

Contrastive Attention Networks for Attribution of Early Modern Print

Nikolai Vogler¹, Kartik Goyal², Kishore PV Reddy¹, Elizaveta Pertseva¹, Samuel V. Lemley³,
Christopher N. Warren³, Max G'Sell³, Taylor Berg-Kirkpatrick¹

¹ University of California, San Diego

² Toyota Technological Institute at Chicago

³ Carnegie Mellon University

nvogler@ucsd.edu, kartikgo@ttic.edu, cnwarren@cmu.edu, tberg@ucsd.edu

Abstract

In this paper, we develop machine learning techniques to identify unknown printers in early modern (c. 1500–1800) English printed books. Specifically, we focus on matching uniquely damaged character type-imprints in anonymously printed books to works with known printers in order to provide evidence of their origins. Until now, this work has been limited to manual investigations by analytical bibliographers. We present a Contrastive Attention-based Metric Learning approach to identify similar damage across character image pairs, which is sensitive to very subtle differences in glyph shapes, yet robust to various confounding sources of noise associated with digitized historical books. To overcome the scarce amount of supervised data, we design a random data synthesis procedure that aims to simulate bends, fractures, and inking variations induced by the early printing process. Our method successfully improves downstream damaged type-imprint matching among printed works from this period, as validated by in-domain human experts. The results of our approach on two important philosophical works from the Early Modern period demonstrate potential to extend the extant historical research about the origins and content of these books.

Introduction

A complete understanding of the content of surviving historical works requires knowledge of the context of their printing. This has long been acknowledged in the humanities, giving rise to *analytical bibliography*, where physical evidence from printed books is used to understand their means of production, revealing incentives for their publication as well as cultural diffusion of language and knowledge. One of the main kinds of such physical evidence is character type-imprints pressed from damaged type pieces. These anomalous type pieces, when linked across printed works, can reveal which works came from the same printing apparatus, as depicted in Figure 2. Historians have, for example, leveraged the analysis of damaged type-imprints as a critical tool in identifying and studying clandestine printers who withheld their identities to avoid harsh penalties under censorship laws (Adams 2010; Como 2018; Warren et al. 2020, 2021).

In this paper, we develop a computational model of the same analytic techniques used by bibliographers. Our method

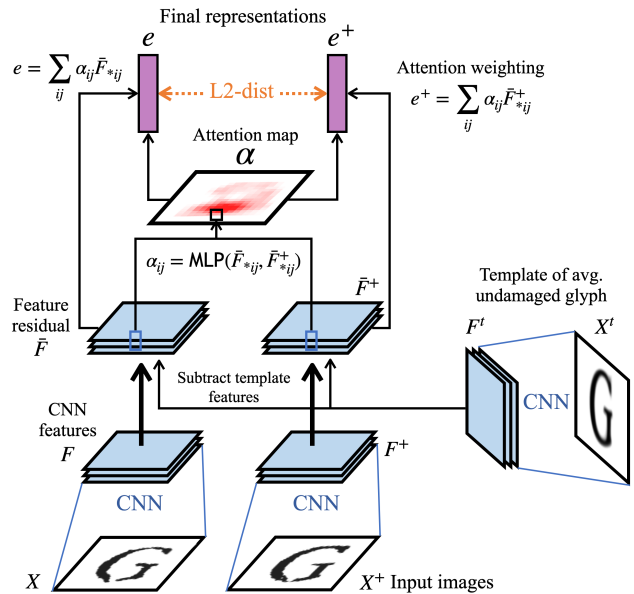


Figure 1: A depiction of the neural architecture behind Contrastive Attention-based Metric Learning (CAML), which we train to identify pairs of damaged type-imprints in early modern books that have similar damage and therefore likely originate from the same damaged piece of metallic type and printing apparatus. Our model uses a joint attention mechanism to attend to residual convolutional features after a template feature map has been subtracted from both inputs.

automatically compares type-imprint extractions across individual books, or larger collections of books, identifying pairs of *damaged* type-imprints that match one another and therefore likely originate from the same, damaged piece of metallic type. Our approach enables *computational* bibliographical analysis that has the potential to scale far beyond the limits of manual analyses – for example, to the tens of thousands of censored early modern works that are of interest to social, cultural, and intellectual historians.

Matching damaged type-imprints across texts is challenging, since the variation in scale, position, font, and inking in scans of printed characters generally dominates the minute deviations due to localized damage that are of actual interest, as shown in Fig. 3. We present a **Contrastive Attention-based**

Metric Learning approach to identify similar damage across type-imprint pairs. Due to the dearth of training data consisting of matching damaged types, we design a specialized data generation process informed by human experts that aims to simulate bends, fractures, and inking variations induced by the early printing process. This lets us produce synthetic training data of image pairs with realistic matching damages. Then we develop a convolutional neural architecture that uses an attention mechanism for localized comparison of character pairs. By fitting this model to our synthetic data, we learn a metric appropriate for identifying similar, damaged type-imprints that generalizes to real data. Our model is sensitive to very subtle glyph deformations, yet robust to various confounding sources of noise associated with digitized historical books. We evaluate our approach against other common methods for image comparison on a downstream damaged type-imprint matching dataset of English early modern (c. 1500–1800) books, a period when print censorship was prevalent. Our model’s inferences are consistent with and support the findings of recent, manual bibliographical analyses of *Areopagitica* (Warren et al. 2020) and *Leviathan Ornaments* (Warren et al. 2021). In a deployment case study, a broader set of type from *Leviathan Ornaments* is matched against type from 138 candidate books and evaluated by an expert bibliographer. Results suggest that our method can help scale printer attribution to the tens of thousands of anonymously printed early modern documents.

