


# Towards a Network of Sixteenth Century Book Illustrations

Germaine Götzelmann   
Karlsruhe Institute of Technology (KIT)  
Karlsruhe, Germany

## Abstract

The following article investigates the reuse of book illustrations in German printed books of the sixteenth century. We utilize a workflow involving data and metadata retrieval, illustration segmentation, and reverse image search to detect identical and nearly identical illustrations in digitized books. The results of these steps are modeled as graph data, which enables us to filter, assess, and categorize them with the aid of graph measures and graph visualization. Drawing on a data sample of books published in Frankfurt am Main, we demonstrate how illustration reuse links the books in question in a complex and interconnected network, which can in turn be analysed and exploited to identify the specific patterns and characteristic features of this Early Modern cultural practice.

Following Johannes Gutenberg's invention of the movable-type printing press, a fundamental change took place in the use and reuse not only of texts, but also illustrations and ornaments within books. Previously, illustrations in manuscripts had to be redrawn and manually copied, whereas those in



*Creative Commons License Attribution 4.0 International (CC BY 4.0).*

In: Tara Andrews, Franziska Diehr, Thomas Efer, Andreas Kuczera and Joris van Zundert (eds.): Graph Technologies in the Humanities - Proceedings 2020, published at <http://ceur-ws.org>

This long paper is based on research presented at "Graph Technologies in the Humanities 2019" (January 18-19, Academy of Sciences and Literature | Mainz, Germany).

block books formed part of a fixed page layout that combined text and images in a permanent arrangement. With the advent of letterpress printing, illustrations could be placed in multiple positions and on multiple pages within the same publication. They could also easily be transferred to other books, which in many cases altered their context (and meaning) completely. While individual cases of re-purposed illustrations have been studied extensively, the wider phenomenon of illustration reuse in the sixteenth century remains elusive. This article seeks to showcase how a graph-based perspective can be employed in a data-driven approach to obtain a much-needed overview of this practice.

If we consider the research topic from the perspective of graphs or networks, we can untangle the following layers of networks:

1. Network(s) of printers and publishers
2. Network(s) of printing materials (wood blocks)
3. Network(s) of books
4. Network(s) of woodcut illustrations

Network 1 is a social, economic, and legal network based mainly on printers' guilds and permissions for setting up a print shop in a city. Printers and publishers were connected by family ties through marriage and inheritance, but also via joint ventures and the purchase of printing materials from other printers, occasionally in the form of bankruptcy assets.

Printing blocks were valuable items that were sometimes bought in large quantities and varieties (i.e. by Christian Egenolff in the aftermath of Heinrich Steiner's bankruptcy in Augsburg (Künast, 2013)). Printers focusing on the production of illustrated/ornamental books had a vested interest in assembling a large pool of printing blocks for reuse in various contexts. This circulation of printing materials constitutes network 2.

Books produced by a network of printers and publishers are interlinked in multiple ways, including re-prints (authorized or unauthorized), intertextuality, and compilation (network 3). As far as book illustrations are concerned, these multifarious entanglements engendered phenomena such as repeated use of the same printing block, recutting of wood blocks, and adaptation of certain elements and motifs from the said illustrations. In this article, we focus specifically on the reuse and recutting of wood block illustrations (network 4).

All of these networks are of potential interest to scholars working in the historical disciplines as well as in the fields of book and literary studies, as they can be linked to questions about the social, political, and economic conditions under which illustrated books were created and produced in a newly emerging context of printing, anonymous readers, book fairs, and 'mass pro-

duction.’ Typically, knowledge about such networks is either the result of tedious and time-consuming individual research and thereby limited in scope (e.g. to one specific topic, type of illustration or genre, printer or block-cutter, etc.),<sup>1</sup> or it is based on quantitative analysis and therefore concerned with the ‘big picture.’ This article assumes the latter approach. In order to analyse illustration reuse in a quantitative way, we have established a workflow that exploits metadata and image data for the creation of an illustration reuse graph. Our workflow consists of seven steps:

1. Metadata retrieval and corpus definition
2. Data retrieval
3. Image segmentation/illustration classification
4. Set-up of an image search engine
5. Creation of search result graphs
6. Categorization of search result graphs
7. Creation of illustration reuse graphs

VD 16, the register of printed works of the sixteenth century published in German-speaking countries,<sup>2</sup> is a collection of over 100,000 metadata records that provides access to a comprehensive and detailed national bibliography of German prints from the period in question. This extensive data pool enables data-driven approaches on a large scale and is the starting point for step 1 of our workflow. Over 68,000 of the records in question link to one or more digitization(s) of the respective books. The data is retrieved (step 2) through the collection of digitized books from library pages and especially via IIF APIs (Snydman et al., 2015). Step 3 addresses the task of segmenting every digitized book page in a dataset, so that illustrations can be identified and described by means of a selector shaped as a polygon or rectangle. While we use a very lightweight segmentation algorithm based on morphological image analysis with *OpenCV*, any segmentation/classification approach will work as long as it results in a collection of images and segmented regions that are either rectangular or can be translated into bounding boxes.<sup>3</sup> At the end of step 3 of the proposed workflow, each illustration (ideally) has been assigned a segmentation region by the illustration segmentation algorithm. For the illustration search, we use the open source VGG Image Search Engine (VISE) in version 2.0.1-beta 1.<sup>4</sup>

---

<sup>1</sup>As an example, see Bodo Gotzkowsky’s seminal studies of the uses and reuses of the works of block-cutter Hans Brosamer.

