

Automated chain line marking and matching in radiographs of Rembrandt's prints

X. Xi*, D. Conathan[†], A. House[‡], W. A. Sethares[§], C. R. Johnson, Jr.[¶]

April 27, 2016

1 Introduction

The sieve-like molds used early paper manufacturing leave characteristic chain line patterns that can help determine if two pieces of paper were made using the same mold [1]. Such *molddates* can help in establishing chronology, suggest paper preferences, and indicate periods of intense activity of an artist, and the study of Rembrandt's prints has occupied a prominent place within this scholarship [2]. There are two basic steps to the automation of the analysis. First, the chain lines must be accurately marked, extracting points and lines from the x-ray images. Second, the markings from many x-rays must be compared to find matches. This paper presents improved algorithms for both of these tasks. Section 2 follows [5] in replacing the classic Radon transform approach of [3] (which requires that the chain lines be straight) with a local method that outputs marked grids of points. The method segments the images, corrects for tilt, filters, conducts a (local) Radon transform and then concludes with a peak detection. The marked grids from a large collection of x-rays of Rembrandt's prints is then used in Section 3 as input

*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY

[†]Department of Electrical and Computer Engineering, University of Wisconsin – Madison, conathan@wisc.edu

[‡]School of Electrical and Computer Engineering, Cornell University, Ithaca, NY

[§]Department of Electrical and Computer Engineering, University of Wisconsin – Madison, WI and the Rijksmuseum, Amsterdam, sethares@wisc.edu

[¶]School of Electrical and Computer Engineering, Cornell University, Ithaca, NY and the Rijksmuseum, Amsterdam, johnson@ece.cornell.edu

to the matching algorithm. Previous approaches were based on line-pairs, which may be problematic due to individual missed lines and to error propagation. The new approach uses a parameterization based on the Hesse normal form which transforms the segmented lines into points. A modified Procrustean algorithm is then applied to optimally match the complete set of points, bypassing issues with the pairwise line comparisons. The performance of the system is evaluated using a dataset of 462 radiographs of Rembrandt's prints [4]. Results and new moldmate matches are presented in Sec. 4 and compares the chain line approach with the more traditional watermark-based approach.

2 Extraction of Chain Lines

Radiographs can be overly light or dark (see Fig. 1(a)), of low contrast, and noisy, making the chain lines difficult to extract in a fully automated way. Accordingly, we adopt a strategy that minimizes human intervention, leaving the human to make only a handful of judgements at the start and at the end of a mostly automated procedure. The full procedure is described in several steps, and shown in outline in Fig. 1.

2.1 Image Preparation

Each radiograph in the Rembrandt data set contains an identification number, the Bartsch number, a ruler, and the image of the paper. We remove any image with fewer than five chain lines. The resolutions of the images may differ, and so it is necessary to find the dpi (dots per inch) for normalization. When the dpi is not specified, it is possible to manually mark two end points on ruler; the dpi is the ratio of the number of pixels along the line to the horizontal distance between the endpoints.

2.2 Image Tilt and Equalization

Rotate the image so that the leftmost chain line is vertical, and apply histogram equalization to heighten the contrast. Reverse the color of the image (if needed) to ensure that the images have bright chain lines on a dark background, as shown in Fig. 1. Part (a) shows the original dim and monochromatic radiograph. The images in (b) (after equalization) are clearer.

