presented more challenges in

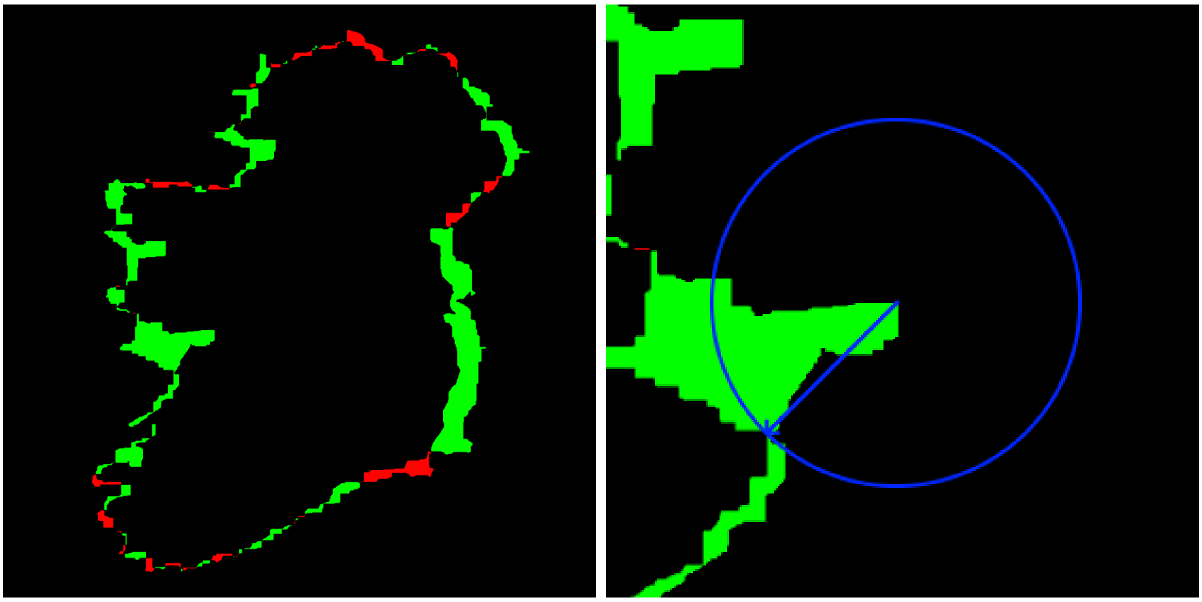

Figure 5.8: Left: This figure shows the area difference between the two contours of Ireland, one from 1764, and one modern day. The green area shows the area which was contained within the map from 1764 but not in the modern day, and the red area shows the contour area which was contained in the modern day map but not in the map from 1764. Right: A visualisation of the Hausdorff distance found in this comparison. The blue arrow represents the furthest distance one would have to travel to get from a point on one contour to the minimum distance to the other contour.

vectorising and labeling the coastline of this map than the other maps used in this project. As will be discussed in , the outline of the shallow water ended up being the contour found in this map.

1904: This map image has a resolution of 5514x6876 and was produced by Edward Standford in 1904. It is a digital composite of four different map sheets and is the most detailed out of the maps analysed in this project. The amount of small text in this image was significantly more difficult for the OCR to recognise, and as such, many text areas were left included in the final contour, as can be seen as the spurious protruding contour areas in Figure 5.10, bottom middle.

The scans of these maps can be seen in Figure 5.9 and each of their respective contours can be seen in Figure 5.10. While the resolution of the original images are not all the same, they are roughly similar. The vectorisation of these images involves normalising these to the same resolution, so that any differences that result from having differing resolutions in that process is minimised. However, the resolution of the images does impact the performance of the OCR software used. It is very likely that by using a lower quality image of each map would reduce the effectiveness of the OCR used in this project. Each of these historical maps will be compared against the modern day map, seen in the bottom right of Figure 5.9 and in Figure 5.7. A table of results for each of these maps can be seen in Table 5.1.

Map Date	**S.A.D.**	**Hausdorff**	**cosine**	**template**
1692	87960	21.63	0.1	0.206
1764	57844	8.94	0.066	0.125
1825	49914	5.39	0.17	0.108
1888	55104	12	0.139	0.119
1904	54610	3.16	0.0998	0.117

Table 5.1: Difference metrics results of a series of historical maps of Ireland to a modern day map of Ireland. The metrics shown here are: the sum of absolute differences (S.A.D.), Hausdorff distance (Hausdorff), cosine similarity (cosine), and template matching (template).

5.2.1 Sum of Absolute Differences

Figure 5.11 shows a bar chart highlighting the results of the comparison of a series of historical maps, compiled in Table 5.1. This Figure focuses on the results of the sum of absolute differences as a metric for comparison. The values shown in Table 5.1 and Figure 5.11 give the area difference between the two contours. This metric proved to be a useful aid to intuitively understand the differences between two maps contours. Having a visual understanding behind the statistic provided by this difference metric makes it uniquely useful. Its ease of understanding of this metric, along with the possibility of a visualisation such as that in Figure 5.8 highlights this metric over the other metrics discussed in this dissertation.

5.2.2 Hausdorff Distance

Figure 5.12 shows a bar chart highlighting the results of this difference metric when applied to a series of historical maps of Ireland against a modern map of Ireland. This chart shows the clear downward trend that the maps used in this comparison have in their inaccuracies when compared to a modern map. This trend is clearer than what can be seen in the other difference metrics used in this dissertation, see Figures 5.11, 5.13, and 5.14. It does however present an unusual spike in the measure for the map from 1888. This highlights one of the likely drawbacks of using the Hausdorff distance as a difference metric; Hausdorff distance is more susceptible to outlier data than the other difference

Figure 5.9: The images shown here are a series of Historical maps of Ireland from different time periods. From the top-left these maps were produced in 1692, 1764, 1825, 1888, 1904, and the bottom-right is a map obtained from Google Maps (accessed 2024). All maps apart from the bottom-right can be found at Stanford Libraries (2024).

metrics in this dissertation. Since it relies on only maximum distances, if there are any outliers in the coastline present, this metric is disproportionately impacted over the other metrics. Assuming the map image was vectorised appropriately, this metric provides a clear and accurate result for measuring the difference between two contours.

Figure 5.8 right shows a visualisation of the hausdorff distance when applied to the map of Ireland from 1764 (see Figure 5.1) when compared to a modern map (see Figure 5.7). In this particular case, the Hausdorff distance found was a result of a vectorisation error, where one map included additional islands, while the other did not. If the contour from the map from 1764 did not include this bay, the Hausdorff distance chosen would have been from elsewhere on the figure shown. Another method similar to Hausdorff distance which is not so disproportionately impacted by these outliers may be more appropriate

Figure 5.10: The six contours seen here are obtained from the maps in Figure 5.9. These were made using the vectorisation process as outlined in this dissertation. Minimal manual corrections have been made to these contours.

for this application. Using an average distance travelled instead of only the maximum, or discounting a top N% of distances may provide a better metric in this scenario. The exploration of such a metric could be explored in future work.