<sup>2</sup>[http://www.gateway-bayern.de/index\\_vd16.html](http://www.gateway-bayern.de/index_vd16.html)

<sup>3</sup>For the basic principles of semi-automatic Layout Analysis and Region EXtraction (LAREX), see Reul et al. (2017).

<sup>4</sup><https://gitlab.com/vgg/vise/-/tree/vise-2.x.y>. For a detailed description of the applica-

For the purposes of this article, we will mostly leave steps 1–4 aside, and focus instead on the graph-related steps 5–7. Steps 1 and 2 are briefly described in relation to a specific application in subsection 2.1. Step 5 begins with the task of querying every segmented illustration region in the image search engine and retrieving all matches for the said region of interest, which results in a number of query results that is identical with the number of images that have been fed into the search engine. These query results are then refined by removing noise and irrelevant results, and the data is put into the format of a search result graph. The procedures and results of this step are described in detail in section 1. While this step of the workflow already limits the result data, further manual work may still be necessary to bring the results in line with a specific research question. We outline this process with regard to sixteenth century book data in subsection 2.2. As a final step, we merge the collected results into a network of illustration reuse, fully enabling our dataset to be used for social network visualization and analysis (subsection 2.3).

## 1 Search Result Graphs

Search results in general are lists created by a search engine as an answer to a specific user input. A visual search is a type of search query that receives image information as input and outputs other image information as search results. Search results are usually ranked in descending relevance to the search query. In our visual search with VISE, the user input is a selected book page and a defined Region of Interest (ROI) on the said page, both of which are derived from the image segmentation in step 3 of our workflow (see Figure 1). The search result is a list of other book pages with matching image features. The ranking is based on the number of relevant matching key features in close proximity to each other (inliers) in a method called *bag of visual words*.<sup>5</sup>

Upon retrieving the complete results of a full  $n \times n$  image search for  $n$  input images, we are faced with two problems: first, the image search is a black box, which means we can never be exactly sure how well it is dealing with a specific dataset until we have a look at the results; and second, while the ranking ensures that the most relevant search results come first, it does not prevent the return of potential matches that have virtually no resemblance to the input ROI. A fixed value score threshold for the number of feature inliers can be used, but given that the number of inliers is highly dependent on the number of key features in both the query image and in the matching image, this is not really a feasible solution for achieving a sharp limitation of results. Were we to use those search results without further filtering, we would either

---

tion of VISE on incunabula data from the fifteenth century, see also Zisserman et al. (2020).

<sup>5</sup>For technical details regarding this approach, see Arandjelovic (2013).

lose a significant amount of valuable results by being too restrictive, or we would end up with a large number of book pages that technically match one another, while in fact having no meaningful connection at all. Therefore, it is necessary to limit the result sets provided by the image search according to various parameters. The main parameters that can be used for this operation are relative inlier score, transformation matrix properties, and the Jaccard coefficient between search result and segmented region.

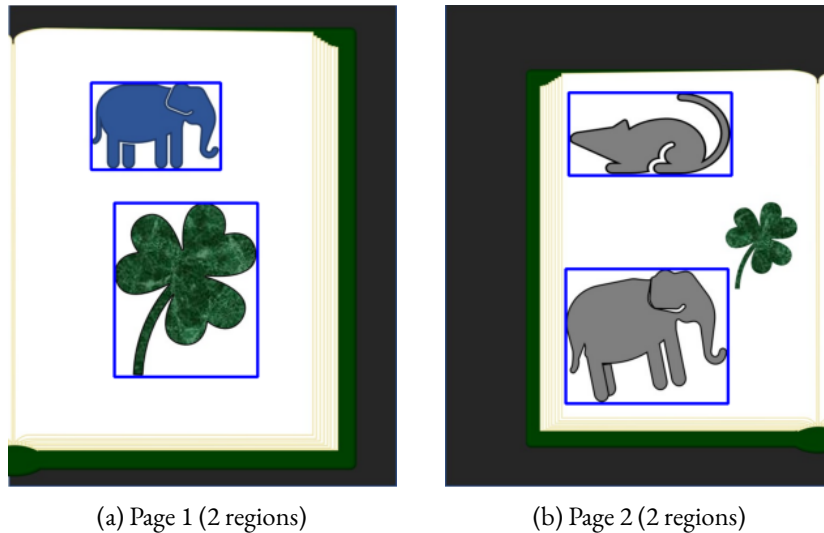


Figure 1: Example pages with segmented illustration regions after workflow step 3

## 1.1 Result Filtering

### 1.1.1 Relative Score Threshold

The relative inlier score takes into account the maximum possible number of matching key features that can be found in a given query region. This is easy to determine, since every query region finds itself as the best match. Based on this maximum score, we can compute the ratio between current score and maximum score for every search result. Empirically, cutting off the results at a threshold of about 0.02 for that ratio removes a reasonable amount of irrelevant matches (for details, see Table 1), making the score parameter the most simple and robust cutting parameter.

### 1.1.2 Transformation Matrix Thresholds

The relation between a query region and a result region is given by the values of a  $3 \times 3$  transformation matrix. It describes the transformation of a given rectangle into a polygon by translation in x and y dimension, scaling in x and y dimension, rotation, and skewing. Translation only describes a shift

between image positions and cannot be used for cutting results. While the digitized data of book pages can be quite heterogeneous and of varying quality, we can assume that we always have a picture of a more or less flat page positioned parallel to the image frame, with the camera or scanner pointing directly at it. While warped pages or slightly rotated books do effect the outcome, we can rule out the rotation of the camera as a factor in most of the illustrations that concern us. This vastly limits the possibilities of plausible skewing between the query region and the result region, and enables us to assess the decomposed transformation matrix in itself without the need to account for the specifics of a given query: skewing is limited to values close to 0. Rotation is expected in ranges around 0,  $(-)$ 90 and  $(-)$ 180 degrees, taking into account that a woodcut might have been printed on a book page in a vertical position instead of a horizontal one, or that a digitized page might not have been rotated properly. Scaling is unlimited, but the ratio between x and y scaling is not. If the scaling ratio differs too much from a 1:1 ratio, the transformation must be considered as distorted, eliminating it from the list of possible results.