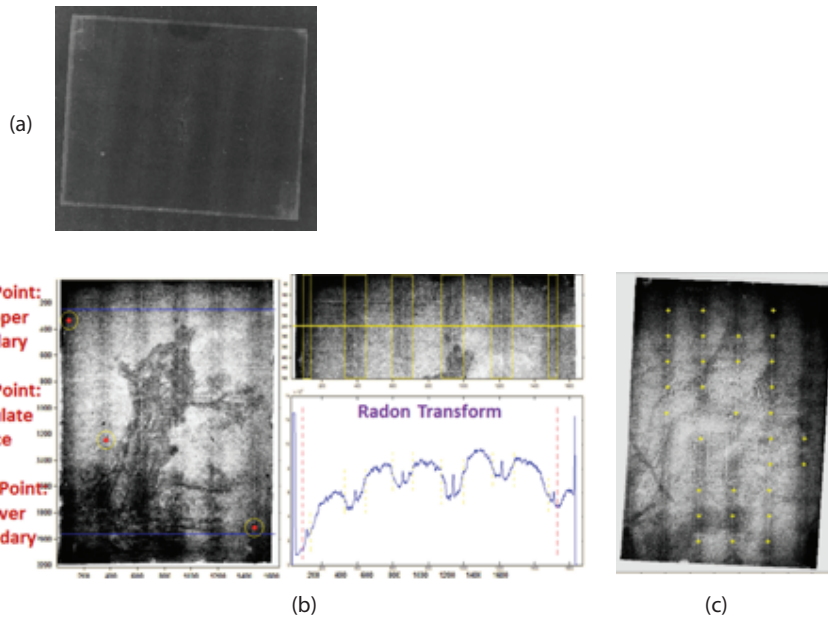


Figure 1: The extraction of the chain lines begins with a manual step of image cropping and rotation to locate the relevant portions of the image as in (a). The operator then chooses three points (as in (b)) which are used by the automated methods (as described in the text and drawn from [5]) to arrive at a collection of keypoints. These are shown in yellow, aligned along each chain line, in (c). The final step requests the user to verify that the derived points are sensible.

2.3 Manual Marking (three points)

Mark three points on the chain lines, as shown in Fig. 1(b). The first point is at the top of the leftmost chain line. The second point is at the bottom of the rightmost chain line. These two points outline the bounding box formed by the chain lines. Mark the third point on the second chain line from the left. This is used to calculate an approximation to the average chain spacing in the image, which can be used to help focus the automated procedure on the correct regions.

2.4 Automated Marking

The program divides the image into a number of horizontal strips (with default ten). The Radon transform is applied in each strip; an example is shown in Fig. 1(b). The highest peaks correspond to the white straight lines at the edge; these are removed. The repeated wide peaks are caused by the broad bright vertical parts in the left image; it is the sharp local peaks that show the locations of the chain lines.

Since the average chain spacing is known (from the manual marks), and since the standard deviation of the chain spacing in any image is quite small (compared to the distribution of average spacing over all images in the dataset [1]), small “regions of interest” are set across the horizontal strip. These contain the bright chain lines. Observe that these local peaks have some characteristic features. The local peaks of interest are narrow and sharp, and have higher magnitude than other needle-like peaks in the image. The local peaks of interest also have relatively similar maximum values when compared with other noisy needle-like peaks. Thus the program selects local peaks which have these characteristics. A threshold H is derived from the magnitude of the local peak on the leftmost chain line in the first image. This is compared to the height of the manually marked point with the height of the points 10 pixels to the left and 10 pixels to the right; the larger height difference is notated h and $H = 0.55h$ is used for the thresholding.

2.5 Verification

This semi-automated chain line detection and marking method works well on more than 95% of the images with over 5 chain lines. An example of a marked up result is shown in Fig. 1(c). Typical failures occur when the method selects a portion of the watermark (instead of the chain line). The final stage of human

intervention is to visually examine the selected points (as in Fig. 1(c)) and to remove those which are obviously in error.

The performance of the marking program is evaluated using the dataset of 1121 Rembrandt radiographs. In this dataset, 299 (22.9%) images are deleted as having too few chain lines. Among the remaining 822 images, 805 images (97.9%) are well marked, and 17 images (2.1%) are rejected. Using the semi-automated method described above, it requires about one minute to extract the chain lines from each image. In contrast, manually marking three point on each chain line requires about seven minutes, even with the use of marking software.