5.2.3 Cosine Similarity

Figure 5.13 shows the results from the cosine similarity from its column in Table 5.1. The results for this metric did not seem to follow any significant trend with the images presented to it, unlike the other three metrics discussed in this dissertation. As a result, it is likely that this different metric is not very appropriate when applied to this context, or that it should be utilised in another way. Despite the results of cosine similarity being inconclusive in comparison to the other metrics discussed, it should not be discounted, as it is likely it is used inappropriately as implemented in this project. As will be discussed in 6.1.2, the way in which cosine similarity could be implemented should be expanded upon in future work.

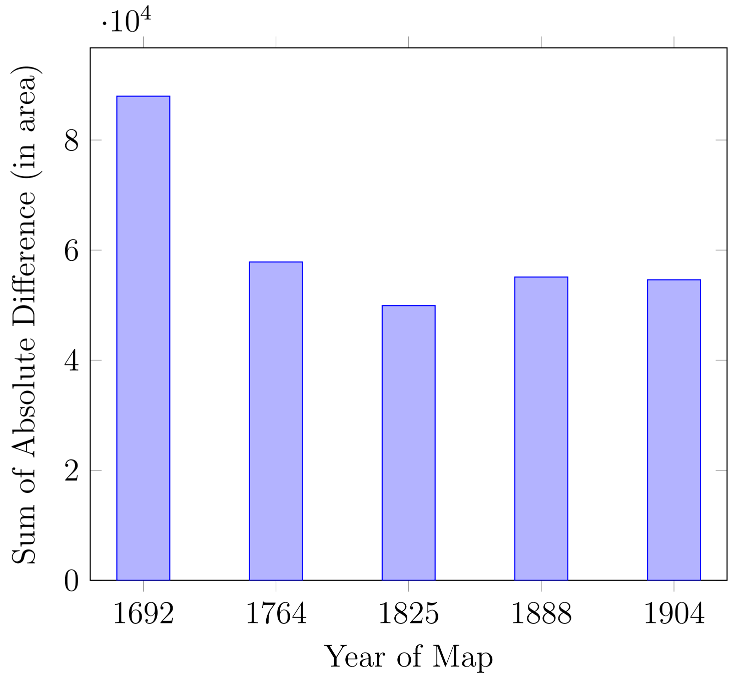

Figure 5.11: Using the values in Table 5.1, this bar chart plotting the results found for the sum of absolute differences was made. A trend showing a lower difference area as time progresses can be seen.

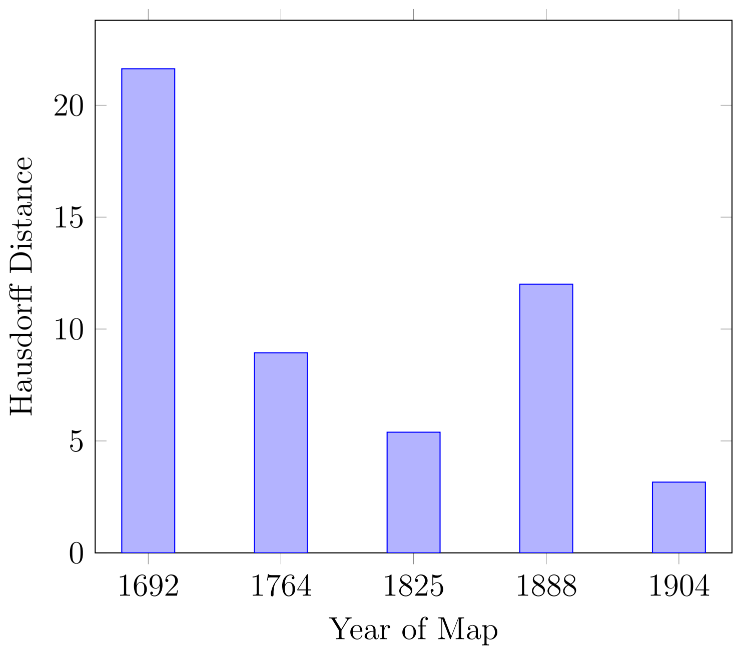

Figure 5.12: Using the values in Table 5.1, this bar chart plotting the results found for the Hausdorff distances was made. A trend showing a lower difference distance as time progresses can be seen.

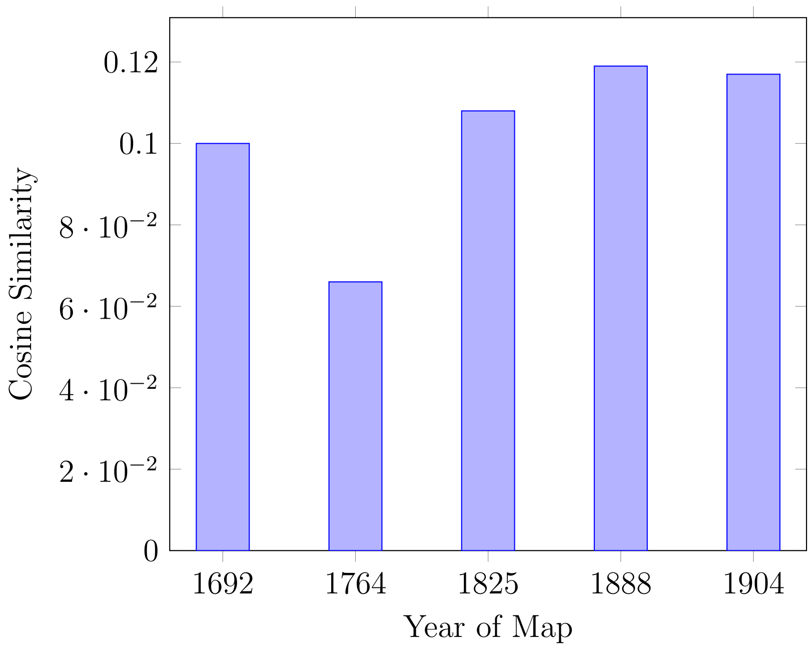

Figure 5.13: Using the values in Table 5.1, this bar chart plotting the results found for the cosine similarities was made.

5.2.4 Template Matching

Figure 5.14 shows a bar chart with the results found by the template matching difference metric when applied to the contours of historical maps of Ireland, with the specific values in Table 5.1. Similar to most of the other difference metrics, a downward trend can be seen, indicating an decrease in the difference between these contours in respect to their creation date. The results for template matching very closely resemble those from the sum of absolute differences, providing additional insight to their similarities. Although in theory they are different metrics with different computation behind them, their proportional similarity shows that their results are influenced by the same features in the differences of each map.