### 1.1.3 Jaccard Coefficient Threshold

While the query regions are defined as bounding box rectangles for an illustration segment, the resulting regions of the image search are parallelogram regions. For a concise result of the whole image search process, we need to match the resulting regions with the ones we have provided as query regions. For that purpose, we calculate the intersection over union value (IoU), also called Jaccard coefficient.<sup>6</sup> This matching is necessary if there are any pages in the corpus that contain more than one illustration region.

Tests have shown that once one of the applied filter parameters indicates an invalid match, it is very unlikely for results with a lower search result rank to be valid. Therefore, the list of results can be cut once an invalid parameter threshold is met.

## 1.2 Graph Creation

Graph creation can be explained by means of two sample pages to which the image search workflow is applied. On sample page 1 (p1), two illustrations were segmented successfully (indicated by blue bounding boxes in Figure 1):

---

<sup>6</sup>The Jaccard similarity coefficient is often used in machine learning to determine the quality of classification – especially object detection – compared to a ground truth. Given a classified region and a region for comparison, the similarity between the two regions is computed by calculating the intersection of both areas and dividing it by the union of both areas. Identical regions result in a value of 1, disjoint regions in a value of 0.

one depicting a blue elephant ( $p_1r_1$ ), the other a large three-leaf clover ( $p_1r_2$ ). Likewise, on sample page 2 (p2), two illustrations were segmented: a gray mouse ( $p_2r_1$ ), and a gray elephant ( $p_2r_2$ ) that is similar to the blue elephant on p1, except for minor changes to its contours and a slight rotation to the left. A third illustration, depicting a three-leaf clover identical in shape and color to the one on page one was not segmented, because the leaf was too small (due to a size threshold for illustration areas used in the segmentation step).

When the image search is applied to the four segmented illustration regions, the following occurs:

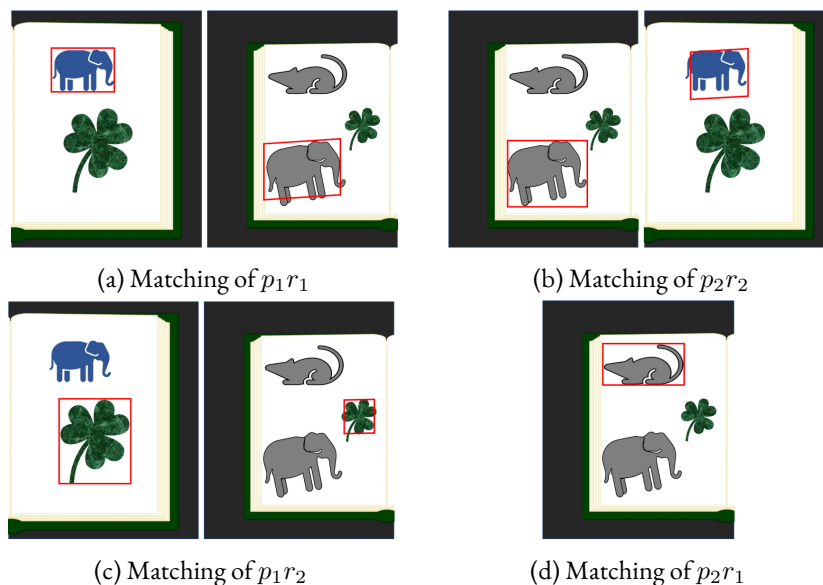


Figure 2: Image search matches for page 1 and page 2

- $p_1r_1$  (blue elephant) finds a match on p2 (Figure 2a). This match has a Jaccard similarity of 0 with  $p_2r_1$  (no intersection) and a Jaccard similarity close to 1 with  $p_2r_2$ . Therefore,  $p_1r_1$  and  $p_2r_2$  are considered a match.
- $p_2r_2$  (gray elephant), on the other hand, finds a match on p1 (Figure 2b). This match has a Jaccard similarity of 0 with  $p_1r_2$  (no intersection) and a Jaccard similarity close to 1 with  $p_1r_1$ . Therefore,  $p_2r_2$  and  $p_1r_1$  are considered a match, making the match reciprocal.
- $p_1r_2$  (three-leaf clover) finds a match on p2 (Figure 2c). This match has a Jaccard similarity of 0 with  $p_2r_1$  and a Jaccard similarity of 0 with  $p_2r_2$ . Therefore, no valid match is detected.
- $p_2r_1$  (mouse) does not find match on any of the available pages (Figure 2d).

