

Mining Spatio-temporal Data on Industrialization from Historical Registries

D. Berenbaum¹, D. Deighan¹, T. Marlow², A. Lee¹, S. Frickel², and M. Howison^{1*}

¹*Data Science Practice, Computing & Information Services, Brown University, 3 Davol Square, Providence, RI 02912, USA*

²*Institute at Brown for Environment and Society, Brown University, 80 Waterman Street, Providence, RI 02912, USA*

Received 07 December 2016; revised 28 July 2017; accepted 02 September 2017; published online 20 November 2017

ABSTRACT. Despite the growing availability of big data in many fields, historical data on socio-environmental phenomena are often not available due to a lack of automated and scalable approaches for collecting, digitizing, and assembling them. We have developed a data-mining method for extracting tabulated, geocoded data from printed directories. While scanning and optical character recognition (OCR) can digitize printed text, these methods alone do not capture the structure of the underlying data. Our pipeline integrates both page layout analysis and OCR to extract tabular, geocoded data from structured text. We demonstrate the utility of this method by applying it to scanned manufacturing registries from Rhode Island that record 41 years of industrial land use. The resulting spatio-temporal data can be used for socio-environmental analyses of industrialization at a resolution that was not previously possible. In particular, we find strong evidence for the dispersion of manufacturing from the urban core of Providence, the state's capital, along the Interstate 95 corridor to the north and south.

Keywords: structured text, historical data, geocoding, page layout analysis, socio-environmental analysis

1. Introduction

In most states in the U.S., detailed registries of manufacturers are compiled annually, dating back to the 1950s. These printed historical registries are a rich source of data on the location, size, and type of industrial activity over time. However, mining those data is not straight-forward using existing OCR tools, because of the registries' structure. To address this, we developed *georeg*, a pipeline for extracting addresses and other business information from historical registries, and tested it with images we scanned from Rhode Island manufacturing registries spanning the 1950s through the 1990s. In these scanned images, *georeg* identifies each heading and the ordering of manufacturer listings, so that we can extract the name, address, business type, and number of employees as tabular data, which are then geocoded to provide latitude and longitude.

With these spatio-temporal data, we are able to perform more detailed socio-environmental analyses of changes in industrialization and locations of potentially hazardous manufacturing sites than was previously possible. Among other findings, the data show a dispersion of industrial sites from Providence, the state's capital and most populous city, to adjoining cities along the Interstate 95 corridor, which runs the length of the state.

While many robust methods exist for digitizing historical documents, there has been less focus on modeling the complex structure in printed texts like directories, which use combinations of headings, font styles, and column and block layouts to organize and present repeated listings of structured data. The two primary components of digitizing printed documents are page layout analysis and optical character recognition (OCR). Page layout analysis provides information about the position and arrangement of blocks of text within a page. OCR converts pixelated images into characters and words.

While OCR makes it possible to extract all text from scanned images of directories, it fails to automatically detect complex structure, even when using more advanced OCR methods that incorporate page layout analysis, such as those in the Tesseract package (Smith, 2009). Tesseract includes an OCR mode that detects tab-stops to extract text that is consistent with the document's column layout. However, this approach does not segment individual text blocks within the columns.

Although there is an extensive literature on page layout analysis (for a review, see Tang et al. (1996) or Mao et al. (2003)), existing approaches to segmenting documents are largely based on physical features such as the relative position and size of lines or blocks of text. In general, physical segmentation does not capture the hierarchy or logical structure of a complex document. The existing approach that is closest to ours is a method created by Ma and Doermann (2003) that segments documents with repeated entries of similar structure, such as phone books and dictionaries. In this method, both physical features (e.g., line structure) and logical features (e.g., word patterns) are used to train a segmentation model, starting

* Corresponding author. Tel.: +1 401-863-6743; fax: +1 401-863-7216.

E-mail address: mhowison@brown.edu (M. Howison).

from sample documents which have already been correctly segmented into repeated entries. However, the segmentation provided by this method is limited in that it does not distinguish between different kinds of entries (e.g. headings vs. listings) and does not capture the structure of data elements within an entry (e.g. name vs. address vs. phone number in a phone book entry).