5.3 Additional Experiments

As a supplementary experiment, the same vectorisation and comparison process was performed on an example set of six constructed images. These are used as a second standard to ensure that the results of comparison as discussed in this evaluation so far have not been a product of only working with maps of Ireland. These example images can be seen in Figure 5.15. The results of vectorisation, i.e. the contours found, can be seen in Figure 5.16. The results of comparing each of these contours against that found in the

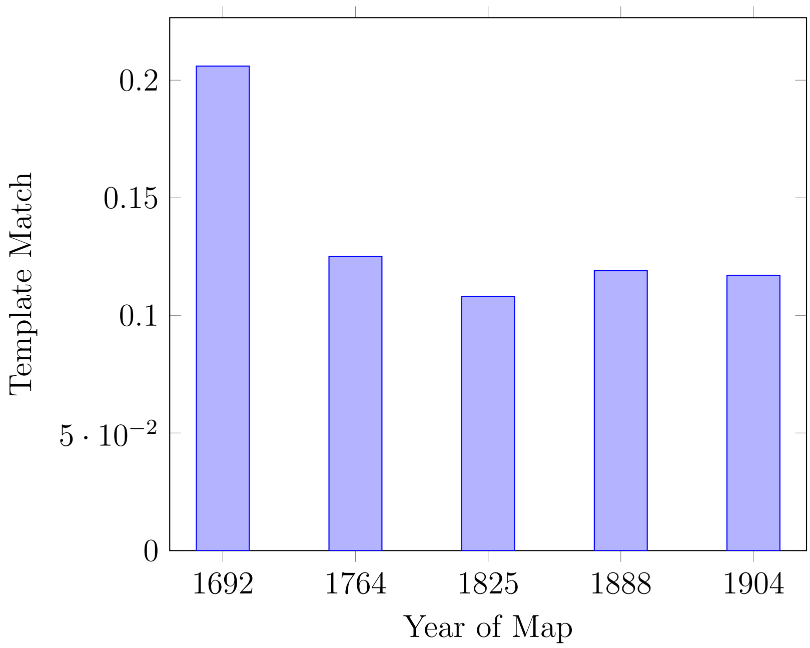

Figure 5.14: Using the values in Table 5.1, this bar chart plotting the results found for the template matches was made. A trend showing a lower difference value as time progresses can be seen.

circle image in the top left of Figure 5.15 can be seen in Table 5.2.

5.3.1 Evaluation of Constructed Images' Vectorisation

Figures 5.15 and 5.16 are an illustration of vectorisation of constructed images. The contours of the lightly distorted, heavily distorted, rounded square, and circle images were extracted easily, the images with the grid of lines, and the text overlayed required a minor amount of manual input to fix. While the lines were successfully removed, it left small gaps in the circle where the lines used to be. As a result, minor edits were required by using the GUI provided by this tool to fix those gaps in the contour. In this experiment, the curve of the circle was not followed in these manual edits, instead straight lines were drawn to correct any disconnected segments of the contour. This is not clearly apparent in the final contour of the image with lines. However, for the final contour of the image with text, this is shown more clearly. Since all of the text present in the image was detected, larger areas of the contour were removed. Without additional information present on the inside of the contour, like is often the case for maps, segments of the contour were left disconnected. Similarly for this image, straight lines were manually drawn from the disconnected areas. This is particularly visible for where the word "Pellentesque" was detected.

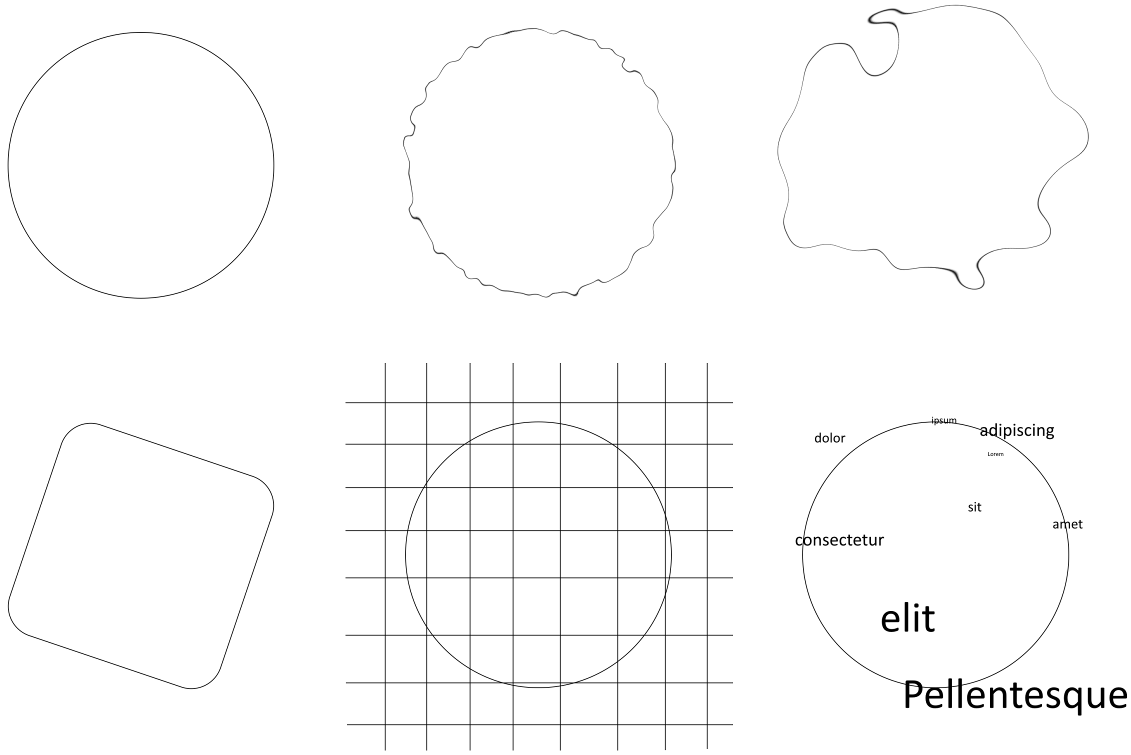

Figure 5.15: The figures presented here are used as further evaluation of the tool presented in this dissertation. Top left shows a perfect circle. Top middle shows a slight distortion effect applied to a circle. Top right shows a heavy distortion effect applied to a circle. Bottom left shows a rounded square. Bottom middle shows a circle which has a grid of lines imposed atop of it. Bottom Right shows a circle with random text placed around it.

Image	**S.A.D.**	**Hausdorff**	**cosine**	**template**
light distortion	27611	8.544	0.243	0.0339
heavy distortion	126376	51	0.298	0.219
square	113759	55.471	0.0877	0.193
lines	8601	1.414	0.00503	0.00179
text	17268	5	0.111	0.01599

Table 5.2: Results of the difference metrics from comparing each of the constructed image contours found in Figure 5.16 against the contour of the top right of that Figure.

