

Hunting for Paper Moldmates among Rembrandt's Prints

C. Richard Johnson, Jr., William A. Sethares, Margaret Holben Ellis, and Saira Haqqi

Early paper manufacturing used sieve-like molds through which paper pulp was drained. Two pieces of paper are called *moldmates* if they were made using the same mold. When there exists a large body of one artist's work on paper, the identification of moldmates can help in establishing chronology, suggest paper preferences, and indicate periods of intense activity of the artist. Rembrandt is an especially good example. With several thousand prints in existence today, the study of Rembrandt's prints has occupied scholars for over two centuries, and the study of his printing papers occupies a prominent place within this scholarship [1]. This paper examines the feasibility of moldmate identification via chain line pattern matching, and conducts a series of experiments that demonstrate how accurately the measurements can be made, how straight and how parallel the lines may be, and provides a rule-of-thumb for the number of chain lines required for accurate moldmate identification using a simplified model. The problem of identifying moldmates among Rembrandt's prints is presented as a pair of image/signal processing tasks; our strategy is to provide basic solutions to these tasks and to then reveal the shortcomings of these solutions in the hopes of encouraging future work in the signal processing community. With the support of the Morgan Library & Museum and the Metropolitan Museum of Art, we have made high resolution data available [2] to facilitate this quest.

LAI D LINES, CHAIN LINES, MOLDMATES, AND WATERMARKS

Before the introduction of the papermaking machine in the early 19th century, handmade paper was created by dipping a *mold* – a porous screen surrounded by a removable wooden frame – into a vat of mascerated and suspended pulp. Prior to 1750, the screen was fabricated from fine densely spaced horizontal rows of *laid* wires held into position by thicker, widely spaced vertical *chain* wires. When the mold was plunged into the vat and lifted out, these wires acted as a sieve, retaining the paper pulp in thinner and thicker accumulations as the water drained away. A modern re-enactment of this

method of paper making can be found online [3]. The grid-like configuration – called a *laid and chain line* pattern – formed by the interference of the wires, is replicated in the final sheet and is easily visible when the blank paper is viewed via transmitted light or imaged via radiography. For example, Fig. 1 shows one of Rembrandt's prints and a beta-radiograph of the same print in the region around the watermark. The chain lines are the whitish vertical features spaced slightly more than 2 cm apart. The laid lines are the closely spaced (and barely perceptible) horizontal lines. Each mold was made by hand and, while at first glance, two molds may appear to be identical, small variations exist between the exact intervals of chains from one mold to the next. Papers having identical laid and chain line patterns are commonly presumed only to occur if they have been formed on the same mold – hence they are called moldmates.

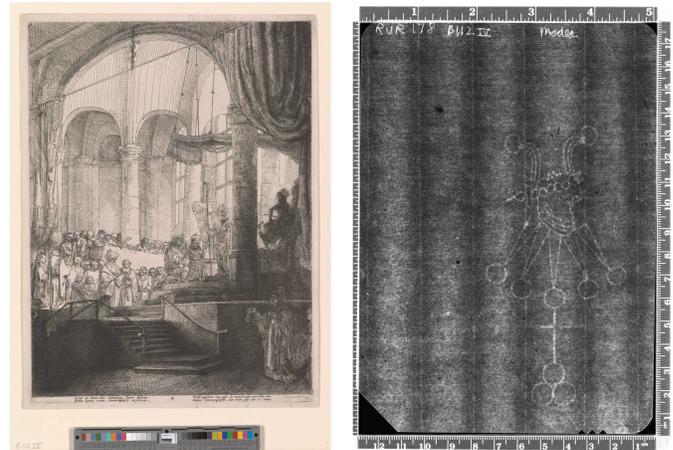


Figure 1. Rembrandt's etching "Medea, or the Marriage of Jason and Creusa," (Bartsch 112iv) is shown in visible light and as rendered, in a detail, by beta-radiograph. The area near the watermark is shown in the beta-radiograph where the chain lines are the vertical features that are spaced a little less than one inch apart (see the top ruler). Thanks are extended to Reba Snyder for providing the beta-radiograph of B112. (Etching Photo Credit: The Morgan Library & Museum, New York. RvR 178. Photography by Graham S. Haber, 2014.)

Often a watermark designating the paper's manufacturer, size, or function was stitched onto the heavier chain lines using a wire bent into a simple shape, for example, a star, shield or monogram. Just like the chain and laid lines, the watermark wire affected the rate and quantity of pulp as it drained through the mold, and left behind a characteristic impression on the paper. Typically, images of watermarks are captured for art historical research using beta-radiography or low-energy x-radiography [1]. For example, the watermark in

C. Richard Johnson, Jr. is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY and the Rijksmuseum, Amsterdam, johnson@ece.cornell.edu

W. A. Sethares is with the Department of Electrical and Computer Engineering, University of Wisconsin, Madison, WI and the Rijksmuseum, Amsterdam, sethares@wisc.edu

M. H. Ellis is with the Conservation Center, Institute of Fine Arts, New York University, and the Thaw Conservation Center, Morgan Library & Museum, mhe1@nyu.edu

S. Haqqi is with the Conservation Center, Institute of Fine Arts, New York University, haqqis@gmail.com

Fig. 1 is known as a “foolscap” and depicts the head of a jester wearing a five-pointed ruff surmounting a cross emerging from three spheres. Twenty-one variants of foolscap watermarks found on Rembrandt’s prints are cataloged in [1], one of which has nine subvariants, with one of these appearing in Fig. 1. This particular subvariant can be found in sixteen different prints by Rembrandt according to [1]. Papers having identical watermarks may also be presumed to be moldmates.