2. Methods

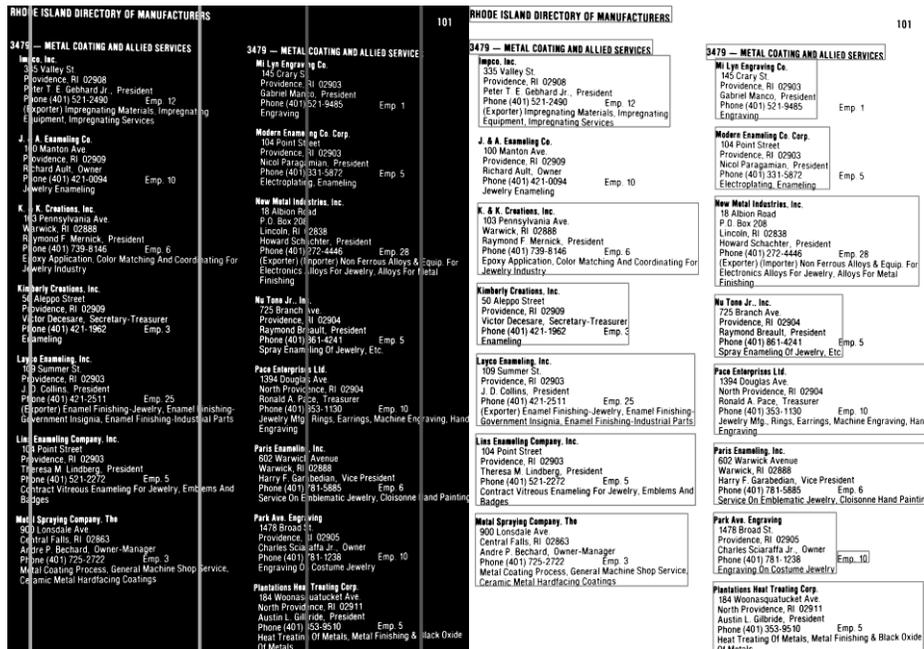
2.1. Scanning

We scanned printed registries through a combination of book scanners and photography. Manufacturing registries for 16 individual years between 1953 and 1994 were accessed at the Brown University Library and the Rhode Island Historical



(a) Dilation/Erosion

(b) Center Clustering



(c) Column Positions

(d) Headings/Records for OCR

Figure 1. A contour merging and column clustering approach to identifying headings and records in the registry pages.

Society. We used the open-source software ScanTailor (<http://scantailor.org>) to trim edges, straighten, de-warp, and convert the images to grayscale.

2.2. Contour Merging

We developed a contour merging algorithm to partition a page image into the areas that corresponds to an individual listing of a manufacturer. Contours are the outermost boundaries of a set of contiguous pixels of the same intensity (Suzuki and Abe, 1985) and we initially identify them using OpenCV (Bradski, 2000). However, these contours frequently outline only the individual characters in the text. We use thresholding, erosion, and dilation to merge nearby contours into contiguous areas of the image that cover all characters in a manufacturing listing (Figure 1a).

After thresholding and inverting the grayscale image to white text on a black background, we repeatedly perform a close operation, which is an erosion followed by a dilation. The close operations eliminate small gaps of black-space between characters and lines of text. Then, we perform fewer iterations of an open operation, which is a dilation followed by an erosion. The open operations de-noise the contours and smooth them, removing small artifacts of white pixels. For the erosion and dilation operation’s kernel shape, we use a rectangle with dimensions that the user can configure.

$$I \ominus K = \{p \mid K_p \subseteq I\} \tag{1}$$

where K_p is the kernel for pixel p . That is, each pixel p remains

white only if every other pixel in its kernel is also white. Similarly, binary dilation is defined as:

$$I \oplus K = \bigcup K_p \tag{2}$$

In some cases, there is no additional space between the lines of text for each manufacturer listing. In these registry formats, a new listing is denoted by indenting the first line. When we merge the contours of these pages, a single merged contour may encompass several listings or an entire column of text. In this case, we use the indentation to identify individual listings and split the contour at each indentation. We identify each indent by following the contour and noting where the x-coordinate stretches beyond a given threshold distance from the left bound of the contour.

2.3. Identifying Manufacturer Listings

After merging the contours, some correspond to manufacturer listings, but others are image artifacts or correspond to extraneous information, like page numbers. The manufacturer listings are bounded within the columns of the page, so we only consider a contour as a candidate for a manufacturer listing if it is aligned with a page column. We use k-means clustering (Lloyd, 1982) to group the contours by their left and right bounds, setting the number of clusters equal to the number of columns in the page (Figure 1b). Column-aligned listings have left and right bounds close to the column boundaries and generally outnumber the extraneous contours, so the clusters converge around the left and right bounds of the columns (Figure 1c). To



Figure 2. The number of manufacturers identified and geocoded in each year’s registry, compared to the estimated number of manufacturers.

filter out extraneous contours, we eliminate those which are greater than a given number of standard deviations away from the cluster centroids.