5.3.2 Evaluation of Constructed Images' Comparison

The results of the comparison of the constructed images can be seen in Table 5.2. This Table shows the comparison of each image against the circle image in the top left of

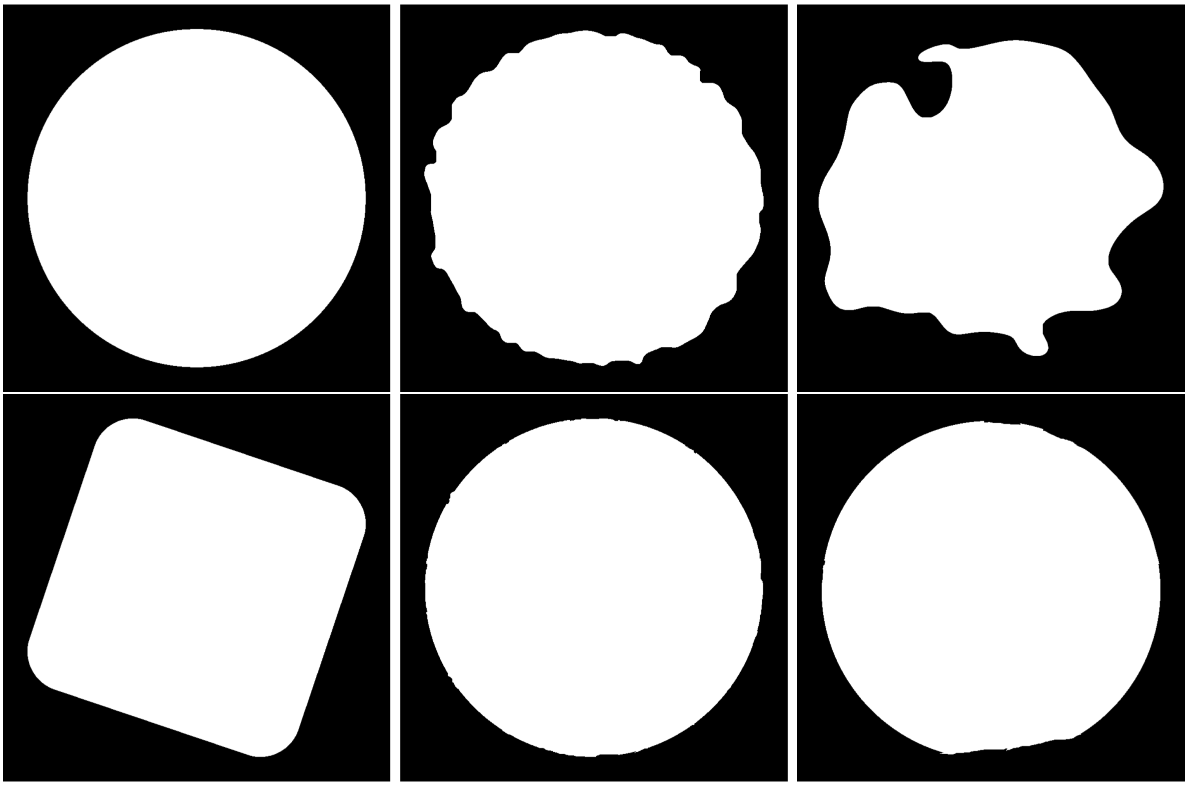

Figure 5.16: The figures here are the result of the vectorisation process as discussed in this dissertation on the images seen in Figure 5.15. Any broken lines present were manually connected by connecting the points that were broken.

Figure 5.15. The results for each metric approximate the expected result from each of the comparisons made. All four metrics find that the lines contour (bottom middle) is the most similar to the original circle image. S.A.D., Hausdorff, and template matching find that the next most similar is the text contour (bottom right), while cosine similarity finds the square contour (bottom left) to be the next most similar. Again, S.A.D., Hausdorff, and template matching find that the light distortion contour (top middle) is the next most similar contour to the circle contour, while cosine finds that the text contour (bottom right) is the next most similar. Lastly, S.A.D., Hausdorff, and template matching find the last two contours, heavy distortion and square, significantly less like the circle contour. Interestingly, cosine similarity finds light distortion and heavy distortion significantly less like the circle contour. This could be because of the concave nature of these contours in comparison to the other contours present, but more work is required to understand how cosine similarity works as a metric in this context.

Overall, the results of these constructed images support the evaluations on the maps of Ireland found in 5.1.1 and 5.1.2.

5.4 Summary

This section began with an extensive examination of the vectorisation process of historical maps as outlined in this project. Each of the steps in this process was presented with figures for each step and their usefulness and drawbacks were discussed. Where any issues were present in these steps, a set of ideas for possible future work to solve said issues were also presented. After the analysis of the vectorisation process, this section continued with an analysis of the comparison process. This began in 5.1.2, where an in depth look at the comparison of the contour constructed through the vectorisation process was compared against a modern day map. After this individual comparison, an exploration of the difference metrics used was done using a series of historical maps of Ireland from different sources and periods. With each metric discussed, any evaluations and findings of those metrics were discussed. Lastly, an additional experiment was conducted on a constructed set of images for the purposes of verifying the results found in the previous experiment on the maps of Ireland.

Chapter 6

Conclusions & Future Work

In conclusion, this project provides a set of methods and a tool which can be used to score the feature differences between historical maps. This project focused mostly on obtaining the coastline data of the primary object in the map, and using that to evaluate the difference in the coastline from another map. Extracting that coastline data required extensive image processing in order to isolate only the coastal border information from the map image. This project also provided an exploration of the use of four difference metrics; the sum of absolute differences, Hausdorff distance, cosine similarity, and template matching.

In relation to the difference metrics, this dissertation found that those metrics discussed had different levels of usefulness in their effectiveness at measuring and visualising the amount of error present in historical maps. As per the implementation in this project, the result of cosine similarity did not appear to accurately reflect the accuracy of the maps explored. On the other hand, Hausdorff distance, the sum of absolute differences, and template matching all followed the trend of the maps discussed increasing in accuracy over time.

While this project's vectorisation was limited to labeling the coastline of a land mass, there are many complexities in extracting that information that must be considered. In this dissertation, the removal of text and additional lines intersecting with the coast was discussed, along with one possible method of that removal. The process of removing text in this project was adequately successful. Although not all of the text was found, a significant amount was found so that the amount of manual effort required is significantly reduced. The method of removing lines from the image was less successful. It was not able to remove curved lines or lines at an angle from the borders of the image. It also interfered with the coastline, which was required to be repaired in subsequent image processing steps. Otsu's Binarization worked very well at producing a binary image from the original. However, for some images, they required to be inverted. This is often caused by having a background or sea colour which is darker than the content of the map itself (see appendix Figure 2).

The results of this project aids us in understanding more about our historical maps. It helps us to understand how mapping technology has changed over time and could influence the development of future technologies. While the scale used within this project assumes that inaccuracies in maps are due to error rather than coastline change, the methods discussed in this dissertation could be extended to consider coastline change on a smaller scale. This project can also help us to utilise historical documents which could contain information which is lost in the modern day. Understanding where error in historical maps occur can help us to correct them, and apply that corrected information in a modern context.