It is possible to create a search result graph directly from the above results. This search result graph is a directed graph where the nodes are segmented illustration regions and the edges connect two segmented regions if they are considered a match. Such a graph can be drawn by following a few general rules (the resulting graph can be found in Figure 3):

1. A segmented region  $p_n r_m$  finds an image search match on  $p_o$ .  $p_o$  contains  $k$  segmented regions ( $k > 0$ ). The Jaccard similarity to one or more of those  $k$  regions is above an IoU threshold (min 0). Region  $p_o r_p$  with the best IoU value is considered a match. A directed edge  $(p_n r_m, p_o r_p)$  is drawn.
2. A segmented region  $p_n r_m$  finds an image search match on  $p_o$ .  $p_o$  contains  $k$  segmented regions ( $k \geq 0$ ). No Jaccard similarity to those  $k$  regions is above an IoU threshold (min 0). No edge is drawn.
3. Optional: All non-reciprocal edges are removed.<sup>7</sup>




Figure 3: Graph result of matches on page 1 and page 2

### 1.3 Visual Search as a Graph Problem

Graph creation and result filtering can be done in any order. Both approaches have their own distinct advantages. Filtering first reduces the amount of data to be processed in the graph creation step tremendously.<sup>8</sup> For large real world datasets with over 50,000 images, filtering reduced the number of results to be processed by over 90% (Table 1). This approach is

<sup>7</sup>Keeping only reciprocal edges simplifies calculation of graph measures, because the graph can then be treated as undirected. It also has the potential to significantly enhance the precision of the result, but may in turn have a detrimental effect on the recall due to valid results being discarded. Experiments have shown that it is not unlikely for valuable matches to be discarded if they are borderline cases for the image search (coloring, skewing, etc.), so this option should be handled with care.

<sup>8</sup>With the exception of filtering by Jaccard coefficient threshold, since calculation of the intersection over union value is mandatory in the graph creation step.



especially advisable if speed is of the essence, or if system resources are limited. Creating the graph first, before refining the result by filtering it with the same thresholds and by removing any edges from the graph, slows down the process. However, it also makes full use of the graph approach for the assessment of the filter threshold parametrization. The characteristics of the created graph provide valuable information on the quality of the search results that cannot be taken into account when filtering the search result data directly.

The ideal search result graph consists of  $n$  fully connected components. A fully connected component is a subgraph where each node is connected by an edge with every other node within it. A fully connected subgraph with  $k$  nodes perfectly describes one matching illustration with  $k$  corresponding segmentation regions on up to  $k$  book pages (see Figure 4). In reality, the result usually differs to a certain degree from the ideal graph form due to errors both in image segmentation and image search. In theory, we would like to calculate the difference between the ideal graph and the real graph as a measure of quality, but besides specified ground-truth data, we know neither the correct number of subgraphs, nor their individual ideal sizes for a given result graph. When it comes to evaluating real world data, however, we can still assess some graph measures in order to estimate whether the graph is reasonably close to an ideal graph structure, even though such proximity does not in itself guarantee an ideal or even passable result.




Figure 4: Fully connected subgraph with 16 matching illustrations

### 1.3.1 Assessment by Graph Measures

Graph measures for quality assessment must be divided into global measures calculated for the whole graph and local measures calculated for individual subgraphs or nodes. Given the sheer size of search result graphs that contain millions of nodes, global graph measures tend to be difficult to calculate efficiently. Some basic filtering is mandatory before starting to look into the finer grained assessment, because without it, it might not even be possible to partition the overall graph into subgraphs. Once basic filtering is done, it becomes significantly easier to apply local graph measures on the subgraphs.

#### Global Measures

In the following section, we demonstrate the global assessment of a basic graph filtering sequence using a string of (conservatively chosen) thresholds. We start with an unfiltered graph and apply thresholds in sequence, so every new filter is applied to the output of the previous filtering. Initially, a relative score filter of 0.01 is applied to the unfiltered graph due to the high number of nodes it contains. This is done to circumvent excessive memory consumption in generic graph analysis tools like *Gephi*. Whenever the edge filtering leaves isolated nodes with no further connections, those nodes are also removed. The following filter chain is then applied in sequence: filtering edges by relative score threshold, then by rotation angle diff mod 90, followed by ratio threshold of scaleX and scaleY, and finally by IoU threshold of the matching regions. For the filtered graphs, we calculate the following global graph measures: number of nodes, number of edges, number of connected components, and the global avg. clustering coefficient for directed graphs<sup>9</sup>. The results can be found in Table 1.

The clustering coefficient of a graph is a measure of how densely the subgroups of a graph are connected. The closer a graph is to a graph consisting of fully connected components as subgraphs, the higher the clustering coefficient. Therefore, the avg. clustering coefficient of a graph is a good global measure with which to assess the distance to an ideal graph structure. Since the image search produces a lot of ‘noisy’ and meaningless results, we expect to be able to filter a significant amount of edges from the initial graph. The number of filtered nodes is dependent on the specific dataset. Ideally, if every segmented illustration had at least one match in the dataset, no nodes would have to be removed in the filtering process. In reality, however, we encounter a varying number of segmented illustrations which are unique within the dataset. These are created in the initial, unfiltered graph and are

---

<sup>9</sup>We use Patrick McSweeney’s implementation in *Gephi*. For details on the implemented algorithm, see <https://github.com/gephi/gephi/wiki/Average-Clustering-Coefficient>.

subsequently removed. The number of filtered nodes therefore can only be assessed with a deeper knowledge of the dataset. Table 1 shows that the number of edges is reduced tremendously over the course of filtering. Over 90% are filtered on the score threshold, while the avg. clustering coefficient also increases. This shows that the relative score filter is crucial to the creation of a meaningful graph. As expected, the number of connected components increases as the edge filtering successively breaks up larger subgraphs into smaller, more precise subgraphs.

