

Few-Shot Connectivity-Aware Text Line Segmentation in Historical Documents

Rafael Sterzinger[✉], Tingyu Lin[✉], and Robert Sablatnig[✉]

Computer Vision Lab, TU Wien, Vienna, AUT
{firstname.lastname}@tuwien.ac.at

Abstract. A foundational task for the digital analysis of documents is text line segmentation. However, automating this process with deep learning models is challenging because it requires large, annotated datasets that are often unavailable for historical documents. Additionally, the annotation process is a labor- and cost-intensive task that requires expert knowledge, which makes few-shot learning a promising direction for reducing data requirements. In this work, we demonstrate that small and simple architectures, coupled with a topology-aware loss function, are more accurate and data-efficient than more complex alternatives. We pair a lightweight UNet++ with a connectivity-aware loss, initially developed for neuron morphology, which explicitly penalizes structural errors like line fragmentation and unintended line merges. To increase our limited data, we train on small patches extracted from a mere three annotated pages per manuscript. Our methodology significantly improves upon the current state-of-the-art on the U-DIADS-TL dataset, with a 200% increase in Recognition Accuracy and a 75% increase in Line Intersection over Union. Our method also achieves an F-Measure score on par with or even exceeding that of the competition winner of the DIVA-HisDB baseline detection task, all while requiring only three annotated pages, exemplifying the efficacy of our approach. Our implementation is publicly available at: https://github.com/RafaelSterzinger/acpr_few_shot_hist.

Keywords: Few-Shot Learning · Text Line Segmentation · Historical Document Analysis.

1 Introduction

In historical documents, text line segmentation is a foundational task of computational document analysis. Approaches to early layout analysis relied on heuristics, such as projection profiles and clustering of connected components [14,18,21]. While effective on clean documents, these rule-based pipelines often fail when faced with the complexities inherent in historical manuscripts, such as bleed-through, elaborate ornamentation, or degraded scripts [15]. As a consequence of these challenges, a shift towards data-driven approaches occurred using deep neural networks such as CNNs, which re-framed the task as an end-to-end learning problem [16]. However, the performance of deep learning models depends on the availability of large, pixel-precise annotated datasets. In the context of

historical documents, the required annotated datasets are either prohibitively expensive and time-consuming to create, or the data itself is simply unavailable.

Given this data scarcity, much research has focused on few-shot learning, where models are trained using only a handful of annotated examples. The challenge of this low-data regime is evident in research where powerful architectures, such as DeepLabV3+[4], achieve poor results when trained on just three pages per manuscript [31]. In response to these challenges, some have employed increasingly complex models, such as transformer-based hybrids [25].

Our approach, however, illustrates the principle of *Occam's Razor*: when a simpler method outperforms more complex ones, additional complexity is not only unnecessary but can be counterproductive, as an increased number of parameters raises the risk of overfitting and reduces generalization. In this work, we make the following contributions:

- We demonstrate that a simple and lightweight architecture is sufficient to significantly outperform more complex models [31] in a few-shot setting. In addition, we incorporate an established patch-level training strategy, which efficiently handles high-resolution imagery and has proven effective in data-scarce domains [6,23] to increase data.
- We propose the usage of a connectivity-aware loss function, originally proposed for neuron morphology reconstruction [12], to the task of text line segmentation, to address the critical importance of structural integrity in text lines: standard pixel-level losses often fail to prevent fragmentation or unintended merging, which severely impact line-level evaluation metrics.
- We perform a thorough evaluation of our methodology on two public competition datasets, U-DIADS-TL [31,2] and DIVA-HisDB [22], under a strict few-shot constraint of only three annotated training pages. Our results establish a new SOTA for the former and demonstrate the approach's generalizability on the latter.

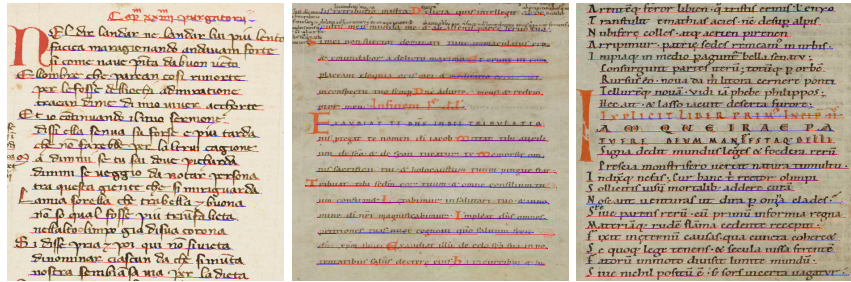
Our final model sets a new baseline on the U-DIADS-TL dataset, achieving relative gains of over 75% in Line Intersection over Union (IoU) and nearly 200% in Recognition Accuracy (RA) compared to the baseline reported by Zottin et al. [31]. Furthermore, on the DIVA-HisDB baseline detection task, our few-shot approach achieves performance competitive or even superior to top-ranked systems from the ICDAR 2017 competition [22], despite using three instead of 20 annotated pages, i.e., 85% less training data.

2 Related Work

In the following, we provide an overview of relevant literature on text line segmentation as well as few-shot learning with an emphasis on its application within historical documents.