Watermarks have received considerable attention in print, drawing, book, and manuscript connoisseurship, with scholars carefully recording watermarks in the works of Shakespeare, Beethoven, Jane Austin, Dürer, and other artists and writers. To date, moldmates have been identified primarily by comparing watermarks. However, extracting and comparing watermarks using computer-generated algorithms is a challenging proposition, as they have complicated shapes (as shown by the foolscap) and can assume many forms and have many variations. Another drawback to the use of watermarks for moldmate identification is that not all prints under consideration have watermarks. Indeed, only approximately one-third of the extant impressions of Rembrandt’s etchings include a full watermark or even a fragment[1]. Accordingly, the use of the chain line pattern has been suggested as way to identify papers made on the same mold [4]-[5]. The difficulty in manually pursuing moldmate identification via chain line pattern matching has prompted the consideration of its automation [6]-[8], though no automated scheme has yet been universally adopted.

CHAIN LINE PATTERN MATCHING AND SIGNAL PROCESSING

The approach taken here is to characterize the problem of identifying moldmates among Rembrandt’s prints – and by extension any handmade *laid and chain line* patterned papers such as those found in manuscripts, printed books, archives, and collections of prints and drawings – as a basic pair of image processing tasks.¹ The first task is to locate the chain lines in images such as Fig. 1; the second is to use the chain line patterns to identify potential moldmates. Our strategy is to provide basic solutions to these two tasks that are sufficient to convince the user community to collaborate in a future imaging campaign providing full sheet images to maximize the extent of the chain line pattern associated with each piece of paper. To help encourage members of the signal processing community to contribute to this topic of paper moldmate identification we point out here some of the limitations of our basic solutions and note that – with the support of the Morgan Library & Museum and the Metropolitan Museum of Art – we have made high resolution images available [2], which will be needed for developing improvements.

Our algorithmic approach to this computational art history task is similar to the strategy in [7], which uses image processing tools (in particular Fourier and Radon transforms) to extract from a suitable image the average chain line separation

¹While considerable progress has been made in the application of signal processing to the technical analysis of canvas supports for paintings [9]-[13], there has been much less focus on the application of signal processing to handmade laid paper.

distance, chain line orientation, number of chain lines, and the sequence of chain line separation distances. As stated in [7]: “All these features are detected under the assumption that the lines are straight and equidistant with respect to each other.” [Here equidistant means parallel.] In [7], the average chainspacing and the chainspace sequence are combined with the laid line density in a similarity measure. Our approach in the present paper differs in that we drop the extraction and use of laid line density as most of our images are of insufficient clarity to assess this feature. Plus, we abandon the assumption that the chain lines are parallel; while we retain the assumption that they are straight. The lead author of [7] modified the straightness assumption to piecewise straightness in [8], with a commensurate increase in the necessary signal processing.

In this paper, we test the straightness of the chain lines in the images in our dataset and find that numerically adequate near-straightness occurs in approximately 90% of our images. Thus, in our quest to establish the ability of chain line pattern matching of a sufficient number of adjacent chainspaces by itself to offer a reduction in a broad library to a manageable number of potential matches for further investigation we assume chain line straightness. We also test our dataset for the occurrence of chain line patterns with at least one non-parallel chain line with a relative angle greater than 0.12 degrees. This is a sizable portion of our dataset (see Fig. 5 for details). We chose to include the possibility of handling images with non-parallel chain lines, for which the location of the measurement of chainspacing on the image relative to its location on the mold needs to be computed, because of the fundamental simplicity of this task. The end result is that we agree with [7] in that the chainspace sequence is “the most discriminative feature”.

Locating Chain Lines

A semi-automatic method of chain line extraction in beta-radiographs is shown in Fig. 2. This begins in (a) with a manual trimming (b) to remove any labels, ruler markings, and blank areas from the source image. The trimmed image is input into a Radon transform (c) and the angle at which the chain lines are (closest to) vertical is given by the column with the greatest variation. The graph (d) shows the standard deviation of the elements in each column, and the angle θ^* with the maximum value gives the angle of rotation. In (e), the trimmed image is rotated by $-\theta^*$ to give the straightened image (f). This is then filtered using an aggressive vertical filter (g) which is $1/3$ the image height and 3 pixels wide, resulting in the filtered image (h). A Hough transform (i) then locates the most prominent lines, and the result is shown in (j), superimposed over the straightened image.

Observe that although the positions and orientations of the lines can be determined automatically in this procedure, both the preparation of the image and the final verification are done manually. It is difficult to control the exposure of beta-radiographs, since under- and over-saturation can occur based on the properties of the paper, which may not be fully known at a time of exposure. As a result, it is common to adjust

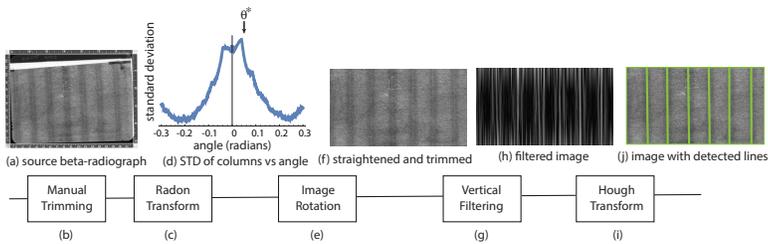


Figure 2. A semi-automatic method of chain line extraction.

the contrast on the images and to trim unwanted portions of the raw data. The images commonly contain a ruler (as in Figs. 1 and 2(a)) which is needed to verify physical dimensions, and there may be regions of all-black or all-white (such as the wedge shaped region in 2(a)) that need to be cropped. While it may be possible to automate some of these actions, mistakes in the preprocessing almost guarantee that the chain lines will not be found properly.