2.4. Parsing Information from Listings

For each potential manufacturer listing, we perform OCR with Tesseract (Smith, 2007) on the sub-image defined by the bounding box of the contour (Figure 1d). We use regular expressions to extract the business name, type, address, number of employees, and other standardized information included for each listing. If the text is more consistent with a heading (for example, it contains only a city name), we retain that information and include it with each subsequent manufacturer listing under the heading.

Some headings are centered across the page, in which case they are not bounded within the page columns and are eliminated from potential manufacturer listings. If such headings exist and contain meaningful information, we OCR all headings centered across the page and note their vertical positions. For each manufacturer listing, we include the information from the closest heading above it (or the last heading on a previous page if there is no heading above it).

2.5. Geocoding

For each manufacturer listing in which an address was identified, we geocode the address to obtain a latitude, longitude, and confidence score. To improve geocoding success, we match the cities to a list of correctly spelled cities within the state. If the match ratio (based on Levenshtein distance) for the best city match is above a given threshold, we replace the city name found by OCR with the matched city name. We geocode each address using a local ArcGIS server at Brown University, but any geocoding service supported by the geopy package could be substituted.

3. Results

3.1. Accuracy

To evaluate the accuracy of georeg, we compared the number of identified manufacturer listings with an estimate of the actual content of the registries. For the estimate, we manually counted the number of listings on ten pages of each year's directory and then multiplied the average number of listings per page by the total page count. Overall, georeg identified 99% of the estimated number of listings across all pages tested (Figure 2). Individual years differ in how closely they approximate the estimated number of listings, but each year's results are within 15% of the estimate.

We also measured the number of listings in which the address was geocoded with a confidence score of at least 75%. georeg successfully geocoded 61% of the estimated number of listings, but the results differed dramatically by year. In 12 of the 16 years, between 62% and 75% of the estimated number of listings were successfully geocoded. However, in 1971, only 26% were successfully geocoded, and in each of the worst four

years fewer than half of the estimated number of listings were successfully geocoded.

Geocoding may fail for a number of reasons:

1. The listing does not contain an address. The listing may be missing an address if a heading or page number, for example, is misidentified as a manufacturer listing. Refining the contour classification method could help reduce these errors, such as by using the shape of the contour or the contents of the text to determine whether it is a manufacturer listing.
2. The characters of the listing are misread during OCR. In the case of OCR errors, we could use additional features in the Tesseract OCR software to increase the accuracy of the characters. In particular, we could train Tesseract for the different fonts which appear in the registries, and we could blacklist characters that never appear in the registries.
3. The address is not successfully extracted by regular expressions. If the regular expressions method does not find the address within the text, we could use known approaches for named entity recognition, in which parts of text are classified into pre-defined categories.