6.1 Future Work

6.1.1 Coastal Change

The change in a lands coastline over time due to geographical effects will also influence how maps are drawn. For this project, it is assumed that this change is negligible in comparison to the scale of the bodies of land and the time-frame used in these comparisons. However, in the future, it is worth exploring how much this would influence the results presented here. Conversely, it is also worth exploring whether this project has any implications on mapping coastline change over time. While outside of the scope of this project, the methods used here could be used to score the difference in coastal erosion and deposition over time in different historical maps. It is likely that this will need to be executed on a much smaller scale, where the coastline is much more clearly defined, unlike the scale of countries as used in this project.

6.1.2 Points of Interest

Extending this project to include points of interest could yield more interesting results than what was explored in this project. These points of interest could be manually or automatically selected, likely of important locations on the map such as capital or important cities, prominent geographic features, or historical landmarks. Since these features generally do not change over time, they could be used as another basis for comparison.

Contour Alignment

Instead of strictly using the selected points of interest for comparison, they could also be used in order to align the coastal contours before their comparison. As discussed in 4.2, before using the contour data in the comparison steps, they are aligned to each other. This ensures that any comparison made is invariant to the position the land was in the map image. In this project, the centroid of the contour boundary was used as this point of alignment. In the future, an exploration of the use of other points as the point of alignment should be conducted.

OCR

A more accurate OCR could aid in automatically finding these points of interest. As an example of how a more accurate OCR could be used, if two maps contain the same text, their positions within the map could be analysed. While this analysis could yield interesting results, as found in Chapter 5, while the OCR used in this project, EasyOCR, yielded appropriate results in finding text boundary data, it showed less accuracy in interpreting the text found within those boundaries. In future work, the use of other OCR software should be explored in their effectiveness in reading the difficult to decipher text found in historical maps.

Cosine Similarity

Using specific points of interest would likely provide better results for the cosine similarity difference metric. As found in Chapter 5, cosine similarity did not correlate with the other difference metrics explored in this project. This indicated that it does not suit the type of data given to it in this implementation. Using a series of points of interest rather than purely the coastline contour would likely yield a more appropriate result when using cosine similarity as a difference metric. Often, techniques such as using a Convolutional Neural Network (CNN) are used in order to automatically produce a feature vector for cosine similarity. While the interpretation of the values of that vector are mostly unknown due to the nature of CNN's, they can be used to find their similarity. The use of a CNN to obtain this feature vector should be explored in future work, and likely provide a better set of data to utilise cosine similarity.

6.1.3 Multiple Islands

In this project, only the main contour in the image was used as the basis for comparison. In future work, this project could be expanded to include the functionality of selecting multiple contour boundaries. This could be done by analysing the individual boundaries separately or together in one vector image. If this were done, Hausdorff distance required to be applied to only two contours. As such, for this difference metric, the individual boundaries are required to be analysed separately. The other difference metrics explored in this dissertation can be analysed together.

6.1.4 Visualisations of Change

This project could be extended in the future to provide further visualisations of the change in maps. This project focused more on the analytical side of measuring the change in historical maps. An extension providing a method to visualise the change more appropriately could be used as an educational tool. This could use the methods used within this project, and expand upon them to provide ways of visualising those results. This could be done using annotated images or with animation.

Bibliography

Brainerd, J. and Pang, A. (2001). Interactive map projections and distortion. *Computers Geosciences*, 27.

Brunelli, R. (2009). *Template matching techniques in computer vision: theory and practice.*

Chen, Y., Carlinet, E., Chazalon, J., Mallet, C., Dum´enieu, B., and Perret, J. (2021). Vectorization of historical maps using deep edge filtering and closed shape extraction. In *Document Analysis and Recognition - ICDAR 2021.*

ESRI (2024). Converting raster surfaces to features. `https://desktop.arcgis.com/en/arcmap/latest/extensions/3d-analyst/converting-raster-surfaces-to-features.htm`. Accessed: 2024-04-13.

GanchimegGanbold (2015). History document image background noise and removal methods. *International Journal of Knowledge Content Development Technology*, 5.

Giraldo Arteaga, M. (2013). Historical map polygon and feature extractor. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on MapInteraction.*

Huttenlocher, D., Klanderman, G., and Rucklidge, W. (1993). Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9).

Iosifescu, I., Tsorlini, A., and Hurni., L. (2016). Towards a comprehensive methodology for automatic vectorization of raster historical maps. *e-Perimetron.*

Jaided AI (2024). Easyocr website. `https://www.jaided.ai/easyocr/`. Accessed: 2024-04-14.

Latecki, L., Lakamper, R., and Eckhardt, T. (2000). Shape descriptors for non-rigid shapes with a single closed contour. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No.PR00662).*

Liu, C., Wu, J., Kohli, P., and Furukawa, Y. (2017). Raster-to-vector: Revisiting floorplan transformation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Liu, H.-C. and Srinath, M. (1990). Partial shape classification using contour matching in distance transformation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12.

Loran, C., Haegi, S., and 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*.

OpenCV (2024a). Opencv homepage. `https://opencv.org/`. Accessed: 2024-04-13.

OpenCV (2024b). Template matching. `https://docs.opencv.org/4.x/d4/dc6/tutorial_py_template_matching.html`. Accessed: 2024-04-14.

OSGeo (2024). Gdal documentation. `https://gdal.org/index.html`. Accessed: 2024-04-13.

Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9.

Pontius, R. and Lippitt, C. (2006). Can error explain map differences over time? *Cartography and Geographic Information Science*.

Potrace (2024). Potrace homepage. `https://potrace.sourceforge.net/`. Accessed: 2024-04-13.

Selinger, P. (2003). Potrace: a polygon-based tracing algorithm.

Serge, B. and Meyer, F. (1992). *Mathematical Morphology in Image Processing*, chapter The morphological approach to segmentation: the watershed transformation. CRC Press.

Smith, M. and Cromley, R. (2006). Coastal survey maps: From historical documents to digital databases. In *UCCGIA Papers and Proceedings, No. 1*.

Smith, M. and Cromley, R. (2012). Measuring historical coastal change using gis and the change polygon approach. *Transactions in GIS*.

Stanford Libraries (2024). David rumsey historical map collection. https://www.davidrumsey.com/. Accessed: 2024-04-14.

Wenyin, L. and Dori, D. (1999). From raster to vectors: Extracting visual information from line drawings. *Pattern Analysis Applications 2*.

Woods Hole Coastal and Marine Science Center (2018). Dsas generated transect. https://www.usgs.gov/media/images/dsas-generated-transect. Accessed: 2024-04-14.

Xu, C., Liu, J., and Tang, X. (2009). 2d shape matching by contour flexibility. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31.