Once the procedure (b)-(c)-(e)-(g)-(i) of Fig. 2 has been run, it is necessary to check that the output is sensible. Typical errors in the algorithm would include failing to find some of the lines or finding too many lines. In either case, the algorithm can be rerun with different thresholds and settings within the filters and transforms. Clearly, the method suggested in Fig. 2 is but one possibility. See [7] for another that begins by exploiting the shadow around the chain lines rather than the chain line itself.

As all approaches will, our method makes several assumptions about the chain line patterns. First, it assumes that all the lines are more-or-less parallel (because otherwise the Radon transform cannot locate a single preferred direction for the step (e) derotation). Second, it assumes that all the lines are genuinely straight, since curved or segmented lines cannot be located by the Hough transform reliably. The open problem of finding more automated, more general, more accurate, and/or faster algorithms is one task we are promoting here to signal processing specialists.

Rather than pursuing an improved or more general algorithm here, consider the question of accuracy. It is easy to look at the superimposed output of Fig. 2(j) and to see that it “looks” correct. It has detected the correct number of lines and they are located plausibly. But what is the “real” answer, and how close does the algorithm come to this answer? Indeed, this becomes crucial when comparing different algorithms, since only by comparing with a ground truth can the accuracy of different candidate algorithms be compared.

Experiment #1: A Test of Accuracy

To see why the task of locating the chain lines may not be as straightforward as it might seem, Fig. 3 enlarges the beta-radiograph B112iv twice. While the left and middle images clearly show the chain lines as vertical linear features, the right hand image shows that these are not simple geometric lines. Rather, they are composed of uneven and blurry patches, they are not of uniform width (often extending several pixels to either side), and they may be interrupted repeatedly.

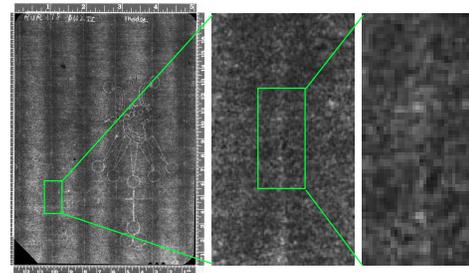


Figure 3. A small region containing a portion of a chain line is extracted from the beta-radiograph in Fig.1. This is then enlarged twice. What appears to be a fairly clear vertical line in the left dissolves into a blur of pixels on the right. Where exactly is the “line”?

To investigate, we asked a group of engineering students at Cornell University taking a course on signal processing algorithms for analyzing art supports [14] to manually identify the chain lines in several beta-radiographs. The subjects were given marking software (a simple graphical user interface built in Mathematica) which allowed placement of marks near the endpoints of the chain lines. A semi-transparent line connected the two endpoints so that the subjects could visually verify their marking. Such manual markings can be used to establish the ground truth of “where the lines really are” and can also be used to assess the agreement (or disagreement) among the subjects.

Since the subjects could choose where to mark, it was not possible to compare the marked locations themselves. Rather, the lines formed by joining the two marked endpoints were parameterized in distance/angle form

$$\rho = x_i \cos(\theta) + y_i \sin(\theta) \quad (1)$$

where the (x_i, y_i) are the Cartesian coordinates of the two endpoints, ρ is the distance from the origin and θ is the angle of the line. With two marked points, (1) is a system of two equations and two unknowns, and hence can be solved for (ρ, θ) . The (ρ, θ) values for each of the lines marked by each of the subjects were tabulated. Altogether, there were twelve subjects who worked with 24 beta-radiographs chosen randomly from our larger set. Each subject measured 4-5 lines on 8 different beta-radiographs and so each chain line was measured independently four times. Over this data set, the mean of the ρ values was 1.06 and the standard deviation was

$$\sigma_e = 0.009, \quad (2)$$

which corresponds to about 5.3 pixels at a nominal resolution of 600 dpi. This shows that despite the patchy nature of the chain lines, subjects can locate the chain lines with reasonable consistency. It also provides a way to assess the quality of an algorithm. Indeed, applying the semi-automatic procedure of Fig. 2 to the same beta-radiographs gives values of ρ that are within one standard deviation of the measured values.

Experiment #2: A Test of Straightness

Both the algorithm of Fig. 2 and the discussion of the geometry of moldmates in Fig. 4 presume that the chain lines are *straight*. Stalduin [8] observes that some chain lines

may be bent, and develops a method that attempts to locate the piecewise segments of such bent lines. Our observation was that the majority of chain lines did not appear to bend, and so we wanted to quantify the straightness of the chain lines. Towards this end, we asked the subjects to manually find the smallest bounding box that encloses the centered spine of each chain line. The subjects and image data were the same as in Experiment #1. The subjects now used the software to mark three points for each chain line, from which the bounding boxes can be calculated. These three points were to be placed in the horizontal center of the vertically-oriented chain line with one each near either end and one at some point of extreme departure from the straight line connecting the marked end points. Again, each bounding box is measured four times, by four different subjects. Fig. 5 shows a histogram of the widths of the chain line bounding boxes. The average width is 7.0 pixels, and the median is 5.3. Both of these numbers are comparable to (2), which is the accuracy to which the measurements are made – hence, over half of the chain lines have width smaller than can be reliably measured.