(a) Latin 14396 (b) Latin 2 (c) Syriac 341



(d) CB55 (e) CSG18 (f) CSG863

Fig. 1: Our few-shot method improves SOTA on U-DIADS-TL (Top) by 200% in RA and 75% in Line IoU, and matches or surpasses DIVA-HisDB (Bottom) top scores using just three annotated pages.

Learning-Based Text Line Segmentation. Introducing Fully Convolutional Networks (FCNs) shifted research toward an end-to-end, data-driven formulation [16]. UNet [19], an encoder-decoder architecture with skip connections, demonstrated that precise pixel-level predictions can be learned. In the context of historical documents, Barakat et al. [3] trained an FCN able to cope with irregular inter-line gaps in medieval Hebrew scrolls. In contrast, Fizaine et al. [10] evaluated a Mask-RCNN against UNet on four public archives and found that instance-aware models handle intertwining marginalia more effectively. More recent transformer hybrids such as SeamFormer [25] employ scribble-conditioned refinement to link diacritics and resolve touching ascenders, achieving SOTA precision on dense palm-leaf manuscripts. Alberti et al. [1] further combined semantic pixel labeling with polygon extraction in a labeling-cutting-grouping pipeline.

To preserve global structure, recent work augments standard segmentation losses with connectivity-aware terms. For instance, such losses have been used

for road extraction [17], for class-imbalance mitigation [20], or have evolved into supervoxel connectivity losses that maintain object continuity with linear complexity and minimal overhead [12].

Few-Shot Learning Approaches. With the shortage of richly annotated datasets, research shifted to text line segmentation under extreme data scarcity. Zottin et al. [31] explored FCNs, the PSPNet [27], and DeepLabV3+ [4] with only three labeled pages per manuscript. In a similar context, De Nardin et al. [6] proposed an efficient framework that converts a handful of high-resolution manuscript pages into several thousand training patches via dynamic instance generation and refines coarse masks with Sauvola thresholding, achieving fully supervised quality on DIVA-HisDB with only two annotated pages per manuscript for layout analysis. An extension replaced Sauvola with adaptive local thresholding while retaining the same pixel-level F-score under a stricter evaluation [7]. Follow-up work demonstrated that a single labeled page per Arabic manuscript can suffice when combined with lightweight augmentation and class-balanced losses [8].

3 Methodology

In the following section, we describe our methodology to enable efficient few-shot learning capabilities for text line segmentation and baseline detection of ancient manuscripts.

3.1 Model Architecture

We adopt UNet++ [28], derived from the original UNet [19], paired with a ResNet34 [13] encoder. Despite its simplicity, the UNet++ consistently outperforms more complex architectures, such as the commonly used DeepLabV3+ architecture [31,6]. We attribute this to better generalization, reduced overfitting, and suitability to the data-scarce nature of historical document segmentation.

In comparison to [31], we opt for training on patch-level as opposed to the whole image, as this has been an established practice in domains where data is scarce, such as in the case of historical documents [6,23]. With accompanying benefits such as better capturing fine-grained text line structures or reducing memory usage, we apply our model to random crops extracted from high-resolution input images during training and overlapping patches during validation and testing. In detail, we extract patches of size 448×448 pixels and perform random rotations between -5 and $+5$ degrees and a random shear transformation between -3 and $+3$ degrees. We incorporate these augmentations to include non-horizontally aligned text lines.

To obtain a prediction for a whole page, overlapping patches are first processed independently, and predictions are stitched back together to form the full-size image. To avoid visual artifacts at border regions, each patch is scaled by a 2D Gaussian function, emphasizing its central region, which ensures smooth and seamless blending across overlapping areas.




Fig. 2: An illustration of the parameters α and β and their effect on the loss using a sample patch of Syriac 341: they control the trade-off between pixel- and structure-level errors, and between split and merge penalties, respectively [12].

3.2 Connectivity-Preserving Loss

As text lines, by their very nature, are continuous, curvilinear structures (cf. Fig. 2a), and maintaining their integrity – avoiding fragmentation or unintended merging – is crucial for accurate text line segmentation and detection. Concerning commonly used evaluation metrics for text line segmentation, minor errors at a pixel level, lead to significant drops in performance at the line-level.

We therefore propose the integration and adaptation of a novel connectivity-aware loss function, initially introduced by Grim et al. [12] for neuron morphology reconstruction. In detail, this loss also penalizes topological errors at a global component level rather than solely at the pixel level.

Formally, the loss function is expressed as:

$$\mathcal{L}(y, \hat{y}) = (1 - \alpha)\mathcal{L}_0(y, \hat{y}) + \alpha \left((1 - \beta) \sum_{C \in \mathcal{N}(\hat{y}_-)} \mathcal{L}_0(y_C, \hat{y}_C) \right) \quad (1)$$

$$+ \beta \sum_{C \in \mathcal{P}(\hat{y}_+)} \mathcal{L}_0(y_C, \hat{y}_C). \quad (2)$$

Here, y represents the ground truth text line segmentation and \hat{y} is the network’s prediction. Furthermore, the sets $\mathcal{N}(\hat{y}_-)$ and $\mathcal{P}(\hat{y}_+)$ identify components that would alter the number of connected components in the predicted text lines.