3 Comparison of Extracted Chain Lines

Each image in the Rembrandt set contains between 5 and 8 chain lines. The data extracted from each image by the technique of Section 2 consists of sets of (x, y) coordinates which indicate points along the chain lines, and each line is represented by up to 10 points. Since the goal is to align images by the underlying lines (and not by the points themselves), the first step was to fit a line to the measured points. MATLAB's *fit* function was used to fit a linear polynomial to the existing data, taking care to filter out missing data.

Each image is represented by a set of lines that are specified by their slope and y -intercept. A metric is needed which can compare the sets of lines to determine their similar or dissimilarity. The metric needs to allow the possibility that the paper used in the sketch could have been rotated or flipped arbitrarily. Moreover, the image could be just a portion of a larger paper, or be a shifted version (for examples, lines 1 through 5 of one image might correspond to lines 2 through 6 of another mage). Accordingly the measure needs to be invariant to such transformations. Throughout, there is an underlying assumption that the distance between two images should represent the minimum distance under all such transformations.

The Cartesian representation of the lines was first transformed into a (ρ, θ) parametrization

$$\rho = x \cos(\theta) + y \sin(\theta) \tag{1}$$

where θ is the angle between the line and the y -axis, and ρ is the distance between the origin and the closest point on the line. This transforms each line in Cartesian space into a point in the (ρ, θ) -plane. The distance metric for the difference between two lines is then defined to be a weighted l_2 distance between the points (ρ^1, θ^1) and (ρ^2, θ^2) in the (ρ, θ) -plane. Two images are then

compared using the l_1 norm (absolute value of the sum) of all the line distances. To be precise, consider two images I_1 and I_2 , each consisting of n lines $I_1 = \{(\rho_1^1, \theta_1^1), \dots, (\rho_n^1, \theta_n^1)\}$ and $I_2 = \{(\rho_1^2, \theta_1^2), \dots, (\rho_n^2, \theta_n^2)\}$. The distance between I_1 and I_2 is defined as

$$d(I_1, I_2) = \sum_{i=1}^n \sqrt{(\rho_i^1 - \rho_i^2)^2 + \lambda(\theta_i^1 - \theta_i^2)^2} \quad (2)$$

for some weight $\lambda > 0$.

In order to optimally align the lines for comparison, it is necessary to find $\min d(I_1, I_2)$ under all allowable transformations, that is, under all transformations that do not distort Euclidean distance. Perhaps the simplest way to achieve this is to make each coordinate zero-mean. Thus

$$\bar{I} = ((\rho_1 - \bar{\rho}, \theta_1 - \bar{\theta}), \dots, (\rho_n - \bar{\rho}, \theta_n - \bar{\theta}))$$

where $\bar{\rho}$ and $\bar{\theta}$ are the means of the ρ_i and θ_i , respectively. Then the distance between I_1 and I_2 is calculated as $d(\bar{I}_1, \bar{I}_2)$, since \bar{I}_1 and \bar{I}_2 both have a mean of $(0, 0)$, meaning that they are optimally aligned in the (ρ, θ) plane. Translations in the θ direction (increasing/decreasing the slope of the lines) correspond to rotations of the image, and translations in the ρ direction (increasing/decreasing the distance between the lines and the origin) correspond to translations of the image in the (x, y) plane. Thus these two transformations preserve the distance between parallel lines and the angles between non-parallel lines, and therefore do not distort Euclidean distances.

In computing $\min d(I_1, I_2)$, it is also necessary to account for arbitrary flips and/or reflections. With respect to the (x, y) coordinates, it is necessary to compare the lines with the y -coordinates negated, the x -coordinates negated, both, and neither. Since there are only 4 possibilities (including the identity), it is easy to exhaustively search over all four reflections and hence find the minimum. Also note that negating either or both of ρ and θ provides a simple way to calculate these reflections.

Finally, this metric assumes that both images have the same number of chain lines n . There is a combinatorial issue when comparing images that differ in their number of chain lines. For example, if one image has 8 chain lines and another has 7, then there are two possible comparisons to make when computing the metric using 7 chain lines, but there are 7 possible comparisons when considering 6 chain lines (and so on). Fortunately, it is easy to iterate over all possible comparisons.