Related Work

Our work fits into a related line of work on computational analysis of historical documents under limited supervision—including, for example, unsupervised optical character recognition (OCR) techniques suited to early modern English historical documents such as *Ocular* (Berg-Kirkpatrick, Durrett, and Klein 2013a; Berg-Kirkpatrick and Klein 2014). Garrette et al. (2015) later proposed an OCR solution for code-switched historic texts using *Ocular*, followed by advances in improved orthographic transcriptions for OCR (Garrette and Alpert-Abrams 2016). Separately, Liu and Smith (2020) investigated code-switching in historic German books with applications to OCR. Neural techniques, such as the *LatinBERT* masked language model for classical Latin language processing (Bamman and Burns 2020), and *Lacuna* for low-resource document transcription for English and Arabic-script documents (Vogler et al. 2022), both utilize self-supervised pre-training to adapt to limited supervised resources. In this work, we instead overcome data scarcity by synthesizing our data using a random augmentation procedure.

Recent work has also used machine learning to discover the origins of historical texts. Assael et al. (2022)’s *Ithaca* uses deep neural networks for geographical and chronological attribution of ancient Greek inscriptions with the help of historians. Closest in theme to our work is Ryskina et al. (2017)’s automatic compositor attribution, which considers the bibliography of the document by predicting the compositors of Shakespeare’s *First Folio* using orthographic and spacing features. While these approaches primarily focus on the text in the documents for computational analysis, we focus on often-overlooked but important complementary features

Books with a shared printer

Revealed by distinctive imprints from a damaged metal stamp

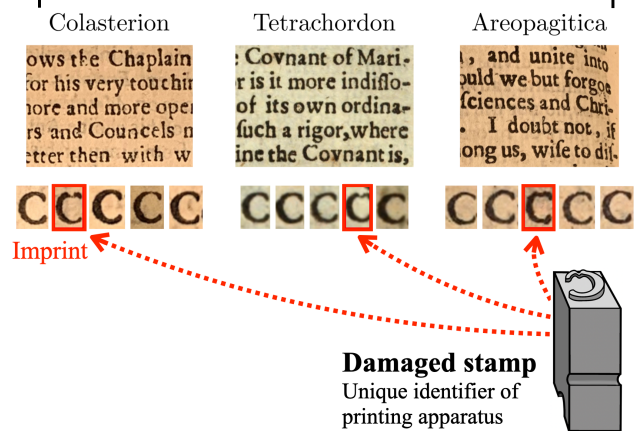


Figure 2: Although the three damaged type-imprints from Warren et al. (2020)’s study, in red, occur with different inking levels and imaging variation, the letters were pressed using the same type piece, evident from the corresponding damage at the top of the glyphs, suggesting that the respective books share the same printer. We are the first to approach this problem using machine learning, which could scale printer de-anonymization to tens of thousands of such books.

related to the *visual* appearance of books for computational attribution of print.

Our work also draws from a body of work on metric learning over images using contrastive approaches and data augmentation (Hadsell, Chopra, and LeCun 2006; Musgrave, Belongie, and Lim 2020; Weinberger and Saul 2009; Zhai and Wu 2019) which are commonplace for many computer vision tasks like self-supervised learning of image representations (Chen et al. 2020), image classification, and object localization (Ki et al. 2021). While these approaches have been shown to be effective for capturing high-level features that characterize semantic variation among images, they are not suitable wholesale for our purpose which is to compare pairs of images based on subtle shape variations while remaining invariant to other dominant sources of variation.

Early Modern Print and Damaged Type Pieces

We focus on a data domain that is highly relevant to the cultural analysis of censorship. From the advent of the printing press in the mid-1400s through the 18th century, printers declined to attach their names to tens of thousands, or about 25%, of all known books and pamphlets. Reasons included controversial content, censorship laws, and piracy (Como 2007; Norbrook 1994; Raymond 2017; Woodfield 1991; Towers 2003; McCabe 1981; McCoog 2016). As early as 1960, analytical bibliographers realized that distinctively damaged metal *type pieces* used to produce the inked *type-imprint* on a book’s page could be recorded and compared across publications to uncover the hidden identities of printers (Mills 1960). Essentially, once a metallic type piece is damaged, whether by being dropped on the ground, warped under pressure, or some other means, the type-imprints produced by

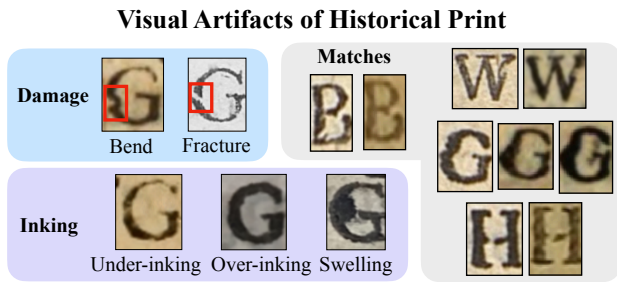


Figure 3: Bending and fracturing, the two main kinds of damage to type pieces in printed works, can be visually observed via the type-imprint left on a book’s page. Three kinds of inking variation may complicate damage matching by either resembling or occluding damage. Under/over-inking generally appear evenly throughout a type-imprint, while swelling is more local. Example match groups are shown on right.

the type piece become unique like a fingerprint as illustrated in Figure 2 and the top of Figure 3. Manual forensic analyses of these most distinctive aberration patterns—bends and fractures, as well as other tell-tale cues identified by bibliographers, have produced evidence connecting anonymously printed works with known printers (Weiss 1992; van den Berg and Howard 2004; Achinstein and Burton 2013; Garrett 2014; Bricker 2016; Lavin 1972; Adams 2010; Como 2012). For example, Charlton Hinman’s pioneering work in the 1960s exhaustively compared all letterforms across 55 copies of Shakespeare’s First Folio using careful notetaking methods to uncover the collation process (Turner 1966; Hinman 1963).