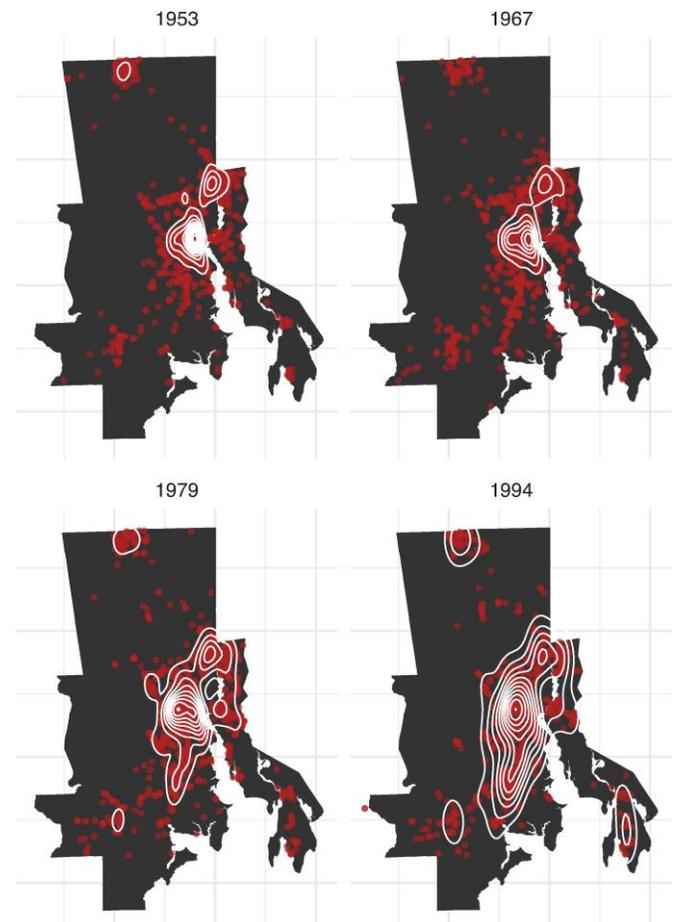


Figure 3. Density contour plots for all active manufacturing facilities.

4. The address is not found in the geocoding database. In this case, we could aggregate results from multiple geocoders to improve coverage.
5. The address is invalid or its location is impossible to determine. The address may be impossible to geocode because it is a post office box or the street no longer exists, for example.

However, despite these challenges, our initial dataset does have sufficient coverage to demonstrate its value for socio-environmental studies.

3.2. Application to Socio-environmental Studies

This new approach to historical data collection has already made substantial contributions to the sociological literature on urbanization and environmental inequality (Elliott and Frickel, 2013; Liévanos, 2015; Mohai and Saha, 2015); and has direct relevance for industrial economics and geography as well (Duranton and Overman, 2005; Martin and Sunley, 2006; Frenken et al., 2015). For example, in their initial series of studies, Frickel and Elliott ((2008); (Elliott and Frickel, 2011; Elliott and Frickel, 2013; Elliott and Frickel, 2015)) also relied on

state manufacturing registries, but practical concerns limited data collection to the seven most polluting industrial sectors in the urban core of their cities of interest. Combining this narrower subset of industrial site data with tract-level data from the U.S. Census, as well as geocoded hazardous site lists from state environmental regulatory agencies, made it possible to conduct longitudinal analysis of demographic and regulatory changes and how those dynamics influence the spatial and temporal accumulation of environmental industrial risk.

By extending this previous approach to include all industrial sites in an entire state over a 41 year period, we obtained several new results:

1. By including annual data for all manufacturing activities in Rhode Island, we are able to study not only deindustrialization in the urban core of Providence, but also the consequent suburbanization of manufacturing as industrial activities gradually shifted to adjacent cities to the north and south.
2. The additional spatial scope allows us to generate stronger evidence that the spatial redistribution of manufacturing has been impacted by changes in the transportation infras-