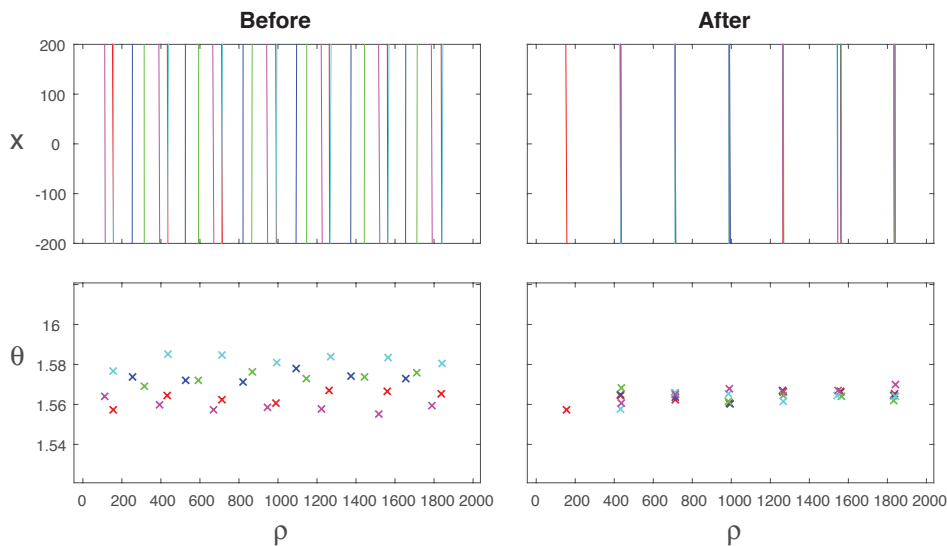


Figure 2: Each color corresponds to a different image. The top row represents the chain lines in the (x, y) plane, and the bottom row represents the chain lines as points in the (ρ, θ) plane. The figures on the right show how the algorithm aligned the images to minimize the sum of distances in the (ρ, θ) plane.

Using this metric, it is possible to compare images to look for possible matches. The simplest way to proceed is to compute the pairwise distance between all images using the metric with a constant “# of chain lines to consider” parameter, and then sort to create a list of possible matches. When conducting this for 6 chain lines, the top match was found to be an almost near-perfect comparison. On closer inspection, this turned out to be a duplicate image in the collection. Figure 2 shows a less-trivial clustering of promising results.

4 Possible Matches in Rembrandt’s Papers

A natural way to assess the results is to compare the ranked list of matches to known matches among Rembrandt’s papers. Since moldmates can often be identified by watermark matches, each of the matches in the ranked list can be classified as a true positive (moldmate match) or a false positive. Identifying watermarks can be done using Hinterding’s catalogue [2]. Cross indexing the Bartsch number and the museum identification code associated with each image, Appendix 2 (“Concor-

dance by Bartsch number” in [2]) narrows down the type, variant, and subvariant of the watermark. The main catalogue can then be used to confirm the watermark classification.

The results of a reduced subset of the data were analyzed since there are fewer images with a greater number of chain lines. For example, with the “# of chain lines to consider” parameter set to 9 chain lines, there are only 12 images, but there are 48 images with 8 chain lines. Using these 48, the program ranked 1128 unique matches. For all 48 images, the watermarks were (manually) identified using the process above. Among these watermark identifications, 12 moldmate matches were found, as well as one pair of duplicate images. Among the 1128, the moldmate matches appeared at rankings 17, 43, 61, 73, 75, 88, 108, 111, 126, 147, 65, and 167. The match of the duplicate image appeared at rank 7. Note that all these matches are in the top 14.8% of the results. Though there are false positives in between, the 15% benchmark suggests a relatively small cutoff point for positive matches, suggesting that with unfamiliar data sets, chain line matching may be useful as one forensic in detecting moldmates.