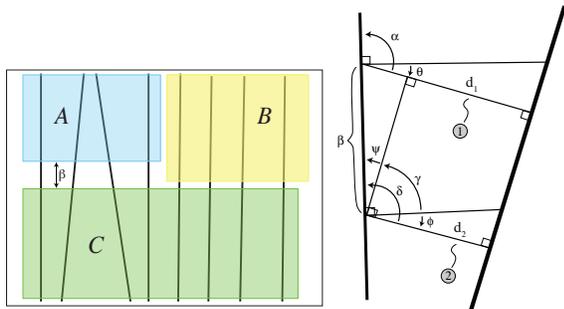


Figure 4. A paper mold is shown in schematized form where the non-parallel chain lines have been exaggerated. The paper has been cut into three pieces labelled A, B, and C. Given a large collection of papers, the goal of moldmate identification is to find those papers that come from the same mold. The geometry of two non-parallel chain lines is shown on the right.

On the other hand, there are isolated chain lines which have bounding boxes with nontrivial width. In this data set [2], the largest is the leftmost chain line in the image 17.37.75, which is measured as 52.7, 27.48, 29.83, and 30.97 pixels by the four subjects. While these values are not particularly consistent, they are all well above the measurement error, indicating a significantly bent chain line. The second largest values are from the left-most chain line of image 17.37.75 (2), which were reported as 40.13, 32.08, 39.5, and 36.2. Interestingly, these two distinct images from our dataset, included in the grouping of [2], are of the same of piece paper and should provide a very close match. However, they did not for the software we subsequently developed, which assumed chainline straightness. They were assessed as somewhat similar but not the closest match.

Experiment #3: A Test of Parallelism

A common assumption is that the chain lines in a single print are parallel [7]. To test this assumption, this third experiment uses a set of $N = 205$ beta-radiographs we received from the Morgan Library & Museum, the Metropolitan Museum of

Art, and the Rijksmuseum to measure the degree to which the lines are parallel, by looking at the difference between the angles of the chain lines. Fig. 5 shows a histogram of the maximum angle that occurs in each print. The mean and median values are 0.4 and 0.36 degrees, so overall the lines tend to be fairly parallel. As will be discussed in succeeding sections, when the lines are not parallel, this can be quite useful as an identifying feature of the print.

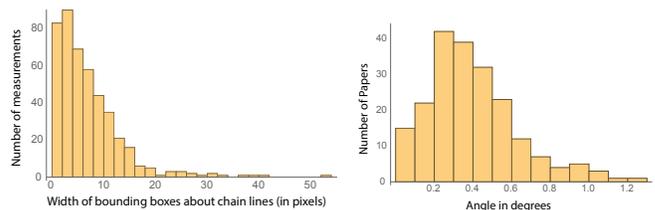


Figure 5. The left hand image shows a histogram of the measurements of the widths of the bounding boxes about the chain lines. The majority of chain lines appear straight, though there are outliers. The methods of this paper will not be effective for these outliers. The right hand image shows a histogram of the maximum angle found between the chain lines in each print gives a measure of how parallel the chain lines are. The majority of papers have fairly parallel chain lines; those that do not may be more readily identified by the methods of this paper.

The Geometry of Moldmates

The dimensions of the molds used in paper making were typically much larger than the sizes of final folded sheets as used for pages in a book or unfolded as stand-alone sheets. Fig. 4 shows an exaggerated schematic of a mold with eight chain lines that has been divided into three sheets labeled A, B, and C. It is worth making a few observations. First, chain line matching cannot show directly that papers A and B come from the same mold since they have no chain lines in common. Second, measurements of chain spacings (inter-chain distances) may show a close match between B and C (because the corresponding lines are parallel), but they will not show a close match between A and C (because the lines have different angles). Third, given two pieces of paper with many chain lines, it is necessary to check all possible subsets for matches. For example, only the final four lines of C can match with the lines of B. Though not obvious from the figure, it is also worth mentioning that it is typically unknown *a priori* whether a paper has been flipped left-right, top-bottom (or both) when taking the beta-radiograph; thus it is also necessary to check all four orientations when attempting to find matches. Finally, the more chain lines that overlap, the more certain one can be that two papers do (or do not) match.

Using the distance/angle form (1) for the detected chain lines allows a geometric calculation to compensate for the angle discrepancy. Essentially, this estimates the distance β in Fig. 4, though it may be pictured geometrically as sliding the chain lines of paper A up and/or down until they best match the chain lines of paper C. The geometry of two non-parallel chain lines is illustrated in the right hand side of Fig. 4. The distances d_1 and d_2 are two measurements of how far apart the lines are, and are related via

$$d_1 = d_2 + \beta \sin(\psi). \quad (3)$$

Because segments ① and ② are parallel, angle α equals angle δ . (Counterclockwise angles are positive while clockwise angles are negative.) Thus, because $\alpha = 90^\circ - \theta$ and $\delta = 90^\circ - \phi$, $\theta = \phi$. Furthermore, because $\gamma + \psi = 90^\circ$ and $\gamma - \phi = 90^\circ$, $\psi = -\phi = -\theta$. A key observation is that the angle off horizontal of the shortest distance line to the right line from any measurement point on the leftmost line is the same.

Given a collection of N laid papers each with M chain lines (presumed to be straight but not necessarily parallel), the first step is to measure (automatically, semi-automatically, or manually) the positions of the chain lines. Using the distance/angle parameterization (1), the M chain lines in paper i can be represented by the vector

$$\{\rho^i, \theta^i\} \equiv \{\rho_1^i, \rho_2^i, \dots, \rho_M^i, \theta_1^i, \theta_2^i, \dots, \theta_M^i\} \in \mathbb{R}^{2M}, \quad (4)$$

where ρ_k^i and θ_k^i represent the distance and angle of each chain line k with respect to the origin of the i th paper. For ease of notation, these are partitioned into vectors ρ^i and θ^i , each in \mathbb{R}^M .