Recent printer attribution work uses digital reproductions in addition to physical copies (Adams 2010; Como 2018). Warren et al. (2020, 2021) go one step further, employing Ocular (Berg-Kirkpatrick, Durrett, and Klein 2013b) historical OCR for automatic extraction of type-imprint images, which they later *manually* match to uncover the printers of important books printed in the early modern period.

This type-imprint matching process remains the most labor-intensive part of the pipeline as every type-imprint image must be compared against an entire set of other type-imprint images. In this work, we focus on developing automatic computational techniques that aim to scale the matching process by filtering the search space and suggesting candidates of matching damaged type-imprints across large collections of digitally archived books.

Our Computational Approach

We use a neural model to learn a distance metric over pairs of type-imprint images. The metric learning framework facilitates treatment of matching as a search and rank problem, in which we embed a type-imprint and use it to *query* and *retrieve* the top- k most similarly damaged type-imprints from the dataset. Supervised training consists of learning to embed matched pairs comparably and unmatched pairs disparately using a margin loss. However, this modeling approach faces two main hurdles:

1. **Spurious sources of deviations between type-imprints**, such as difference in scale, position, inking (Fig 3), font




Figure 4: Our random data augmentation procedure models realistic bends, fractures, and inking to generate paired training data for learning to match damaged type-imprints. We sample two random undamaged images, perturb them with either a similar bend or a similar fracture (e.g., fracture here, green borders) using V character classes’ parameter sets, and apply one of three different types of random inking noise independently to each image (bottom, green borders).

variation, character class due to imperfect historical type-setting and OCR errors, and digital imaging noise, mislead traditional metric learning approaches from isolating variation due to the shape of the underlying type pieces.

2. The **lack of labeled data** complicates learning in the supervised metric learning framework. Collecting high quality image scans of early modern books requires access to rare book libraries, identifying and annotating anomalous damage is labor-intensive, and manual pairwise comparison of every type-imprint is unfeasible at scale.

In the following sections, we describe how we overcome these challenges. First, we present our novel metric learning model (Figure 1), which is able to overcome the confounding noise. Its restricted attention mechanism is well-suited to identifying shared local deviations in a pair of candidate images. Then, after describing the preparation of a collection of character type-imprint images, we describe how, by using

domain-specific knowledge about the type piece damages, we generate realistic synthetic pairs of matching damaged type-imprints to train our model using a supervised objective.

Contrastive Attention-Based Metric Learning

As seen in the three matching character type-imprints in Figure 2, identical type-imprints can be most easily recognized by the corresponding local deformations to character shape. We use an attention mechanism to allow our model to jointly focus on *corresponding locations* across a pair of images. We introduce our Contrastive Attention-based Metric Learning (CAML) model, shown in Figure 1, which computes a weighted attention map from the convolutional features of a type-imprint image pair to allow for more local comparison. At its core, our model takes as input a pair of type-imprint images we are interested in comparing and outputs a *distance* between the pair. As shown in Figure 1, X is an image of a type-imprint of G with local damage at the bottom of the glyph and X^+ is another type-imprint image we would like to compare to. In this case, X^+ also exhibits similar local damage as X , so the desired output of the model, the L2-distance between the final image embeddings e and e^+ , should be small.

Template feature-map residuals Most observable damage on a type-imprint image manifests as local deviations from an exemplary character shape. In order to aid the model in identifying this local deviation from a standard, undamaged character shape, we additionally input a grayscale template image X^t computed via the pixel-wise average over all training set images of the character. We pass these three images through a deep convolutional neural network to get output feature representations F, F^+ , and $F^t \in \mathbb{R}^{d \times h \times w}$ of the input images and the template image. Here h and w correspond to the height and width of the activations/feature-maps and d is the depth of the feature-maps. We encourage the model to focus on deviation between input features and template features by computing *residual feature-maps* \bar{F} and $\bar{F}^+ \in \mathbb{R}^{d \times h \times w}$ via an element-wise difference between the input images’ feature-maps and the template’s feature-map: $\bar{F} := (F - F^t)$ and $\bar{F}^+ := (F^+ - F^t)$.

Contrastive Attention Mechanism Assuming that images X, X^+, X^t have already been aligned (see Dataset section) so that the same pixel location in both images represents the same location on the type-imprint, we would like our model to be able to compare and contrast the same specific visuospatial features between residual feature-maps \bar{F} and \bar{F}^+ . Our proposed contrastive attention mechanism, composed of a multi-layer perceptron (MLP in Fig. 1) that computes an attention map $\hat{\alpha} \in \mathbb{R}^{h \times w}$ on the concatenated residual feature-maps $[\bar{F}; \bar{F}^+]$, enables this comparison. Each value $\hat{\alpha}_{ij}$ of $\hat{\alpha}$ represents a score for each spatial location, or “feature pixel”, on the *joint* location dimension of the residual feature-map: $\hat{\alpha}_{ij} = \text{MLP}([\bar{F}_{*ij}; \bar{F}^+_{*ij}])$.