Appendix

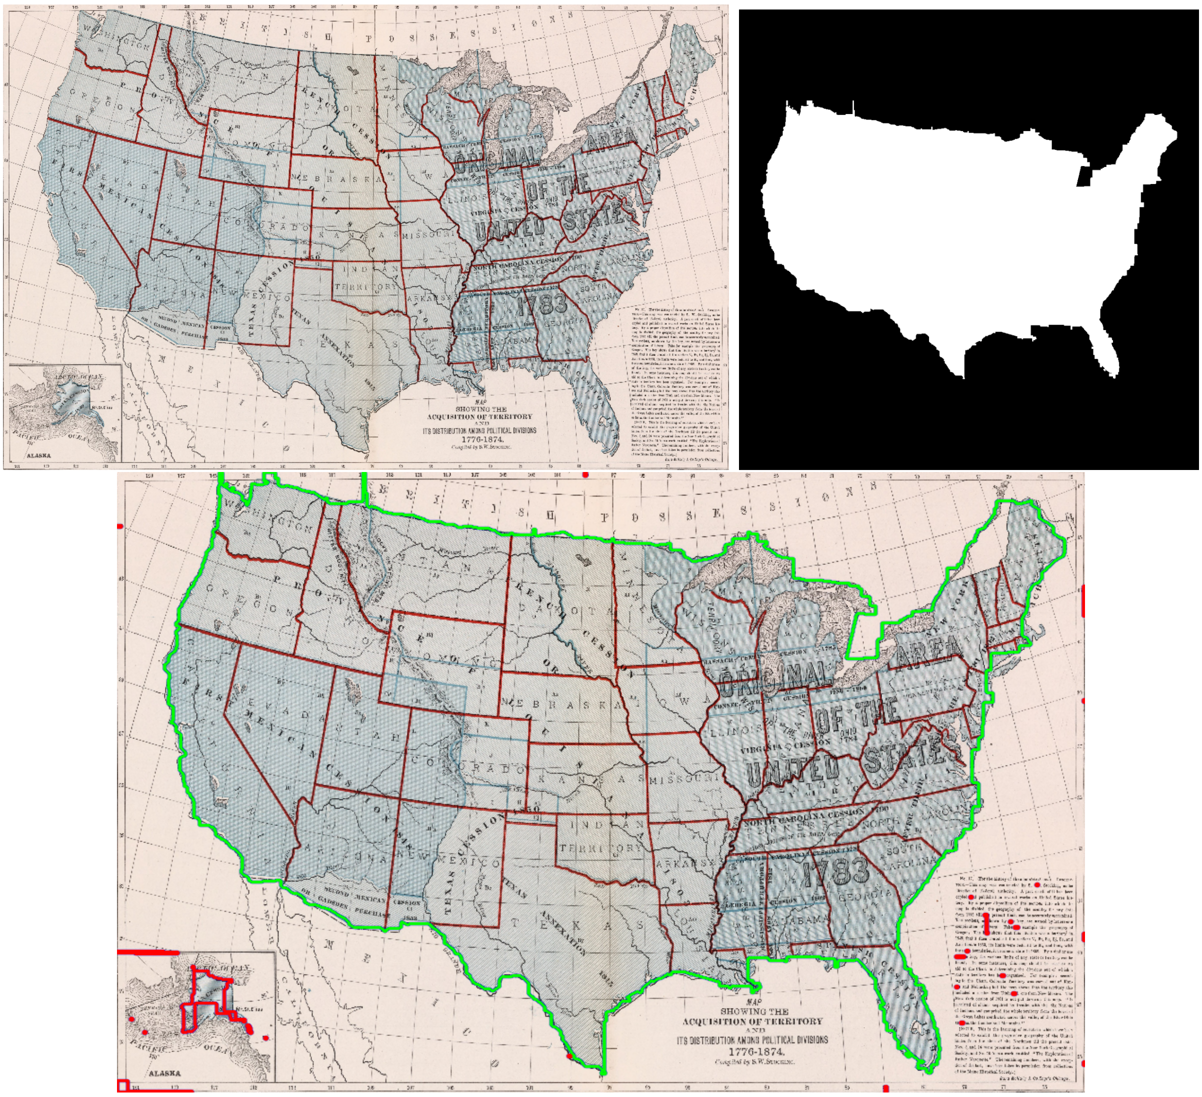

Figure 1: These three figures show the vectorisation process as outlined in this dissertation applied to a map of the USA. Some minor manual corrections were made on the coastline at California, Texas, and Dakota. This map was produced by Walter R. Houghton in 1880 (Stanford Libraries, 2024).