Consider two pieces of paper i and j that are separated by an unknown distance β on the same mold (for example, paper A and the first four chain lines of C in Fig. 4). Applying the logic of Eqn. (3) to each of the successive pairs of chain lines suggests estimating β using a Least Squares procedure

$$\min_{\beta} \left\| \begin{pmatrix} \rho_1^j - \rho_1^i \\ \rho_2^j - \rho_2^i \\ \vdots \\ \rho_M^j - \rho_M^i \end{pmatrix} - \beta \begin{pmatrix} \sin(\frac{\theta_1^j + \theta_1^i}{2}) \\ \sin(\frac{\theta_2^j + \theta_2^i}{2}) \\ \vdots \\ \sin(\frac{\theta_M^j + \theta_M^i}{2}) \end{pmatrix} \right\|. \quad (5)$$

Though the measurements θ_ℓ^j and θ_ℓ^i are similar, they are not identical, which is why they are averaged. Rewriting (5) using $Y_{i,j}$ for the first vector and $X_{i,j}$ for the second, the β that minimizes the least squares error $\|Y_{i,j} - \beta X_{i,j}\|_2 = (Y_{i,j} - \beta X_{i,j})^T (Y_{i,j} - \beta X_{i,j})$ is $\beta_{i,j}^* = (X_{i,j}^T X_{i,j})^{-1} X_{i,j}^T Y_{i,j}$.

Identifying Moldmates

The geometric analysis of the previous section aligns the chain lines of two papers (as well as possible) by offsetting them a distance $\beta_{i,j}^*$. Moreover, the value achieved at this optimum

$$d_{\beta^*}(i, j) = \|Y_{i,j} - \beta_{i,j}^* X_{i,j}\|_2 \quad (6)$$

provides a measure of the dissimilarity between the chain lines of the papers i and j . A value of zero would mean that the shifted versions overlay exactly while a large value implies that the two sets of chain lines are very different. Moldmates should have small $d_{\beta^*}(i, j)$ while unrelated papers should have larger values.

Initial experiments showed that applying (6) directly to the problem of finding moldmates can lead to some undesirable false matches. Fortunately, many of these can be removed by considering more than just the value of $d_{\beta^*}(i, j)$. We have found four such modifications useful. First, if the difference between any of the angles is too large, any small value of $d(i, j)$ is accidental; such accidents can be removed from consideration by placing a threshold on $\max |\theta^j - \theta^i|$. Second,

if all of the individual angles are small, then the computation of β^* is irrelevant, and the distance

$$d(i, j) = \|\rho^j - \rho^i\| \quad (7)$$

is more appropriate. Third, a β^* value that is larger than about 18 in (a typical size for the frame) is impossible; any such false matches can be removed by placing a threshold on β^* . Finally, if the maximum of all the $|\rho^j - \rho^i|$ is small, a distance of zero is reported to encourage consideration of this potential mold match. These can be written succinctly in two parts:

if $\max_j |\theta^j| < 0.005$,

$$\delta(i, j) = \begin{cases} \infty & \max |\theta^j - \theta^i| > 0.02 \\ 0 & \max |\rho^j - \rho^i| < 0.01 \\ d(i, j) & \max |\rho^j - \rho^i| > 0.01 \end{cases}$$

if $\max_j |\theta^j| \geq 0.005$,

$$\delta(i, j) = \begin{cases} \infty & \beta_{i,j}^* > 18 \\ 0 & \max |\rho^j - \beta_{i,j}^* \sin(X_{i,j}) \rho^i| < 0.01 \\ d_{\beta^*}(i, j) & \max |\rho^j - \beta_{i,j}^* \sin(X_{i,j}) \rho^i| \geq 0.01 \end{cases} \quad (8)$$

where $X_{i,j}$ is defined as in (5). The various thresholds (0.005, 0.01, 0.02, etc.) are representative and may require fine tuning for specific data sets. The threshold on the differences in angle and distance are set so they are larger than the differences resulting from manual marking errors. The range of measurement point separation is set by typical mold dimensions. The threshold for declaring the chain lines parallel (so that beta need not be estimated) depends on the shift causing changes in the distance/spacing vector that exceeds the threshold that designates a match. For a modest β of 10 inches and a reasonable distance threshold of 0.02 inches, the sine of the relative angle should be less than 0.002. For such small values, this corresponds to the angle threshold for declaring non-parallelism of 0.002 radians or 0.12 degrees. By this measure a large fraction of our images contain non-parallel chain lines.

HOW MANY CHAIN SPACES?

The objective in moldmate identification is not to return a single answer with the most likely fit, since this is not what the paper conservator or art historian would find most useful. Rather, the goal is to reduce a large library of chain line pattern images to a small number so that the expert user can investigate further, with the expectation that if a match exists, it will likely be in this smaller set. One issue is how many adjacent chain spaces are needed to achieve this goal. Currently, the majority of the beta-radiographs available to us are limited to the vicinity of watermarks, which typically do not cover the entire print. This is due to the earlier emphasis on watermark matching and the technical limitations of beta-radiography [15]. An assessment of the required number of adjacent chain spaces to achieve reliable matching would be helpful to persuade museums to undertake the expense in equipment and personnel costs to mount a campaign to acquire full-print images revealing its entire chain line pattern. This section addresses that need by building a simple statistical model of the chain-line process and assessing the probability of error of the model as a function of the number of chain

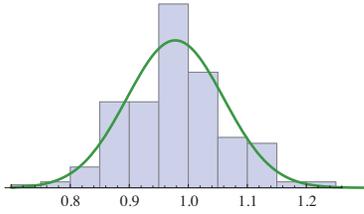


Figure 6. A histogram of the mean chain line spacing of the N sheets in the database is reasonably modeled as a Normal probability distribution with mean $\mu = 0.977$ and standard deviation $\sigma_m = 0.082$, as shown in green.

spaces. Specifically, the analysis calculates the probability that a sheet of interest will be confused with one or more of the existing sheets and the results provide guidelines for the number of chain lines that need be present in order to reliably detect moldmates. This provides a rule-of-thumb for the real problem.