After normalizing these attention scores, we obtain a weighed attention map $\alpha \in \mathbb{R}^{h \times w}$ over such feature pixels. We compute the final d -dimensional output image embeddings e and e^+ by attending to \bar{F} and \bar{F}^+ respectively:

$e := \sum_{ij} \alpha_{ij} \bar{F}_{*ij}$, and $e^+ := \sum_{ij} \alpha_{ij} \bar{F}^+_{*ij}$. Finally, the model returns the Euclidean distance between the embeddings $\|e - e^+\|_2$, representing the distance between the type-imprint image pair.

Training CAML We train CAML with the popular triplet loss (Weinberger and Saul 2009), which operates on an anchor/query embedding e along with the embedding e^+ of a candidate image that matches the anchor and a non-matching candidate image’s embedding e^- . This results in the following loss: $\max(\|e - e^-\|_2 - \|e - e^+\|_2 + m, 0)$, which focuses on minimizing the Euclidean distance between the anchor and the positive matching images’ embeddings, and maximizing the distance between the anchor and the non-matching images’ embeddings, such that the positive and negative examples are separated by a margin of at least m . We sample negative examples uniformly at random from our batch.

Synthesizing Damaged Type-Imprint Pairs

Given pairs of known matches and non-matches, we can train CAML. However, annotated data of this form is extremely sparse and difficult to collect. In this section, we describe how we first create a small collection of *unpaired* damaged character type-imprint images from early modern books of both known and unknown printer origins. Then we describe a random automatic process for synthesizing *paired* supervised match data from this collection for training CAML.

Extracting Type-Imprint Images from Books

We obtain page image scans from 38 different English books printed from the 1650s–1690s by both known and unknown printers of historical interest. The materials were printed mainly in London after its civil war when the explosion of news, pamphlets, and cheap print are said to have led to England’s ‘reading revolution’ (Como 2018; Sharpe 2000; Achinstein 1994; Raymond 2003; Zaret 2000). In contrast to resources like Early English Books Online (EEBO), used in Mak (2014), our images are higher resolution to facilitate fine-grained comparison of anomalous type-imprints.

Next, we use the Ocular OCR system (Berg-Kirkpatrick, Durrett, and Klein 2013a) to extract segmented character images from the color page photographs and select a subset of 16 capital letters, which tend to exhibit the most recognizable damaged compared to their relatively infrequent occurrence in printed English. Then, we align the resulting character images using learned rotation, offset, size, and scale random variables from Goyal et al. (2020)’s recently proposed generative font clustering model for typographical analysis of early modern printing, which significantly reduces the variance in size, offset, skewness, and rotation.

Supervised Data via Paired Damage Synthesis

In order to train our matching system, we require a dataset consisting of pairs of matching damaged characters — we propose a technique for synthesizing such pairs. Inspired by work in learning disentangled representations of morphologically perturbed MNIST images (Castro et al. 2019) and synthetically shifted images of characters from historical books for font discovery (Goyal et al. 2020), we design a detailed

random process for perturbing character images with *paired* realistic damage and inking variation observed in printing press era books. The generating process is depicted in Fig. 4 and described in this section.

First, we sample two random, undamaged letter type-imprint images from different books from the annotated collection described above (as depicted at the top of Fig. 4). By sampling the type-imprint images from different books instead of using the same book or even the same, duplicated character type-imprint image,¹ we force our model to discriminate between the synthesized local damage patterns instead of specious features arising from the book scanning process itself. We take the union of the sampled type-imprint pair and extract its skeleton (henceforth referred to as *union skeleton*), which makes it easier to produce a bend/fracture on the same location of the pair.² We randomly apply one of two kinds of damage—bends or fractures—followed by inking noise, as described below.

Bends In the case of a bend (Fig. 3, top left), we first sample the *length of the segment* to be bent and the *amount to shift* the segment (both as relative percentages of skeleton height) from the respective character class distributions (Fig. 4, right). Next, we iteratively sample for the best set of midpoints and endpoint locations on the *union skeleton* and map them back to the original images using the nearest neighbor on each skeleton. We morph the sampled *segment length* by drawing a Bezier curve between the midpoint and a shifted midpoint computed with the sampled *shift amount*. We increase the bent skeleton’s thickness until it reaches the unbent character’s thickness, as measured by a Euclidean distance transform.

Fractures We sample a center point on the *union skeleton* and map it back to the original images in a similar fashion as in the bends. For each type-imprint in the pair, we locally erode the image by a sampled *percentage of its measured mean character thickness* in a small window around the center point on the character skeleton. Then, we sample a random *angle* and draw a circular brush stroke through the center point at this angle, similar to Castro et al. (2019), with sampled *thickness* proportional to mean character thickness. Finally, the image is locally dilated back to its original thickness amount. By placing the brush stroke between local erosion/dilation, the sharp edges introduced by the circular brush stroke are removed and the damage looks less artificial.