In attempting to improve the results, either by pushing the moldmate matches further up toward the top rankings, or by removing the false positives between them, the weight λ in the cost function $d(I_1, I_2)$ of Eqn. (2) was altered. Recall that λ adjusts the relative importance of the chain line spacing and the angles θ between the lines. The results above used $\lambda = 1000$. Increasing λ to 1200, the moldmate matches then appear in the top 15.3% of the rankings. At 5000, the matches appear in the top 35.7%. Adjusting in the other direction, decreasing λ to 800, the value goes to 14.9%. Decreasing further to 100, 0.1, 0.01, and even 0.001, the value holds steady at 15.2%. Fig. 3 shows a graph of λ values vs. the top percent of rankings in which the moldmate matches appear.

A relatively small increase in λ results in a sharp increase of the range of the moldmate rankings, and decreasing λ to nearly 0 has a negligible effect, suggesting that the angle θ could be removed from the cost function, so that the cost measured only the difference in the line’s shortest distance to the origin. Indeed, with $\lambda = 0$, the moldmate matches appear in the top 15.2% of rankings, though the particular rankings shift around slightly. Since θ represents the angles of the lines, the assumption that chain lines are parallel explains the negligibility of θ in d . Likewise, because differences in ρ between lines corresponds to differences in the chain line spacing, the importance of ρ in the cost function supports the hypothesis that chain line spacing is the most pertinent parameter.

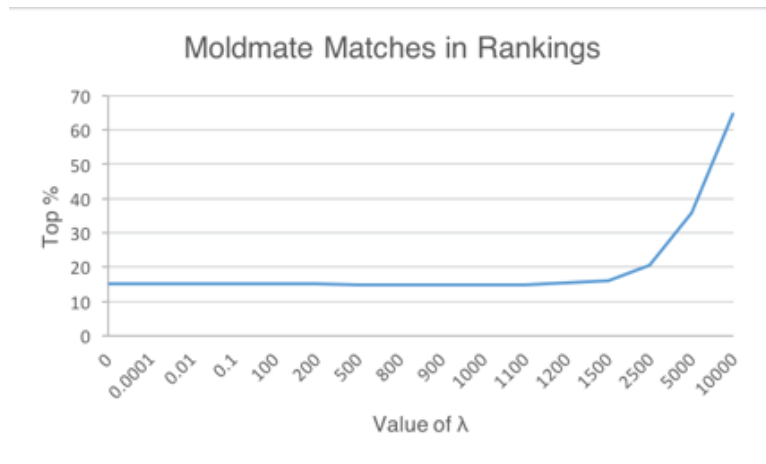


Figure 3: λ values vs. the top percent of rankings in which the moldmate matches appear

References

- [1] C. R. Johnson, Jr., W. A. Sethares, M. H. Ellis, S. Haqqi, R. Snyder, E. Hinterding, I. van Leeuwen, A. Wallert, D. Christoforou, J. van der Lubbe, N. Orenstein, A. Campbell, and G. Dietz, “The Application of Automated Chain Line Pattern (CLiP) Matching to Identify Paper Moldmate Candidates in Rembrandt’s Prints” in *New Directions in the Study of Rembrandt and His Circle*, in press. <http://people.ece.cornell.edu/johnson/CLiP-Herst.pdf>
- [2] E. Hinterding, *Rembrandt as an Etcher: The practice of production and distribution*, chapter 1 (“Rembrandt’s papers”), in vol. 1 of 3, Sound & Vision, 2006.
- [3] C. R. Johnson, Jr., W. A. Sethares, M. H. Ellis, and S. Haqqi, “Hunting for paper moldmates among Rembrandt’s prints,” *IEEE Signal Processing Magazine*. June, 2015.
- [4] RKD Image Sharing, <http://english.rkd.nl/Services/image-sharing>
- [5] Xuelei Xi, “Semi-automated chain line marking in radiographs of Rembrandt’s prints,” M. Eng project, Cornell University, 2016.