In the past, attempts have been made to use the average inter-chain spacing as a way of characterizing sheets of laid paper for the purpose of identifying mold-mates [6]. This section begins by conducting a statistical analysis of this procedure (using data gathered as in the previously discussed experiments) to show why the mean value alone is unlikely to provide a useful characterization, which agrees with observed behavior in experiments reported in [7]. The analysis is then extended to consider the use of $M + 1$ chain lines (M inter-chain spaces), and a simple rule-of-thumb is derived that suggests how large M must be in order to reliably distinguish mold-mate matches to a candidate sheet of paper from among a universe of $N \gg M$ sheets.

The database of images of chain line patterns in laid papers (provided by the Morgan Library & Museum, the Metropolitan Museum of Art, and the Rijksmuseum and used in these experiments) includes $N = 205$ sheets. The mean chain spacing, calculated as $\rho_{i+1} - \rho_i$ where i ranges over all the chain lines in a given image, is $\mu = 0.977$ inches with standard deviation $\sigma_m = 0.082$. Considered as a collection of random samples, this is plausibly Gaussian, as can be seen in the histogram of Fig. 6. To model the question of interest, consider a collection of N random variables $\bar{m}_i \sim \mathbf{N}(\mu, \sigma_m^2)$ each of which is measured in the presence of some small error

$$m_i = \bar{m}_i + \epsilon_i \text{ where } \epsilon_i \sim \mathbf{N}(0, \sigma_\epsilon^2). \quad (9)$$

A new candidate element $m^* \sim \mathbf{N}(\mu, \sigma_m^2)$ is measured from the same distribution as the m_i , and we wish to know how many of the m_i lie close to this candidate m^* . (These will be the potential moldmates that must be subjected to closer examination.) This also requires quantifying ‘‘close to.’’ For simplicity, consider two elements close if they lie within ℓ standard deviations of the measurement error, that is, if they lie within $\pm \ell \sigma_\epsilon$. This is shown schematically in Fig. 7, where the shaded region S lying between the lines defined by $m^* - \ell \sigma_\epsilon$ and $m^* + \ell \sigma_\epsilon$ represents the probability that elements of the data set lie close to the measured m^* .

Let

$$f(x, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (10)$$

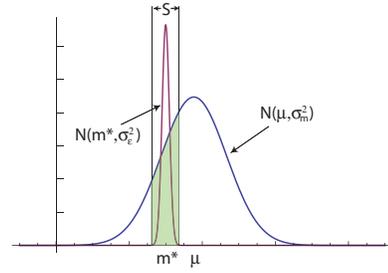


Figure 7. The means of the chain spacings are considered as elements of $\bar{m}_i \sim \mathbf{N}(\mu, \sigma_m^2)$. The region S shows the region in which elements are ‘‘close to’’ the test element m^* and the probability p_S represents the percentage of elements in the database close to the candidate.

be the normal density function. The probability represented by the shaded region S is

$$\begin{aligned} p_S(m^*, \sigma_m, \ell \sigma_\epsilon) &= \int_{m^* - \ell \sigma_\epsilon}^{m^* + \ell \sigma_\epsilon} f(x, \mu, \sigma_m) dx \\ &= \frac{1}{2} \left[\text{Erf}\left(\frac{m^* + \ell \sigma_\epsilon}{\sqrt{2}\sigma_m}\right) - \text{Erf}\left(\frac{m^* - \ell \sigma_\epsilon}{\sqrt{2}\sigma_m}\right) \right]. \end{aligned} \quad (11)$$

Some observations:

- (i) The probability p_S is independent of μ since both the data set and the new candidate element are assumed to have the same mean.
- (ii) As the measurement error $\sigma_\epsilon \rightarrow 0$, the probability of being close goes to zero. In words, the more accurately the measurements can be made, the greater the distinguishing power of the method. Conversely, larger σ_ϵ cause $p_S \rightarrow 1$.
- (iii) The probability p_S does depend on the particular m^* . For $m^* \ll \mu$ or $m^* \gg \mu$, p_S is small.
- (iv) Conversely, $m^* = \mu$ implies that $p_S = \text{Erf}\left(\frac{\ell \sigma_\epsilon}{\sqrt{2}\sigma_m}\right)$. With an eight-to-one ratio of σ_m to $\ell \sigma_\epsilon$, this is 0.1.

Case (iv) may be interpreted to say that roughly 10% of the elements of the database will be considered close to the candidate. Since each must be considered in four orientations, this approximately quadruples (to 40%) the percentage of sheets that will be considered potential matches. One ameliorating factor is that the measurement of the mean values are not independent of the number of chain spacings M ; larger M cause smaller effective σ_ϵ . If each independent chain line is measured with an error σ_{ϵ^*} , the average of M is effectively drawn from $\mathbf{N}(\mu, \frac{1}{M}\sigma_{\epsilon^*}^2)$ and the effective standard deviation is $\sigma_\epsilon = \frac{\sigma_{\epsilon^*}}{\sqrt{M}}$. Perhaps the best use of the mean value of the chain spacings is in the cases described by (iii) where m^* deviates from the mean μ . In these cases, p_S will be small and there will be fewer possible matches to consider. From another point of view, these estimates suggest that the mean value may be better used as a method of ruling out mold matches (in those cases when p_S is small) than as a method of locating mold matches in general.