Inking In early printing press era books, natural variation in a type piece’s inking level can cause significant visual differences in type-imprints produced from the same underlying metal type piece, as shown in Fig. 3. Under-inking can create superficially similar effects as damage, whereas over-inking can occlude actual damage to the metal type piece. Either the whole image could be over or under-inked or local sections of glyphs could exhibit inking variations. We over- or under-ink both undamaged and artificially damaged character type-imprint images *independently* by dilating or

eroding the image by a sampled percentage *amount* of the type-imprint’s mean thickness, which evenly thickens or thins the character (Fig 4, left). Similar to Castro et al. (2019), for local inking, we perturb a character type-imprint image by first uniformly sampling a random location on the character skeleton, sampling a *strength* magnitude to increase inking, and *radius* amount, followed by warping of image coordinates around the sampled location. Modeling these inking variations in our perturbation generating process exposes our learned matching models to this deceptive real-world noise.

Setting Parameters Instead of learning these damage and inking parameters, which would be difficult given the limited labeled data, we consult humanities scholars with expert domain knowledge. We set parameters of the truncated normal distributions controlling the image operations (shown in Fig. 4) through multiple iterations of tuning, sampling, and evaluation. In each round, we display comparisons between hundreds of samples of real and synthesized damages and have annotators evaluate the fidelity of the generated samples.

Experimental Setup

In this section, we first describe the real-world datasets containing a small amount of manually identified matches and damaged types in previous bibliographic studies that we use to evaluate our approach on downstream matching of type-imprint images. We use these datasets to construct various scenarios, ranging from highly optimistic to more realistic settings, pertaining to the quality of the candidate sets we query against for our matching experiments. Then, we present strong baselines that are prevalent for computational image comparison for empirical comparison against our proposed approach.

Ground Truth Evaluation Datasets

We use two different hand-curated datasets from recent bibliographical studies that manually identified and matched damaged type-imprints for attribution of two major early modern printed works (Warren et al. 2020, 2021).

***Areopagitica* validation set** We collect a small *validation* set of the manually identified type-imprint matches used in the study for printer attribution of John Milton’s anonymously printed *Areopagitica* (Warren et al. 2020). Specifically, we focus on four uppercase characters D, F, G, and M whose damaged type-imprints were compared across books from known printers of interest for this study. Overall, this dataset contains 128 total match groups with 159 pairwise queries.

***Leviathan Ornaments* test set** We construct our *test* set from the expert-curated set of matches manually identified in a recent bibliographical study, in which Warren et al. (2021) established that the book was printed in a single print shop by John Richardson in 1695–1696, thus refuting the attribution of Malcolm (2008) to a different printer, John Darby in earlier work by amassing evidence from damaged type-imprints matched manually within *Leviathan Ornaments* and across other books by known printers from the suspected time period of its printing. Overall, this dataset contains 217 total match groups with 858 pairwise queries.

¹In fact, training on paired data synthesized from the same character type-imprint image does not generalize well to real datasets.

²We use *scikit-image* morphology (van der Walt et al. 2014).

For the purpose of damaged type-imprint retrieval, we treat the known damaged type-imprints in *Areopagitica* and *Leviathan Ornaments* as queries against which we attempt to retrieve the matching damaged type-imprints from candidate sets of interest. For the *Areopagitica* validation set, we simply match against the other *queries*—i.e. every candidate is a match for at least one query. Note, this setup is unrealistically easy, and as we will see in experiments, simple baselines do unrealistically well. However, we find it useful to leverage *Areopagitica* for validation and early stopping. *Leviathan Ornaments* represents our main test evaluation set. Here, we conduct experiments with two different setups that include different types of realistic candidates sets: (1) **Strong negative** (Lev-Strong), in which the candidate set is composed of both ground truth positive matches and 1000 random type-imprint images from other books printed by Darby and Redmayne, who were both suspected of being the printer behind *Leviathan Ornaments* until Warren et al. (2021)’s study, and (2) **Mix negative** (Lev-Mix), where the candidate set consists of both ground truth positive matches, 500 random images from Lev-Strong, and 500 random images from books printed by Robert Everingham, who used type pieces in a different font and was never considered to be involved in the printing of *Leviathan Ornaments*. These test setups are more realistic because they contain a larger variety of negative candidates and are based on the actual data bibliographers have combed through to identify matches. The latter setting, Lev-Mix, is possibly the most difficult because it includes the most variability in negative examples, making spurious matches using simple methods more likely.

Baseline Models

In order to show the effectiveness of our approach on matching subtle damages between type-imprints in the presence of multiple confounding sources of noise, we compare it to other prevalent approaches for image comparison.

Image L2 Distance (L2) Instead of learning a metric over the set of training images, this method compares the aligned (see Dataset section) query images to the candidate images by computing the Euclidean distance between the two images and outputs a ranked list of candidates for each query image.

Embedding Triplet Loss (Emb) In a large metastudy on metric learning for images, Musgrave, Belongie, and Lim (2020) report that classic embedding based triplet losses are highly effective and closely match the state-of-the-art on metric learning tasks. Therefore, we compare our method to an *embedding based triplet loss* approach as described in Weinberger and Saul (2009). This approach is trained with the same loss as CAML and is similar to the approach described in the CAML section, but the images (anchor, match, non-match) for triplet loss training are processed independently of each other by a CNN with no attention mechanism to produce their respective embeddings.

For fair comparison, the architecture, artificial data creation and negative mining strategies are the same as our approach. The one *major difference* between our approach (CAML) and this baseline approach is that our model considers both the anchor and the candidate jointly via the de-

Method	Syn Valid	Areo	Lev-Strong	Lev-Mix
L2 Aligned	5.53	43.21	39.91	40.99
Stacked BCE	4.62	1.78	1.10	0.96
Emb	5.72	40.76	33.53	33.00
CAML	35.50	39.42	58.15	56.08

Table 1: Recall @ 5 results micro-averaged over 4 character classes for *Areopagitica* ground truth validation set, and 16 character classes for both the synthetic validation set and *Leviathan Ornaments* ground truth test set. Both the Emb and CAML models use convolutional residual features.

scribed attention mechanism to yield the relevant embeddings whereas this baseline learns the embeddings for each triplet image independently. As observed in the results, this difference makes our model more suitable for metric learning over character images with very subtle local variation.