	relScore >0.01	relScore >0.025	rotation <10°	scaleRatio >0.75	IoU >0.6
Nodes	59,616	56,901	56,774	56,774	54,776
Edges	1,701,595	790,527	777,313	747,893	672,182
Conn. Comp.	635	5,529	5,732	6,645	8,970
Cluster. Coeff.	0.355	0.654	0.659	0.667	0.719
Edge filter %	0.8397	0.9255	0.9267	0.9295	0.9366
Edge filter per step %	0.8397	0.9255	0.9267	0.9295	0.9366
Node filter %		0.0455	0.0477	0.0477	0.0812

Table 1: Graph filtering chain (initial graph size: 59,616 nodes and 10,610,440 edges)

### Local Measures

Iterating over the subgraphs of the search result graph, we can supplement those global measures with more detailed local measures. In practice, even if the graph is not completely broken up into perfect subgraphs, individual components never contain more than a few thousand nodes and, on average, considerably less. We can identify fully connected components by two local measures: either by their density, or by the size of their maximum clique. The maximum clique of a graph is the largest fully connected subgraph of a given graph, so if we want to have an ideal search result subgraph of  $n$  nodes, the maximum clique size is also  $n$ . The clique problem is an np-complete problem and therefore impossible to solve in polynomial time on arbitrary graphs, but it can nevertheless be computed for our specific graph shape of small subgraphs. Calculating the fraction of the maximum clique size and the number of nodes in a subgraph gives us a measure of closeness between the subgraph and an ideal – i.e. fully connected – one. A similar measure is provided by graph density. Graph density calculates the ratio of edges in a graph with respect to the number of edges possible. Due to the complexity of this operation, some implementations of the maximum clique size are

only provided for undirected graphs<sup>10</sup>, which is why we must treat our graph as such. The simplified result might not always represent the true shape of a subgraph, but using a simpler density measure means that it can be computed for either directed or undirected graphs.

A third measure of interest is the calculation of a minimum dominating set in a subgraph. A dominating set of nodes in a graph is a set of nodes from which all other nodes in the graph can be reached by traversing a single edge. In a minimum dominating set, the number of nodes in that set is as small as possible for the given graph. The set itself does not have to be unambiguous – multiple minimum dominating sets may exist in one graph. In an ideal, fully connected graph, the size of the minimum dominating set is one, and every node of the graph qualifies as a dominating set, because every other node can be reached directly. Again, the minimum dominating set size can be computed for both the directed or undirected variant of our subgraphs, providing different measures. Density and minimum dominating set size are not directly correlated, since even a very sparse graph can have a close to ideal dominating set size (if we have a graph of  $n$  nodes (with a large  $n$ ), where one node is connected to all other nodes, which in turn are not connected among each other at all, we have a sparse graph and a minimum dominating set size of 1). The dominating set size is especially interesting as a complementary graph measure because it gives us a much more detailed idea of a specific graph shape. It is also useful for identifying the most ‘characteristic’ nodes in a search result, since it contains the nodes that are most important for a graph’s connectivity.

### 1.3.2 Visual Assessment

In addition to allowing the calculation of graph measures for quality assessment, arranging the search results in a graph structure enables the user of the workflow to employ generic graph visualization tools for visual assessment. It is a fast, flexible, and easy way to identify unusual graph characteristics (subgraphs with an extraordinarily large number of nodes, or subgraphs that differ from the ideal fully connected component structure). Moreover, metadata information can be used for graph partitioning/node coloring, so that, for example, nodes can be colored by book, by place of publication, by printer, etc. This can help to identify underlying problems and may be useful for understanding deviations from the ‘ideal’ structure explained above. In Figure 5 a subgraph is shown with a (directed) density of 0.384. Differing significantly from a fully connected component structure, the graph is separ-

---

<sup>10</sup>See, for example, the Python module networkx: [https://networkx.github.io/documentation/stable/reference/algorithms/generated/networkx.algorithms.clique.graph\\_clique\\_number.html](https://networkx.github.io/documentation/stable/reference/algorithms/generated/networkx.algorithms.clique.graph_clique_number.html).

ated quite clearly into two strongly interconnected areas with relatively weak links between the two. We examine two nodes from the minimum dominating set in the graph, one from the left area and one from the right. The segmented illustration regions can be found in Figure 6. Both feature the same image of a man boarding a ship, so they are overall correctly placed in the same subgraph. Upon closer examination, however, it becomes clear that the prints are not identical: one woodcut is a recut of the other and exhibits slight variations in details like the facial features of the figures.<sup>11</sup> While the recut is similar enough to the original to trigger a match in the image similarity search, the differences are visible in the shape of the resulting subgraph, which in turns divides the graph into two quite separate clusters. This is a case where the graph structure clearly reflects an important aspect of historic production processes. Outliers such as these cannot satisfactorily be evaluated by simple graph measures, but are quick and easy to assess with the help of visual analysis.

## 2 Application to Sixteenth Century Book Illustrations

The workflow described above can be used to research networks of sixteenth century books and their illustrations. In this section, we will define a sample dataset and further elaborate on its characteristics as well as workflow steps 6 and 7, which lead us from a search result graph to a network of book illustrations.

### 2.1 Fine-Tuning the Dataset: Illustrated Books from Frankfurt am Main

Taking the VD16 register as our starting point, we can use the bibliographical metadata to select a specific dataset to which our workflow can then be applied. In addition to bibliographical data concerning a work’s author, title, printer, or place and year of publication, the VD16 also collects more material-oriented information, such as book size and ornamentation (“Buchschmuck”). The data concerning ornamentation is organized as follows: TE (Titeleinfassung – title frame), TH (Titelholzschnitt – title woodcut), TK (Titelkupferstich – engraving on the title page) for ornamental elements on the title page; H (Holzschnitt – woodcut), K (Kupferstich – engraving)

---

<sup>11</sup>See, for example, the man on the bow of the ship (especially his hat and face) and the beards/faces of the men on the right, as well as other details.