Fortunately, more information is available in the chain spacings than just the mean value. The simplest way to model the chain spacings is to presume that the sequence of interchain intervals is

$$\bar{m}_i = \{m_{i,1}, m_{i,2}, \dots, m_{i,M}\} \text{ for } i = 1, 2, \dots, N, \quad (12)$$

where each component $m_{i,j}$ is independent of $m_{i,k}$ for $j \neq k$ and where \bar{m}_i is independent of \bar{m}_j for $i \neq j$. Following the logic of (9)-(11), the probability $p_S(\bar{m}^*, \sigma_m, \ell\sigma_\epsilon)$ of accidental matches from the database to a candidate vector m^* is

$$p_S(m_{i,1}^*, \sigma_m, \ell\sigma_\epsilon) p_S(m_{i,2}^*, \sigma_m, \ell\sigma_\epsilon) \cdots p_S(m_{i,M}^*, \sigma_m, \ell\sigma_\epsilon) \\ = (p_S(m^*, \sigma_m, \ell\sigma_\epsilon))^M. \quad (13)$$

Unfortunately, this exponential decay of 0.1^M represents an overly optimistic scenario in which each chain-space element is independent of the others even within a single sheet.

A more realistic model observes that each chain space sequence has a mean value m^* and a small variance $\sigma_{\bar{m}}^2$ about that mean. Let m^* be the mean value drawn from $\mathbf{N}(\mu, \sigma_m^2)$ as before. The chain spacings are then

$$\{m^* + \hat{m}_1, m^* + \hat{m}_2, \dots, m^* + \hat{m}_M\} \quad (14)$$

where $\hat{m}_i \sim \mathbf{N}(\mu, \sigma_{\bar{m}}^2)$ and where the standard deviation $\sigma_{\bar{m}}$ is the variation in the chain spacings within a given sheet (which is less than the variation in the complete data set). Assuming again that the m^* is independent of the \hat{m}_i and that \hat{m}_i is independent of \hat{m}_j for $i \neq j$, the desired probability $p_S(m^*)$ can be factored as

$$p_S(m^*, \sigma_m, \ell\sigma_\epsilon) p_S(\bar{m}_2, \sigma_{\bar{m}}, \ell\sigma_\epsilon) \cdots p_S(\bar{m}_M, \sigma_{\bar{m}}, \ell\sigma_\epsilon) \\ = p_S(m^*, \sigma_m, \ell\sigma_\epsilon) (p_S(\bar{m}, \sigma_{\bar{m}}, \ell\sigma_\epsilon))^{M-1}. \quad (15)$$

Since $\sigma_\epsilon < \sigma_{\bar{m}} < \sigma_m$, the probability $p_S(m^*)$ is larger than in (13), giving a more pessimistic (though also more realistic) assessment. Using the value $\sigma_{\bar{m}} = 0.034$ (the average of the standard deviations), $\ell\sigma_\epsilon \approx 0.01$ (the average measurement error from (2)), this is

$$p_S = (0.1)(0.25)^{M-1}. \quad (16)$$

By $M = 3$ (four chain lines), there is about $4p_S$ chance of false matches. The factor 4 occurs because of the need to inflate the number of sheets by four to consider all the possible rotations and reflections. By $M = 4$, $4p_S$ drops to 1% and continues to decrease exponentially as M increases, reducing the number of potential false matches by a factor of four with each additional chain line. This leads us to seek matching patterns of 5 (or more) adjacent chain lines with 4 (or more) chainspaces. A comparison of the predictions of the three sets of modelling assumptions is given in Table I.

model	fraction
#1 average chain line spacing	0.4
#2 independent sequence of chain lines	$4(0.1)^M$
#3 dependent sequence of chain lines	$4(0.1)(0.25)^{M-1}$

Table I

THE THREE MODELS MAKE INCREASINGLY MORE REALISTIC PREDICTIONS OF THE FRACTION OF IMAGES THAT WILL BE REGARDED AS "CLOSE," AS A FUNCTION OF THE NUMBER OF CHAIN LINES

A PAIR OF MOLDMATES IDENTIFIED VIA A CHAIN LINE PATTERN MATCH

Applying the distance measure $\delta(i, j)$ of (8) to the dataset of [2], we observed that there was a potential match between

Rembrandt's *Medea, or the Marriage of Jason and Creusa*, dated 1648 (B112iv, Fig. 1) and a left-right flipped version of *The Artist's Mother in Widow's Dress and Black Gloves*, B344. A reproduction of the etching and its beta-radiograph are shown in [16]. What is particularly intriguing about this match is that the latter print's authorship has been questioned for many years. Most scholars after Adam Bartsch, the indefatigable 18th century cataloguer of Rembrandt prints, agree that B344 is by a pupil of Rembrandt, who was most likely copying an earlier print, dated circa 1631, *The Artist's Mother Seated at a Table, Looking Right* (B433). This chain line pattern match of papers in restrikes from around 1650, provides support for the conclusion that the pupil's print originated in Rembrandt's studio. Presuming creation of B344 around 1635 narrows the list of potential copyists to pupils in Rembrandt's studio at that time. The close match between the (shifted) chain lines of these two images (with the proper flips) is shown in detail in [16]. Though this match was identified solely from the close similarity of their chain line patterns, it can be verified [1] by the match of the watermarks, though B344 has only a fragment of the watermark that is fully contained in B112iv.