Stacked BCE To investigate the effect that our attention-based architecture design and contrastive loss function has on performance, we train the same convolutional neural network feature extractor to classify positive and negative pairs (Zhai and Wu 2019) of images with a binary cross-entropy loss. Instead of encoding the 1-channel images separately, we stack the images on separate channels before inputting to the CNN.

Quantitative Results

In this section, we describe our quantitative results on the ground truth evaluation data presented in Experimental Setup. We report all-pairs *Recall@k=5* micro-averaged over the character classes in each dataset configuration in our evaluations so every damaged type-imprint serves as the query image once. We train each model for 60 epochs and early stop using the best *Areopagitica* validation set recall.³

In Table 1, we compare results using the L2 baseline on the aligned images, Stacked BCE classifier, Triplet Embedding model, which is equivalent to CAML without the spatial attention mechanism, and our proposed CAML method. First, we observe that CAML outperforms the nearest baseline approach by at least 15–18 points on recall @ 5 on the *Leviathan Ornaments* test settings. The next closest methods for this setting are the L2 baseline and the Emb model, which both tend to lack the ability to focus on local damage similarities. We hypothesize that character positioning and inking dominates most of the distance for these methods. This explanation is plausible when considering the poor synthetic validation performance, which is a dataset consisting of type-imprints from different underlying fonts to force models to focus on local similarities in lieu of font shape. Also, this is supported by the high recall on the *Areopagitica* validation set—all of the characters in it are hand-identified ground truth matches meant to be as similar as possible to make a convincing bibliographical argument (Warren et al. 2020) and no other misleading candidate images need to be filtered out. The Stacked BCE model, while capable of driving its classification loss near zero on the training set, has a lot of difficulty ranking unseen data and performs worst of all methods.

³Code located at <https://github.com/nvog/damaged-type>.

Residual	Syn Valid	Areo	Lev-Strong	Lev-Mix
None	50.62	33.63	62.88	61.78
Input	47.62	38.75	52.33	49.18
CNN	45.00	43.88	67.81	64.79

Table 2: Micro-averaged recall @ 5 independent ablation study comparing different residual methods in the CAML model. Results on D, F, G, M character subset.




Figure 5: CAML attention for 3 query/candidate pairs.

Ablation We perform two kinds of ablative studies on our model. The first study (Table 2) concerns with how we obtain the residuals of the image representations through subtracting the template character features. We compare 3 settings: (1) no template residual, (2) an input residual where the template image pixels are subtracted from each input image’s pixels, and (3) CNN feature residual (our final model), where the learned template image’s CNN features are subtracted from the learned input image’s CNN features before performing spatial attention. We observe that providing information about template shape of the character class is important for our model’s success as shown by the superior performance of the *CNN residual* variant. Interestingly, not using any residuals performs better than subtracting the templates from the input image, suggesting that the *input residual* variant suffers from information loss before CNN feature computation—perhaps due to misaligned characters or inking variation.

Additionally, we perform a small ablation analysis to study training CAML without global inking synthesis such as Thinning and Thickening from Fig. 4. Compared to our best model which uses this augmentation, most of the decrease in performance is on synthetic validation (3 points R@5) and the Lev-Mix setting (0.68 R@5).

Attention Visualization In Figure 5, we visualize the spatial attention weights on a few top-ranked G type-imprint images from the Lev-Mix test set. We compare two pairs (left, middle) using the same ‘Query 1’, while the third pair (right) uses a different ‘Query 2’. The highest attention weights are clearly located in locations with noticeable damage, along with other areas that help the model differentiate glyph shapes. Weights appear mostly similar for the pairs with the same query, but adapt drastically for the other pair.

Deployment Study: CAML for Attribution

For analyzing how our model can aid bibliographical research, we design a deployment and expert evaluation case study to confirm the manual printer attribution of *Leviathan Ornaments* (Warren et al. 2021). Using an unannotated set of 138 books scanned by various libraries and printed by 4 suspected printers of *Leviathan Ornaments* and 1 unrelated printer, we aim to simulate the entire process of human-in-the-loop printer attribution, from query selection to matching. Until now, such a process has required much manual effort.

Printer	Top-1		Top-5		Top-10	
	Lo	Hi	Lo	Hi	Lo	Hi
Richardson	7	11	17	23	40	31
Richardson & Holt	0	0	2	0	2	0
Darby	13	3	21	7	28	9
Redmayne	0	0	0	0	0	0
Everingham	0	0	5	1	9	1

Table 3: CAML deployment findings aggregated across matches for known printers on *Leviathan Ornaments*. Expert is asked to annotate both low (lo) and high (hi) confidence matches returned by system. The most high confident matches are bolded, confirming Warren et al. (2021)’s manual attribution to the printer John Richardson.