<sup>12</sup>Digitized by Staatsbibliothek Berlin, see <http://resolver.staatsbibliothek-berlin.de/SBB0001C26800000000>.

<sup>13</sup>Digitized by Bayerische Staatsbibliothek, see <http://mdz-mbn-resolving.de/urn:nbn:de:bvb:12-bsb00070159-3>.




Figure 5: Subgraph with 2 separated parts. Nodes of a minimum dominating set in red.



(a) Woodcut on VD16 H 2662, p. 113<sup>12</sup>



(b) Woodcut on VD16 H 3866, p. 62<sup>13</sup>

Figure 6: Recut images of a man boarding a ship

and RL (Randleiste – ornamental frame) for decorative elements within the book; and D (Druckermarke – printer’s mark) for the insignia of printers and publishers. While engravings do not play a significant role as an illustration technique in the sixteenth century, over 26,000 books are tagged with H for woodcuts and over 23,000 with TH for title woodcuts. This information plays a critical role in our data selection, since it allows us to limit our focus on books containing woodcuts or title woodcuts. Frankfurt am Main, which became important to the printed book trade following the establishment of Christian Egenolff’s print shop in the 1530s (Kulturvereinigung Hadamar, 2002), is among the top 10 printing centers listed in the VD16. It is in fact the city where the most books featuring woodcuts were printed (close to 2,600), followed by Wittenberg and Strasbourg. This is especially impressive since Frankfurt only turned into a major hub of the emerging printing industry in the second half of the sixteenth century, while Wittenberg and Strasbourg had played an important role more or less from the beginning of the century. Frankfurt’s prolific output and its many active printers and printing dynasties make the city an ideal candidate to study the reuse of illustrations in a quantitative way. The VD16 contains 1886 entries with ‘Frankfurt/Main’ as their place of publication which are both labelled as containing woodcuts and linked with a digitized reproduction, thereby enabling them to be used in an image search workflow.

Since nearly a quarter of the books with woodcuts from Frankfurt are not linked to a digitized version in VD16, our application does not claim to provide a complete result for the corpus in question. The main goal of the following example is to show if and how the results of network creation are meaningful and coherent. At this point, it is to be understood as a study of technical feasibility rather than as an accurate visualization of historical practices. It does, however, provide a basis for further inquiries along the same lines. We will present the results of our approach for around one thousand digitized books in subsection 2.3, which yielded around 53,700 pages that were segmented as containing one or more illustration region(s) in workflow step 3, and were subsequently ingested into a VISE search engine in workflow step 4. The resulting search result graph consists of around 9,100 subgraphs.

## 2.2 Fine-Tuning the Definition: Woodcuts vs. Illustrations

It became clear both during data retrieval and after initial use of the image search that the VD16 metadata is not detailed enough to meet the requirements of our specific research question. In VD16, the label ‘H’ is assigned to all non-textual elements in a book outside of the title page that are produced

by inserting a wood block into the page layout for printing.<sup>14</sup> This includes all kinds of decorative elements such as ornaments and initials, as well as images, such as botanical illustrations, mathematical figures, etc. While this categorization appears quite clear-cut, the term ‘illustration’ itself is actually rather fuzzy. Narrowly defined, illustrations are

all pictures in a book that are meant to accompany a literary text to clarify content and plot for the reader, to strengthen the text’s significance and impact. Portraits, factual drawings [Sachzeichnungen], depictions of cityscapes, technical representations, botanical images etc. are to be called figures. (Wendland, 1991)<sup>15</sup>

Our goal is to find a middle ground. We are interested in all images that support, clarify or accompany a text. Whenever such an image is reused, some kind of knowledge transfer takes place, which is precisely what makes the reuse and recontextualization of such materials so interesting. Yet our experience has shown that if initials and ornamentation are included, the majority of books printed in Frankfurt are somehow connected with one another, making the resulting network incredibly large, noisy, and outright unusable for further analysis. Therefore, the decision was made to leave purely decorative elements as well as printer’s marks aside, while botanical and medical illustrations, cityscapes, and other explanatory figures are included. There are, of course, borderline cases such as historiated or inhabited initials (i.e. decorative letters that contain identifiable figures or scenes from a story).

However, regardless of the specific definition of illustration being used, an image similarity search cannot be aligned with such fuzzy, *a priori* criteria with acceptable precision. It is a clear advantage of the graph approach that the visual assessment described in subsection 1.3.2 is not only useful to assess the graph quality, but can also be utilized to further specify the data of interest for a given research question. To do so, a lightweight graphical user interface has been implemented to show the subgraphs of the created search result graph in a tabular view (see Figure 7). This enables a researcher to quickly assign tags to subgraphs such as ‘initial,’ ‘printer’s mark,’ ‘decoration,’ etc. Multiple tags chosen from a simple predefined list of terms (to limit erroneous user input) can be assigned to one subgraph. While this system is currently only used to define the subset of data for a specific research interest, it can of course be expanded to meet the demands of various annotation tasks. The list of terms can easily be expanded into more com-

---

<sup>14</sup>The only exception are printer’s marks, which are labelled separately.

<sup>15</sup>Author’s translation.



plex controlled vocabularies for the purposes of classification and analysis, e.g. through the addition of *Icon Class*<sup>16</sup> information to illustrations.