CONCLUSIONS AND OPEN QUESTIONS

This paper has highlighted the problem of moldmate identification of laid paper, and provided a basic argument for the feasibility of the project. The potential for automated procedures to identify moldmates among collections of similar handmade laid papers is exciting to paper conservators and graphic art curators. Using our first-pass signal processing strategies and even with a modestly sized database, we were able to identify a moldmate pair of Rembrandt prints that was previously unrecognized by the owner.

This paper has attempted to clearly display the moldmate problem as a set of signal processing tasks with the goal of making the problem accessible to the signal processing community where those who are not expert in the technical analysis of paper may contribute. One key to this is a description of the various signal processing tasks that must be completed. The other major component is making the images and datasets publicly available [2] to help "fuel" this cross-disciplinary effort.

Some of the problems highlighted in this paper are straightforward applications of signal processing techniques while others may require significant effort to realize. For example, a wonderfully complex problem is the automated extraction and comparison of watermarks (such as that in Fig. 1). Such shapes are much more complex than the simple straight-line we have chosen to attempt to extract and the automated comparison between such unparameterized shapes is not straightforward. On the other hand, there are many tasks that might benefit from more advanced signal processing techniques, more clever implementations of the algorithms, and/or more thoughtful metrics. For instance, the proposed algorithm for chain line extraction (in Fig. 2) can no doubt be improved, both in its accuracy and in its range of applicability, reducing the manual portions with automated extraction techniques. Similarly, the distance measure (8) can no doubt be expanded or improved.

As shown in [16], such chain line pattern matching software can also be used to assess concerns paper experts have regarding the possibility that differences in the wetting, pressure, and drying of intaglio printing and conservation treatments encountered separately by moldmates can distort the chain lines into distinctly different patterns.

The analytic contribution of this paper is to propose a simple model where the statistics of the chain line dataset can be used to estimate the number of chain lines needed to distinguish laid papers. For example, three or four chain lines are inadequate to reliably locate moldmates, and the rule-of-thumb relates the certainty of any such identification with the number of chain lines, given the experimental and computational errors inherent in locating the chain lines.

REFERENCES

- [1] 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.
- [2] RKD Image Sharing, <http://english.rkd.nl/Services/image-sharing>
- [3] University of Iowa, “Chancery Papermaking,” <http://now.uiowa.edu/2013/05/center-book-releases-short-film-about-papermaking>
- [4] A. Stevenson, “Chain-Indentations in Paper as Evidence,” *Studies in Bibliography*, vol. 6, pp. 181-195, 1953-54.
- [5] D. L. Vander Meulen, “The Identification of Paper without Watermarks: The Example of Pope’s ‘Dunciad’, ” *Studies in Bibliography*, vol. 37, pp. 58-81, 1984.
- [6] J. C. A. van der Lubbe, E. P. van Someren, M. J. T. Reinders, “Dating and authentication of Rembrandt’s etchings with the help of computational intelligence,” *Proc. International Cultural Heritage Informatics Meeting (ichim 01)*, pp. 485-492, September 2001.
- [7] M. van Staalduinen, J. C. A. van der Lubbe, E. Backer, and P. Paclik, “Paper Retrieval Based on Specific Paper Features: Chain and Laid Lines” in *Multimedia Content Representation, Classification and Security*, Lecture Notes in Computer Science, vol. 4105, pp. 346-353, 2006.
- [8] M. van Staalduinen, “Content-based Paper Retrieval Towards Reconstruction of Art History,” *PhD Thesis*, TU Delft, 2010.
- [9] C. R. Johnson, Jr., D. H. Johnson, N. Hamashima, H. S. Yang, and E. Hendriks, “On the Utility of Spectral-Maximum-Based Automated Thread Counting from X-Rays of Paintings on Canvas,” *Studies in Conservation*, vol. 56, pp. 104-114, 2011.
- [10] D. H. Johnson, C. R. Johnson, Jr., A. G. Klein, W. A. Sethares, H. Lee, and, E. Hendriks, “A Thread Counting Algorithm for Art Forensics,” *IEEE DSP Workshop 2009*, Nov. 2009.
- [11] L. van Tilborgh, T. Meedendorp, E. Hendriks, D. H. Johnson, C. R. Johnson, Jr., and R. G. Erdmann, “Weave Matching and Dating of Van Gogh’s Paintings: An Interdisciplinary Approach,” *The Burlington Magazine*, vol. CLIV, pp. 112-122, February 2012.
- [12] W. Liedtke, C. R. Johnson, Jr., and D. H. Johnson, “Canvas Matches in Vermeer: A Case Study in the Computer Analysis of Fabric Supports,” *Metropolitan Museum Journal*, vol. 47, pp. 99-106, 2012.
- [13] D. H. Johnson, C. R. Johnson, Jr., and R. G. Erdmann, “Weave Analysis of Paintings on Canvas from Radiographs,” *Signal Processing (Special Issue on Image Processing for Art Investigation)*, vol. 93, pp. 527-540, March 2013.
- [14] C. R. Johnson, Jr., *Art Support Analysis Algorithms* Cornell University, ECE4210, Fall 2014.
- [15] D. Kushel, “Radiographic Methods Used in the Recording of Structure and Watermarks in Historic Paper” in *Fresh Woods and Pastures New: Seventeenth Century Dutch landscape Drawings from the Peck Collection*, pp. 116-123, Chapel Hill, 1999.
- [16] 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>