Query Selection and Matching First, we start with query selection, which involves choosing hundreds of type-imprints exhibiting anomalous bends and fractures in *Leviathan Ornaments*, the anonymously printed book of interest. Instead of manually identifying these by sorting through tens of thousands of characters, we train a CNN-based classifier with the same architecture as our matching models on thousands of unpaired, labeled images from other 17th-century books to rank all *Leviathan Ornaments* type-imprints by damage intensity. An expert bibliographer then selects a few hundred of the top results to create a query set of 246 images. Next, using the L2 distance on CAML’s image embeddings, we generate the top-10 ranked candidates for each query image against a set of 518,891 total type-imprint images.

Annotation Findings We setup an annotation interface and ask a bibliographical expert to annotate matching type-imprints among the retrieved candidates as either not a match, or as a *high* or *low* confidence match without access to book/printer name or scoring information. In Table 3, we present aggregated match counts by summing up the counts for each printers’ books. Of the 246 unique damaged type-imprint queries, at least one match was found in 80 of them (32.5%). Among the suspected candidate printers, John Richardson has the largest amount of high confidence matches. In contrast to John Darby, who Malcolm (2008) attributed the book to, John Richardson has 22 more high confidence matches in the top-10. While John Redmayne printed another clandestine version of Hobbes’ *Leviathan*, zero matches surfaced in our investigation. Robert Everingham’s books, which tend to use different sets of fonts entirely, only surfaced a single high match. Ultimately, evidence strongly suggests that John Richardson did, in fact, print *Leviathan Ornaments*.

Conclusion

We demonstrate that machine learning can be successfully applied to printer attribution by learning to match damaged type-imprint images in early modern books. By attributing early modern print at scale, we can begin to uncover the hidden figures behind the tens of thousands of clandestinely printed works from the period, which amounts to roughly 25% of documents. Ultimately, this translates into more opportunities to discover significant historical networks of early print media involving printers, authors, and arguments.

Acknowledgments

This project is funded in part by the NSF under grant 1936155, and by the NEH under grant HAA-284882-22. We would like to thank Ciaran Evans, DJ Schuldt, Kari Thomas, Laura DeLuca, and the rest of the Print & Probability team, along with Berg Lab, and the anonymous reviewers for their assistance and helpful feedback.

References

- Achinstein, S. 1994. *Milton and the revolutionary reader*. Princeton N.J.: Princeton University Press. ISBN 978-0-691-03490-4.
- Achinstein, S.; and Burton, B. 2013. Who Printed Milton's Tetrachordon (1645)? *The Library*, 14(1): 18–44.
- Adams, D. R. 2010. The Secret Printing and Publishing Career of Richard Overton the Leveller, 1644–46. *Library*, 11(1): 3–88.
- Assael, Y.; Sommerschild, T.; Shillingford, B.; Bordbar, M.; Pavlopoulos, J.; Chatzipanagiotou, M.; Androutopoulos, I.; Prag, J.; and de Freitas, N. 2022. Restoring and attributing ancient texts using deep neural networks. *Nature*, 603(7900): 280–283.
- Bamman, D.; and Burns, P. J. 2020. Latin bert: A contextual language model for classical philology. *arXiv preprint arXiv:2009.10053*.
- Berg-Kirkpatrick, T.; Durrett, G.; and Klein, D. 2013a. Unsupervised Transcription of Historical Documents. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 207–217. Sofia, Bulgaria: Association for Computational Linguistics.
- Berg-Kirkpatrick, T.; Durrett, G.; and Klein, D. 2013b. Unsupervised Transcription of Historical Documents. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 207–217. Sofia, Bulgaria: Association for Computational Linguistics.
- Berg-Kirkpatrick, T.; and Klein, D. 2014. Improved Type-setting Models for Historical OCR. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 118–123. Baltimore, Maryland: Association for Computational Linguistics.
- Bricker, A. 2016. Who was “A. Moore”? The Attribution of Eighteenth-Century Publications with False and Misleading Imprints. *The Papers of the Bibliographical Society of America*, 110(2): 181–214.
- Castro, D. C.; Tan, J.; Kainz, B.; Konukoglu, E.; and Glocker, B. 2019. Morpho-MNIST: Quantitative Assessment and Diagnostics for Representation Learning. *Journal of Machine Learning Research*, 20(178).
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Como, D. R. 2007. Secret Printing, the Crisis of 1640, and the Origins of Civil War Radicalism. *Past & Present*, 196(1): 37–82.
- Como, D. R. 2012. Print, Censorship, and Ideological Escalation in the English Civil War. *Journal of British Studies*, 51(4): 820–857.
- Como, D. R. 2018. *Radical Parliamentarians and the English Civil War*. Oxford: Oxford University Press. ISBN 978-0-19-954191-1. OCLC: 1089237699.
- Garrett, C. E. 2014. How T. S. Became Known as Thomas Sherman: An Attribution Narrative. *The Papers of the Bibliographical Society of America*, 108(2): 191–216.
- Garrette, D.; and Alpert-Abrams, H. 2016. An Unsupervised Model of Orthographic Variation for Historical Document Transcription. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 467–472. San Diego, California: Association for Computational Linguistics.
- Garrette, D.; Alpert-Abrams, H.; Berg-Kirkpatrick, T.; and Klein, D. 2015. Unsupervised Code-Switching for Multilingual Historical Document Transcription. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1036–1041. Denver, Colorado: Association for Computational Linguistics.
- Goyal, K.; Dyer, C.; Warren, C.; G'Sell, M.; and Berg-Kirkpatrick, T. 2020. A Probabilistic Generative Model for Typographical Analysis of Early Modern Printing. In *Proceedings of 2020 Annual Conference of the Association for Computational Linguistics*.
- Hadsell, R.; Chopra, S.; and LeCun, Y. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, 1735–1742. IEEE.
- Hinman, C. 1963. *The printing and proof-reading of the first folio of Shakespeare*. Oxford: Clarendon Press.
- Ki, M.; Uh, Y.; Choe, J.; and Byun, H. 2021. Contrastive Attention Maps for Self-supervised Co-localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2803–2812.
- Lavin, J. A. 1972. The Printer of “Hamlet” Q3. *Studies in Bibliography*, 25: 173–176.
- Liu, S.; and Smith, D. 2020. Detecting de minimis Code-Switching in Historical German Books. In *Proceedings of the 28th International Conference on Computational Linguistics*, 1808–1814. Barcelona, Spain (Online): International Committee on Computational Linguistics.
- Mak, B. 2014. Archaeology of a digitization. *Journal of the Association for Information Science and Technology*, 65(8): 1515–1526.
- Malcolm, N. 2008. The Making of the Ornaments: Further Thoughts on the Printing of the Third Edition of Leviathan. *Hobbes Studies*, 21(1): 3–37.
- McCabe, R. A. 1981. Elizabethan Satire and the Bishops' Ban of 1599. *The Yearbook of English Studies*, 11: 188–193.
- McCoog, T. M. 2016. “Guiding souls to goodness and devotion” : clandestine publications and the English Jesuit mission. In Bela, T.; Calma, C.; and Rzegocka, J., eds., *Publish-*