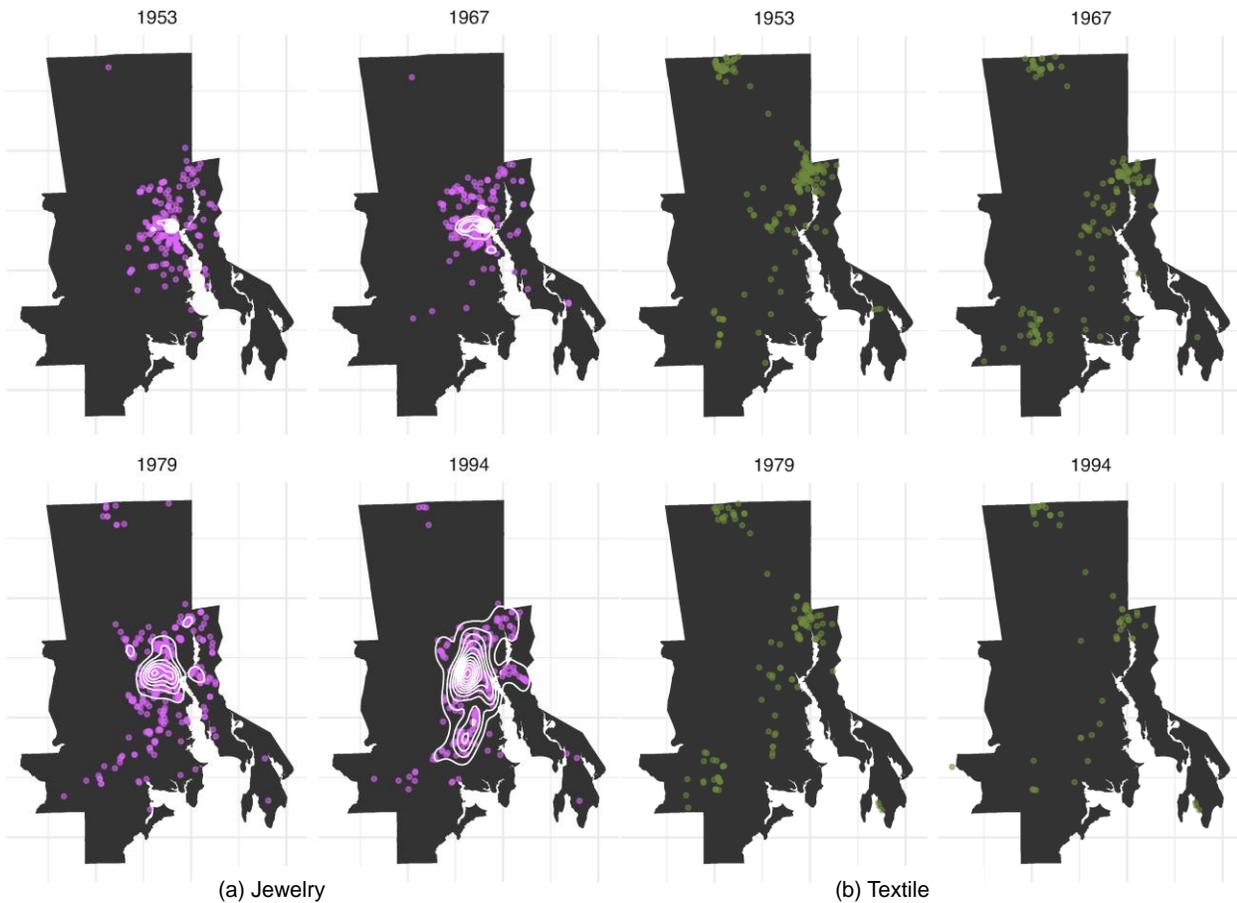


Figure 4. Density contour plots for active manufacturing facilities in the jewelry and textile industrie.

structure of the state. In 1953, manufacturing concentrated in the urban core of downtown Providence near the large rail terminals and shipping port; by 1979, manufacturing had begun to relocate southward following the newly constructed Interstate 95 corridor (Figure 3).

- By mining data for all manufacturing sites, the new approach has allowed greater flexibility in the exploration of sector-specific clustering and distributional changes over time. For example, preliminary analysis of the spatial and temporal distribution of jewelry and textile manufacturing (Figure 4) - two historically important industries in Rhode Island - shows that jewelry manufacturing exhibits a distinct pattern of clustering while textiles have a more uniform and widespread distribution.

These preliminary findings illustrate the advantages of our comprehensive approach to data mining, which allows for richer empirical analysis of socio-environmental change and greater opportunities for theory development than was possible in prior research.

4. Conclusion

We have demonstrated a successful approach for mining and geocoding structured text in historical registries, and shown its utility for analyzing industrialization in the US state of Rhode Island. Previous studies have shown that structured text can be reliably mined in an automated way from printed directories (Ma and Doermann, 2003). However, our pipeline is the first, to our knowledge, to demonstrate the comprehensive integration of automated page-layout analysis with optical-character recognition, content analysis, and geocoding, to produce tabulated data ready for socio-environmental analysis.

Previous studies have also shown that remote-sensing data can be used to study the dynamics of urbanization (Li et al., 2015). A key difference between our approach and a remote-sensing approach is that we are able to extract standardized information, including employee counts and industry sector codes. In particular, sector codes allow us to study the dynamics within a particular sector, such as jewelry or textile manufacturing, which are historically important industries in Rhode Island.

The approach we have presented is also scalable. Although we have shown a specific application to manufacturing registries from Rhode Island, the same principles apply to other US state's registries and to other printed historical documents with similar structure, such as phone books, almanacs, and city directories. We are currently adapting georeg to these additional data sources. In particular, we believe phone books will be an ideal supplement, adding commercial, non-profit (e.g. schools, hospitals, churches) and residential data to provide a comprehensive history of land use in urban areas.

Availability: Source code for georeg is freely available for non-commercial use at <https://github.com/brown-data-science/>

georeg. All results presented in this paper were obtained using the version of georeg at commit hash cacda11.

Acknowledgement. This research was conducted using computational resources at the Center for Computation and Visualization, Brown University.