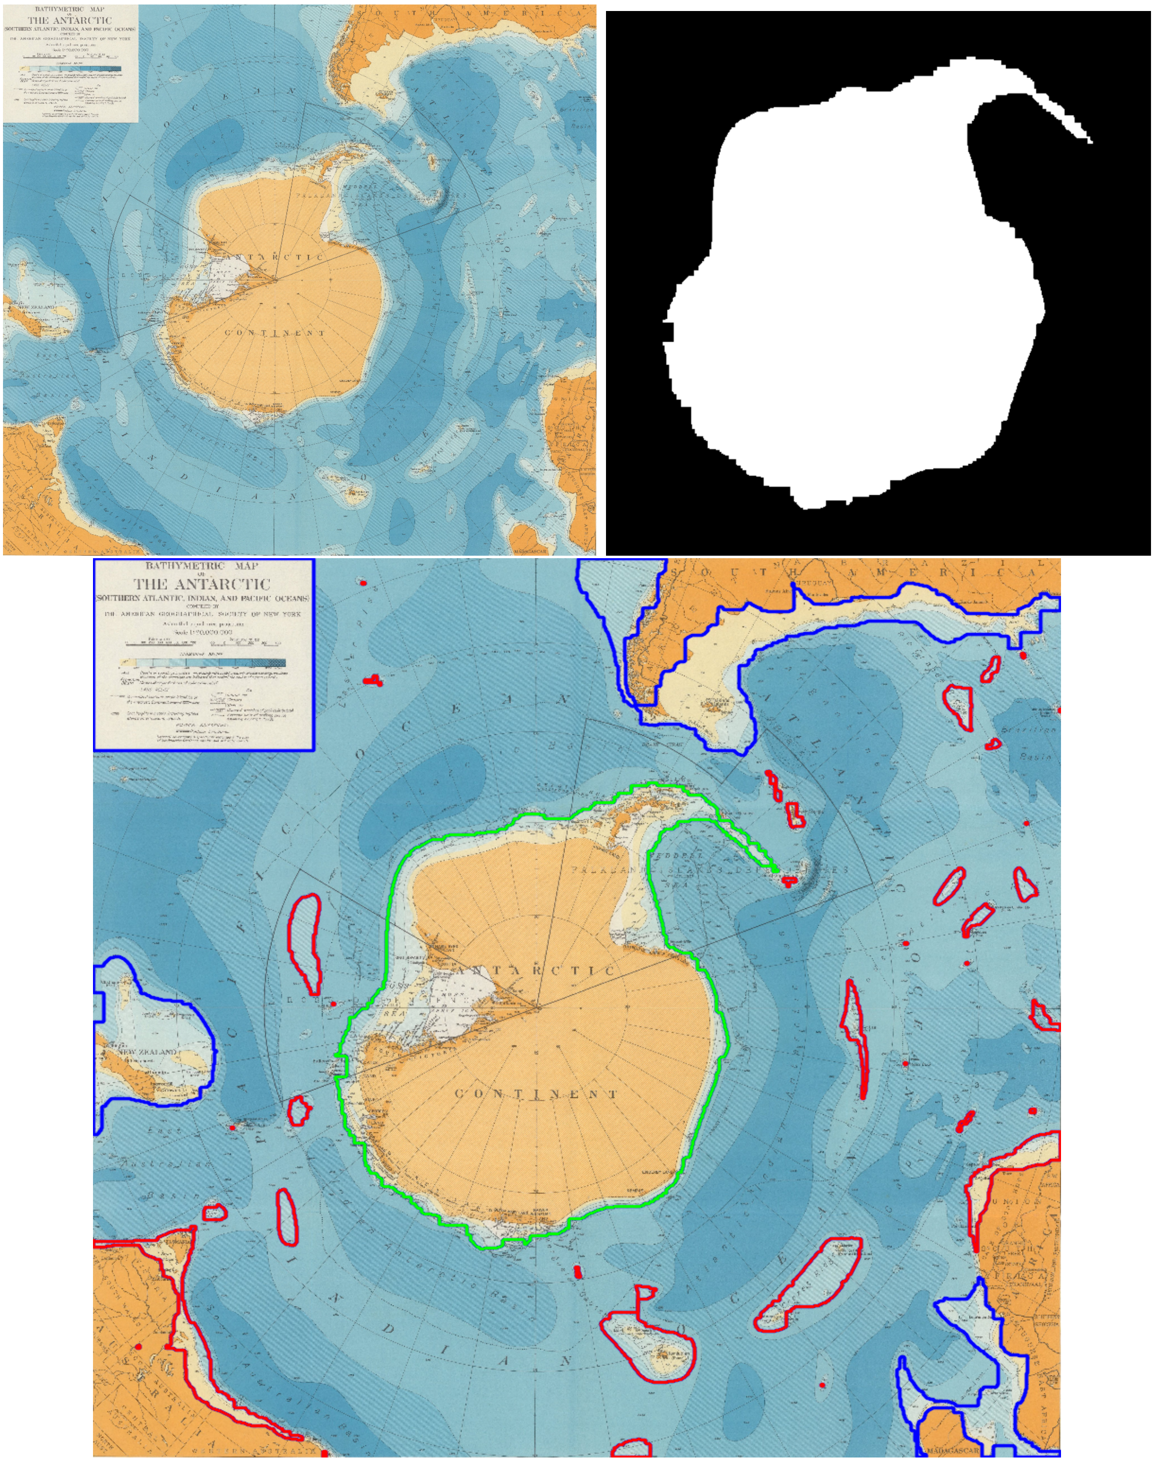

Figure 2: These three figures show the vectorisation process as outlined in this dissertation applied to a map of Antarctica. The result of threshold required to be inverted, since the grey-scale value of the water is darker than that of land. No manual corrections were made other than selecting the boundaries of the image. This map was produced by W.L.G. Joerg and the American Geographical Society in 1930 (Stanford Libraries, 2024).