- ing subversive texts in Elizabethan England and the Polish-Lithuanian Commonwealth*, 93–109. Leiden: Brill. OCLC: 953683964.
- Mills, J. C. 1960. Detective in the Book World. *Graphic Arts Review*, 23.
- Musgrave, K.; Belongie, S.; and Lim, S.-N. 2020. A metric learning reality check. In *European Conference on Computer Vision*, 681–699. Springer.
- Norbrook, D. 1994. Areopagitica, Censorship, and the Early Modern Public Sphere. In Burt, R., ed., *The Administration of Aesthetics*, volume 7 of *Censorship, Political Criticism, and the Public Sphere*, 3–33. University of Minnesota Press, new edition. ISBN 978-0-8166-2365-5.
- Raymond, J. 2003. *Pamphlets and pamphleteering in early modern Britain*. Cambridge: Cambridge University Press. ISBN 0-521-81901-6 978-0-521-81901-5.
- Raymond, J. 2017. Censorship in Law and Practice in Seventeenth-Century England: Milton's Areopagitica. In Hutson, L., ed., *Oxford Handbook of English Law and Literature, 1500-1700*. Oxford: Oxford University Press.
- Ryskina, M.; Alpert-Abrams, H.; Garrette, D.; and Berg-Kirkpatrick, T. 2017. Automatic Compositor Attribution in the First Folio of Shakespeare. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 411–416. Vancouver, Canada: Association for Computational Linguistics.
- Sharpe, K. 2000. *Reading revolutions: the politics of reading in early modern England*. New Haven (Conn.: Yale University Press. ISBN 978-0-300-08152-7. OCLC: 1014626495.
- Towers, S. M. 2003. *Control of Religious Printing in Early Stuart England*. Boydell Press. ISBN 978-0-85115-939-3.
- Turner, R. K. 1966. Reappearing Types as Bibliographical Evidence. *Studies in Bibliography*, 19: 198–209.
- van den Berg, S. J.; and Howard, W. S. 2004. G. M. Revealed: Printer of the First Attacks on The Doctrine and Discipline of Divorce. *Milton Quarterly*, 38(4): 242–252.
- van der Walt, S.; Schönberger, J. L.; Nunez-Iglesias, J.; Boulogne, F.; Warner, J. D.; Yager, N.; Gouillart, E.; Yu, T.; and the scikit-image contributors. 2014. scikit-image: image processing in Python. *PeerJ*, 2: e453.
- Vogler, N.; Allen, J.; Miller, M.; and Berg-Kirkpatrick, T. 2022. Lacuna Reconstruction: Self-Supervised Pre-Training for Low-Resource Historical Document Transcription. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 206–216. Seattle, United States: Association for Computational Linguistics.
- Warren, C. N.; Williams, P.; Rijhwani, S.; and G'Sell, M. 2020. Damaged Type and Areopagitica's Clandestine Printers. *Milton Studies*, 62(1): 1–47.
- Warren, C. N.; Wiscomb, A.; Williams, P.; Lemley, S. V.; and G'Sell, M. 2021. Canst Thou Draw Out Leviathan with Computational Bibliography? New Angles on Printing Thomas Hobbes' "Ornaments" Edition. *Eighteenth-Century Studies*, 54(4): 827–859.
- Weinberger, K. Q.; and Saul, L. K. 2009. Distance metric learning for large margin nearest neighbor classification. *Journal of machine learning research*, 10(2).
- Weiss, A. 1992. Shared Printing, Printer's Copy, and the Text(s) of Gascoigne's "A Hundreth Sundrie Flowres". *Studies in Bibliography*, 45: 71–104.
- Woodfield, D. B. 1991. *Surreptitious printing in England, 1550-1640*. New York: Bibliographical Society of America. OCLC: 704200450.
- Zaret, D. 2000. *Origins of Democratic Culture: Printing, Petitions, and the Public Sphere in Early-Modern England*. Princeton, N.J.: Princeton University Press.
- Zhai, A.; and Wu, H.-Y. 2019. Classification is a strong baseline for deep metric learning. *British Machine Vision Conference*.