References

- Bradski, G. (2000). The OpenCV Library. *Dr. Dobbs's J.*, 25(11), 384-386.
- Durant, G., and Overman, H.G. (2005). Testing for localization using micro-geographic data. *Rev. Econ. Stud.*, 72(4), 1077-1106. <https://doi.org/10.1111/0034-6527.00362>
- Elliott, J.R., and Frickel, S. (2011). Environmental dimensions of urban change: Uncovering relict industrial waste sites and subsequent land use conversions in Portland and New Orleans. *J. Urban Aff.*, 33(1), 61-82. <https://doi.org/10.1111/j.14679906.2010.00533.x>
- Elliott, J.R., and Frickel, S. (2013). The historical nature of cities A study of urbanization and hazardous waste accumulation. *Am. Sociological Rev.*, 78(4), 521-543. <https://doi.org/10.1177/0003122413493285>
- Elliott, J.R., and Frickel, S. (2015). Urbanization as socioenvironmental succession: The case of hazardous Industrial Site Accumulation. *Am. J. Sociol.*, 120(6), 1736-1777. <https://doi.org/10.1086/681715>
- Frenken, K., Cefis, E., and Stam, E. (2015). Industrial dynamics and clusters: A survey. *Reg. Stud.*, 49(1), 10-27. <https://doi.org/10.1080/00343404.2014.904505>
- Frickel, S., and Elliott, J.R. (2008). Tracking industrial land use conversions: A new approach for studying relict waste and urban development. *Organ. Environ.*, 21(2), 128-147. <https://doi.org/10.1177/1086026608317799>
- Li, X.W., Liang, C., Shi, J.B., and Li, M.D. (2015). Spatiotemporal dynamics and urban land-use transformation in the rapid urbanization of the Shanghai metropolitan area in the 1980s-2000s. *J. Environ. Inf.*, 20(2), 103-114.
- Liévanos, R.S. (2015). Race, deprivation, and immigrant isolation: The spatial demography of air-toxic clusters in the continental United States. *Soc. Sci. Res.*, 54, 50-67. <https://doi.org/10.1016/j.ssresearch.2015.06.014>
- Lloyd, S.P. (1982). Least squares quantization in PCM. *IEEE Trans. Inform. Theory*, 28(2), 129-137. <https://doi.org/10.1109/TIT.1982.1056489>
- Ma, H., and Doermann, D.S. (2003). Bootstrapping structured page segmentation. *Proc. SPIE 5010*, 179-188. <https://doi.org/10.1117/1.2.476058>
- Mao, S., Rosenfeld, A., and Kanungo, T. (2003). Document structure analysis algorithms: a literature survey. *Proc. SPIE 5010*, 197-207. <https://doi.org/10.1117/12.476326>
- Martin, R. and Sunley, P. (2006). Path dependence and regional economic evolution. *J. Econ. Geogr.*, 6(4), 395-437. <https://doi.org/10.1093/jeg/lbl012>
- Mohai, P., and Saha, R. (2015). Which came first, people or pollution? Assessing the disparate siting and post-siting demographic change hypotheses of environmental injustice. *Environ. Res. Lett.*, 10(11), 115008. <https://doi.org/10.1088/1748-9326/10/11/115008>
- Smith, R.W. (2007). An overview of the tesseract OCR engine. *Proc. of the 12th International Conference on Document Analysis and Recognition*, 629-33. Los Alamitos, CA, USA. <https://doi.org/10.1109/ICDAR.2007.56>
- Smith, R.W. (2009). Hybrid page layout analysis via tab-stop detection.

- Proc. of the 10th International Conference on Document Analysis and Recognition*, 241-245. Barcelona, Catalonia, Spain. <https://doi.org/10.1109/ICDAR.2009.257>
- Suzuki, S., and Abe, K. (1985). Topological structural analysis of digitized binary images by border following. *Comput. Vision, Graphics, Image Process.*, 30(1), 32-46. [https://doi.org/10.1016/0734-89X\(85\)90016-7](https://doi.org/10.1016/0734-89X(85)90016-7)
- Tang, Y.Y., Lee, S.W., and Suen, C.Y. (1996). Automatic document processing: A survey. *Pattern Recogn.*, 29(12), 1931-1952. [https://doi.org/10.1016/S0031-3203\(96\)00044-1](https://doi.org/10.1016/S0031-3203(96)00044-1)