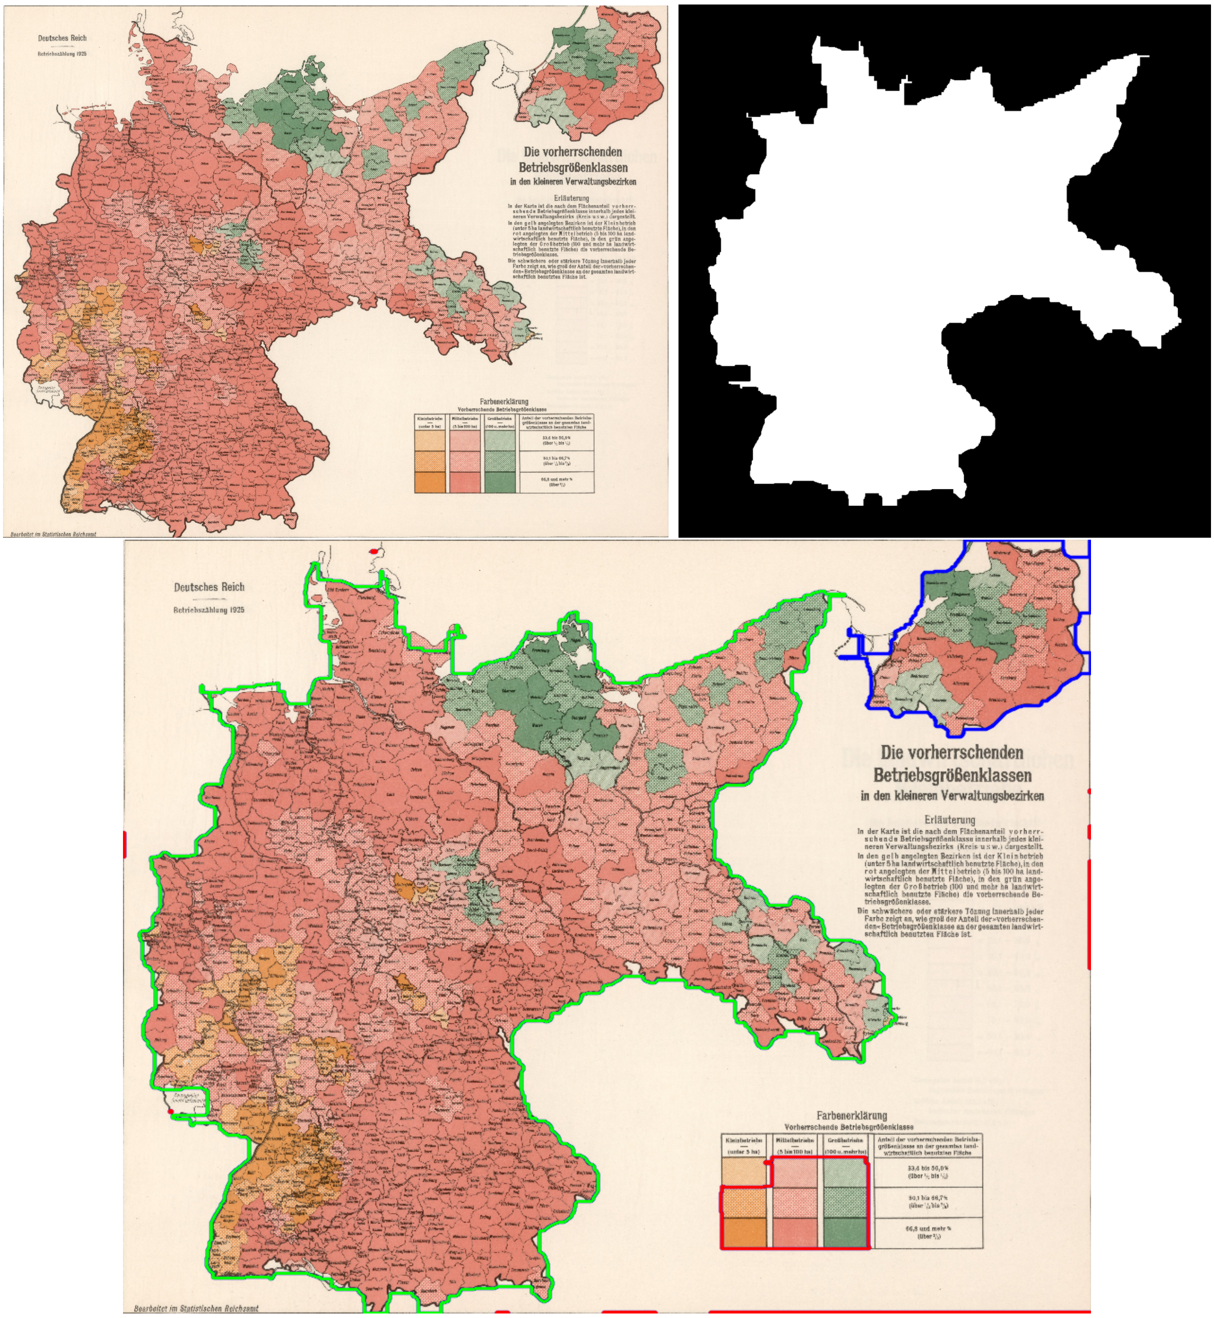

Figure 3: These three figures show the vectorisation process as outlined in this dissertation applied to a map of Germany. No manual corrections were made other than selecting the boundaries of the image. Notably all of the text on a solid colour background was found using EasyOCR in this map, but any text on a dotted background was not found. This map was produced by Statistischen Reichsamt (Reich Statistical Office) in 1925 (Stanford Libraries, 2024).

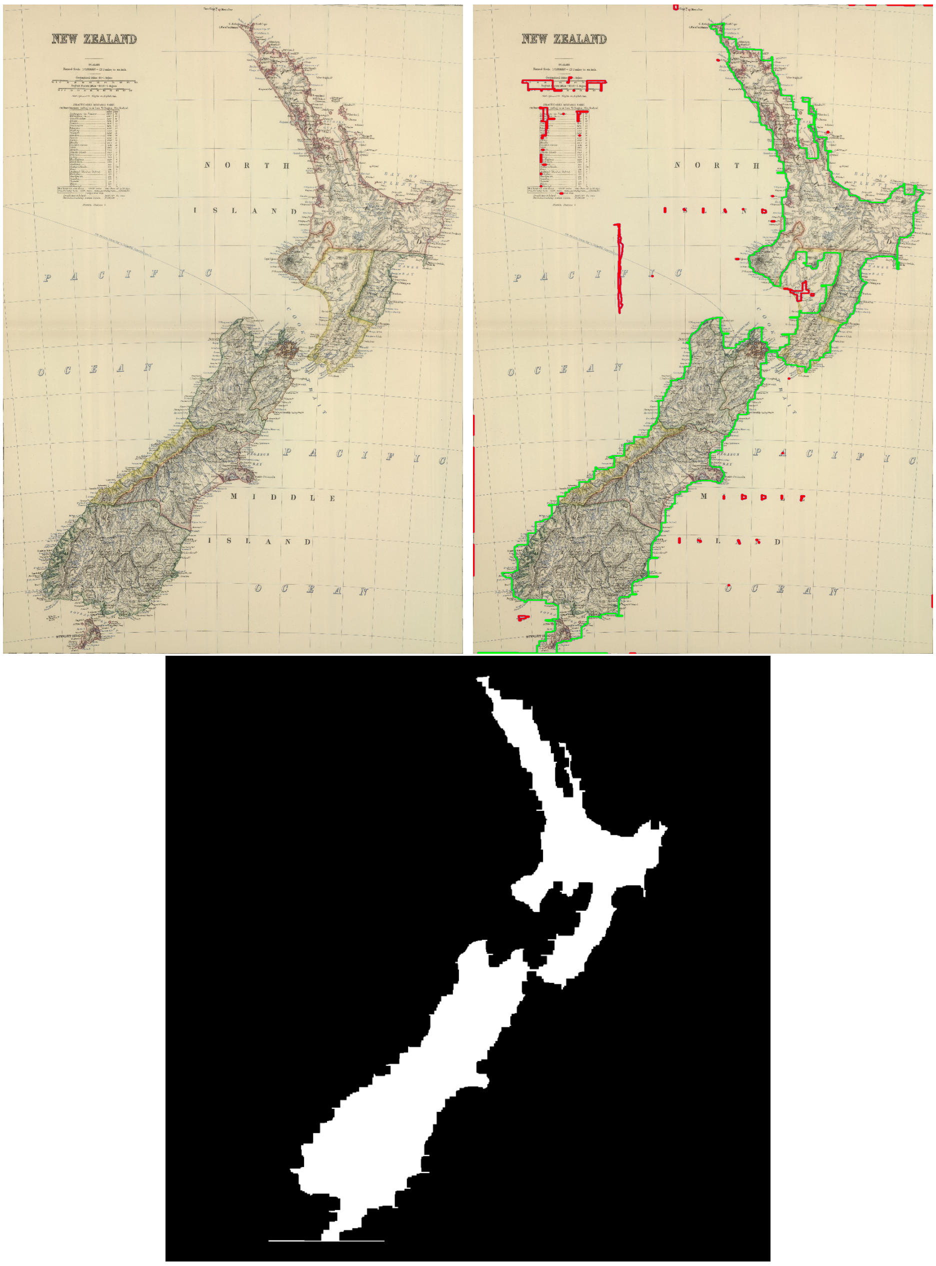

Figure 4: These three figures show the vectorisation process as outlined in this dissertation applied to a map of New Zealand. No manual corrections were made other than selecting the boundaries of the image. Not all of the text was found on the coastline of this map. This is more prominent at the south of the map, where there is more spurious protruding areas of the contour. The very bottom of the contour also connects with the map border, resulting in a horizontal line appearing in the contour image. The two separate islands were connected automatically, but unfortunately, this may not be the case for every map of New Zealand used with this tool. A section of the contour of the north island was also not detected appropriately, as can be seen in the inner red section being disconnected from the rest of the island. This map was produced by Alexander Keith Johnston in 1879 (Stanford Libraries, 2024).