Images (first 6)	↓ Bookcount	Density	Max Clique %	Dominating Set %	Tags
	22	0.465	0.405	0.19	EDIT
	22	0.947	0.853	0.029	• signete EDIT
	22	0.947	0.893	0.036	• decoration EDIT
	22	0.801	0.704	0.074	EDIT
	22	1	1	0.038	EDIT

Rows per page: 5 36-40 of 9168 < PREV 1 7 8 9 1834 NEXT >

Figure 7: Graphical user interface for manual tagging of subgraphs

This approach allowed 77 different printer’s marks from our dataset to be identified,<sup>17</sup> and 185 subgraphs to be tagged as initials.<sup>18</sup> Combining automatic segmentation and search with manual categorization narrows the figure of several thousand book pages down to a manageable number of subgraphs that can then be used to accommodate more finely tuned research requirements. This versatility makes our approach suitable for a wide range of research topics: not only does it allow for manual categorization on large datasets, but it also makes it possible to conduct complex algorithmic assessment or analysis on one representative image per subgraph instead of all images in the larger dataset. To ensure representativity, graph measures can again be utilized, for example, by choosing one image per subgraph from the respective minimum dominating set.

### 2.3 Fine-Tuning the Representation: Illustration Reuse Graph

Following the creation and optional categorization of the search result graph, the next step is to transform the subgraphs that identify identical illustra-

<sup>16</sup><http://www.iconclass.nl/home>

<sup>17</sup>The identification of common printer’s marks was greatly aided by the collection of Wendland (1984).

<sup>18</sup>The total number of initials would be significantly higher, but the workflow frequently (and erroneously) puts different initials in the same subgraphs due to their high structural similarity and their low number of distinct visual features.

tions into a network of books connected by identical illustrations. The result of this step is a graph in which each node represents one book in our dataset. An edge is drawn between node a and node b if they share at least one illustration. The edge is weighted by the number of shared illustrations between the books, so that if one illustration is found in book a and book b, the edge weight is 1, whereas if 42 illustrations are found in both book a and b, the edge weight is 42. The input that is required for such a node merging process to transform a search result graph into a network of illustration reuse is shown in Figure 8. Meanwhile, the merged output appears in Figure 9. Isolated book nodes with no connections to other books are removed after merging.




Figure 8: Example search result graph, nodes colored by book




Figure 9: Example of a merged book graph with weighted edges

It is possible to apply graph layout algorithms on the created network that take into account the weighted edges between the books. Thus, two books are attracted to each other if they are connected with a (highly) weighted edge and repel each other if no edge exists between them. This way, books with similar illustrations cluster together in subgroups. Bibliographical metadata from VD16 can then be used to color nodes and edges by partitioning. For example, book nodes can be colored according to the respective printers involved, in order to show how illustrations circulated within one print shop, or were transferred to other businesses. Edges can be colored according to the difference of publication dates between books, ranging from a ‘hot’ color for illustrations that are reused within a short period of time (within a few

years or even in the same year) to a ‘cold’ color for those that are reused over a longer temporal duration. Figure 10 shows an illustration reuse network for Frankfurt am Main before and after application of a graph layout and coloring of nodes and edges by metadata. It contains 735 books connected by at least one shared illustration. The data has been processed and manually categorized according to the steps described above. The network depicts a high amount of illustration reuse both within and between particular print shops. Node degrees (and therefore illustration reuses involving individual books) range from 1 to 148 (the latter being VD16 P 3553, a *Naturalis historia*), the time span between the first and the last use of the same illustration varies between 0 and 67 years (the latter showing the reuse of an illustration of medical blood-letting). The type of clustering that occurs here is indicative of a ‘functional’ use of illustrations: it is hardly surprising that botanical or medical books are far more likely to share illustrations with each other than literary texts or religious treatises. It is still evident, however, that the functional clusters are not completely separated, but remain connected by singular instances of reuse that transcend categorizations by functionality or ‘genre.’ Clearly, further exploring the illustrations in question and the books (re)using them would be a worthwhile pursuit, and herein lies one of the main advantages of a quantitative approach like the one proposed here: far from being limited to showing the connections between books by one specific printer, dealing with one topic, or belonging to one genre, it is capable of detecting weaker links between the clusters which are much harder to discover by traditional means.

We will now examine one of the subgroups in which literary works such as *Fortunatus*, *Herzog Ernst*, Thüring von Ringoltingen’s *Melusine*, Elisabeth von Nassau-Saarbrücken’s *Hug Schapler*, and other well-known and widely disseminated narrative texts are clustered together. Historically, this book cluster began to emerge with the activities of Weigand Han in Frankfurt, who illustrated multiple print runs of *Fortunatus* with a collection of woodcuts produced by Hans Brosamer, which were then quickly reused in conjunction with other narrative texts in the 1550s. The network as it stood at the end of the first half of the sixteenth century is represented in Figure 11a. At this point, significant reuse of the illustrations in question had already taken place, albeit only within the ‘genre’ of entertaining narrative texts, and only in Weigand Han’s own print shop (Han mostly printed under his own name, but on one occasion used the imprint of ‘heirs of Hermann Gölferich’). In the second half of the sixteenth century, however, the reuse of illustrations became more diverse, and went beyond the group of narrative texts for which the images had initially been created. For example, a generic

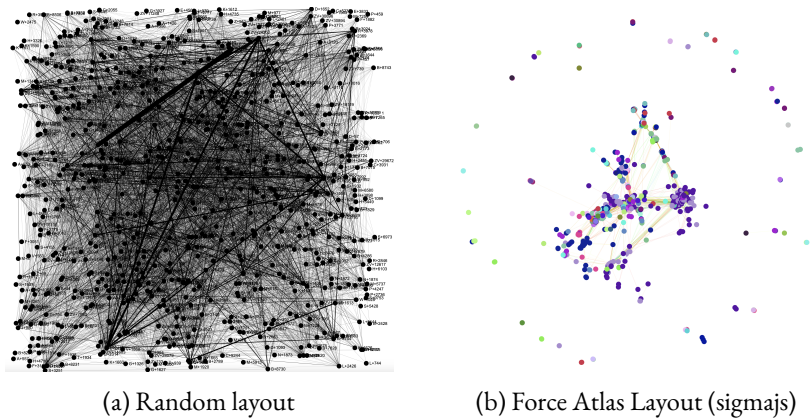


Figure 10: Illustration reuse network before and after applying a graph layout and coloring by metadata

illustration of a griffin saw frequent reuse as a title woodcut for a group of cautionary ‘devil’s books’ (‘Teufelsbücher’). Generic social scenes, such as an image of people sitting at a table talking, or of two knights riding in a tournament, were also recontextualized, as were images of torture and execution. Entry J 627 in the VD 16 register plays an extraordinary role in linking the cluster of narrative texts to another group of books. Printed in 1580 by Weigand Han’s son, Hartmann Han, it is a German version of John Mandeville’s *Travels*, a fictional travelogue detailing the author figure’s journey through the Middle and Far East. The text exhibits a fascinating mixture of literary story-telling and biblical references, and the book’s illustrations follow the same pattern, with pictures that typically accompanied literary texts being intermingled with images otherwise used to illustrate religious texts.

Our case study shows that all four network types outlined in the beginning play a significant role in the reuse of book illustrations in the sixteenth century; a practice characterized by a multitude of social, economic, and material connections between the various actors and publications involved. Using the approach outlined here, all of these links can be explored simultaneously.

### 3 Conclusion

As we have demonstrated, graph technologies can be employed at every step of the way from ranked visual search results to a fully-fledged network of illustration reuse. All individual parts of the data workflow can be assessed and analysed by means of graph measures and graph visualization. Our method allows us to apply a quantitative approach to large-scale datasets, while at the same time enabling us to take full advantage of the benefits of visual search.

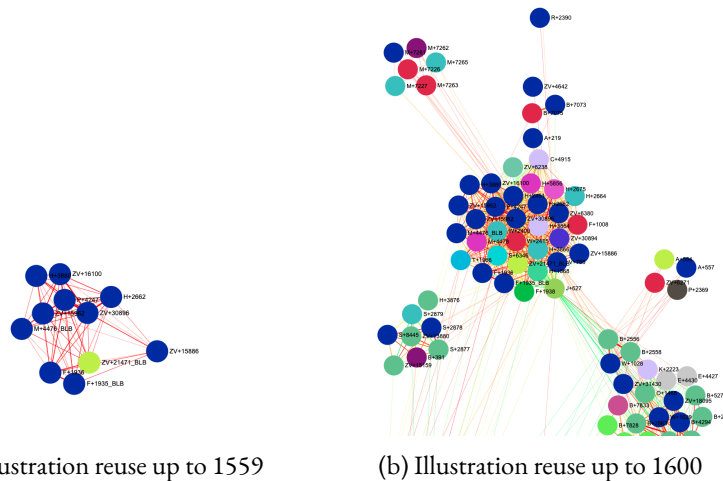


Figure 11: Illustration reuse network of literary texts (partial view of Figure 10b)

The application of our workflow to digitized books from the sixteenth century and their bibliographical metadata has shown that it delivers viable and useful results with a high degree of versatility and transparency. Even if some of the links between books and printers from this time period are bound to go undetected as long as gaps remain in the digitized data, our approach is an expedient step towards a more complete picture of the historical practice of illustration reuse. The method we propose allows researchers to tap into a wealth of cultural heritage data that has yet to be explored with the aid of big data, and helps bridge the gap between generic image similarity search and highly individualized research interests.

## References

Arandjelovic, R. (2013). *Advancing Large Scale Object Retrieval*. PhD thesis, University of Oxford, <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.640168>.

Kulturvereinigung Hadamar, editor (2002). *Christian Egenolff 1502–1555: Ein Frankfurter Meister des frühen Buchdrucks aus Hadamar*. Glaukos, Limburg.

Künast, H.-J. (2013). »*Getruickt zu Augspurg*«. De Gruyter, Berlin/Boston, MA, DOI: 10.1515/9783110935080.

Reul, C., Springmann, U., and Puppe, F. (2017). LAREX – A semi-automatic open-source Tool for Layout Analysis and Region Extraction on Early Printed Books. *arXiv*, 1701.07396, <https://arxiv.org/abs/1701.07396>.

- Snydman, S., Sanderson, R., and Cramer, T. (2015). The International Image Interoperability Framework (IIIF): A Community & Technology Approach for Web-Based Images. In *Archiving 2015: Final Program and Proceedings*, pages 16–21.
- Wendland, H. (1984). *Signete. Deutsche Drucker- und Verlegerzeichen 1457–1600*. Schlütersche Verlagsanstalt und Druckerei, Hannover.
- Wendland, H. (1991). Illustration. In *Lexikon des gesamten Buchwesens*, volume 3. Hiersemann, Stuttgart.
- Zisserman, A., Dutta, A., Malaspina, M., and Dondi, C. (2020). The Use and Reuse of Printed Illustrations in 15th-Century Venetian Editions. In Dondi, C., editor, *Printing R-Evolution and Society 1450-1500*, pages 839–869. Edizioni Ca' Foscari, DOI: 10.30687/978-88-6969-332-8/030.