It is important to note that in our implementation, we use an *inverted ground truth* representation. Consequently, $\mathcal{P}(\hat{y}_+)$ captures false negative regions that cause a text line to appear fragmented (a false split), while $\mathcal{N}(\hat{y}_-)$ captures false positive regions that bridge two distinct ground truth lines (a false merge). Moreover, \mathcal{L}_0 refers to a standard pixel-level objective, such as, e.g., the binary cross-entropy loss. In our methodology, we adapted the loss by adding the Dice loss [24] and weighing the two terms to be of the same magnitude:

$$1 - \frac{2 \sum_i^N y_i \hat{y}_i}{\sum_i^N y_i + \sum_i^N \hat{y}_i} + \frac{10}{N} \sum_i^N \mathcal{L}(y_i, \hat{y}_i). \quad (3)$$

Finally, the hyperparameters α and β provide crucial control over the connectivity: $\alpha \in [0, 1]$ balances the influence of the conventional pixel-level loss against the topological penalties, with higher values prioritizing the correction of connectivity errors. As this is vital for maintaining the precise structure of text lines in ancient manuscripts prone to fragmentation or merging, we select $\alpha = 1$. Concurrently, $\beta \in [0, 1]$ allows for differential weighting between false merge and false split errors in text line segmentation; a $\beta < 0.5$ prioritizes penalizing merged lines, while a $\beta > 0.5$ emphasizes correcting fragmented lines, enabling tailored error correction based on the specific challenges of the segmentation task. Interestingly, for both our selected datasets, we found that focusing exclusively on penalizing false merged lines yielded the best performance, and thus we chose $\beta = 0$. To better illustrate the impact of these hyperparameters, Fig. 2 shows the penalized error in yellow overlaid on the prediction.

Due to the high computational cost of calculating the connected components and matching them, we utilize the structure-aware loss solely to fine-tune a model that has been pretrained with only the Dice loss. For similar reasons, we do not employ the structure-aware loss for model selection during the fine-tuning stage. Instead, we utilize the mean IoU as a proxy function to evaluate performance during validation.

4 Datasets

Our experiments are grounded in publicly available datasets from recent academic competitions focused on advancing text line segmentation for historical documents. These benchmarks, including the ICDAR 2017 HisDoc-Layout-Comp [22] and the ICDAR 2025 FEST track [30], emphasize challenges such as complex layouts, document degradation, and the latter, critically, few-shot learning constraints where only a handful of annotated examples are available for training. We use two such datasets: U-DIADS-TL[31,2] and DIVA-HisDB [22], a widely-used dataset for historical layout analysis. While our primary development targets the proposed few-shot paradigm of U-DIADS-TL, we also adapt and evaluate DIVA-HisDB in a few-shot setting to test the generalization of our method.

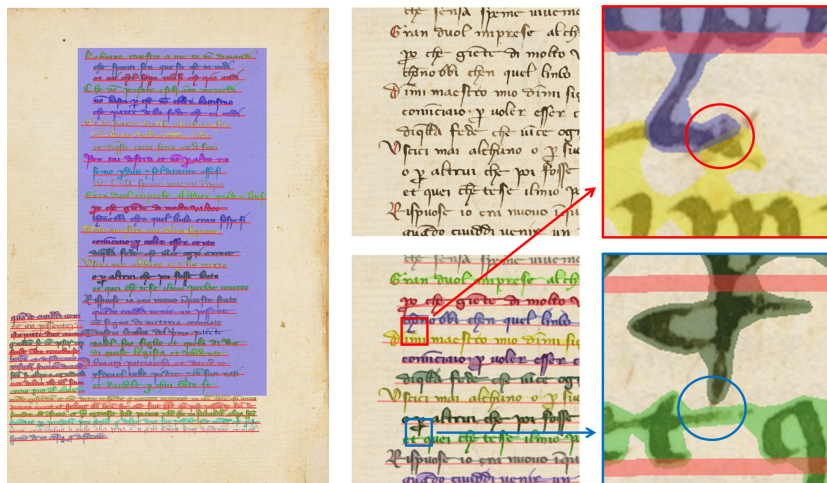


Fig 3: Example excerpts of the CB55 manuscript in DIVA-HisDB illustrating the issue of overlapping text line segments.

4.1 U-DIADS-TL

The U-DIADS-TL dataset is the official benchmark for the FEST 2025 competition and is explicitly designed for few-shot text line segmentation [31]. It is a subset of the larger U-DIADS-Bib [29] collection, which provides a rich resource for multi-class layout analysis.

Technical Details. The dataset contains 105 high-resolution color images from three distinct medieval manuscripts: Latin 2, Latin 14396, and Syriac 341. These sources feature diverse scripts, two- and three-column layouts, interlinear glosses, and significant degradation. For each manuscript, the data is strictly partitioned into 3 pages for training, 10 for validation, and 22 for testing, thus requiring effective methods to handle such low data regimes. The ground truth consists of pixel-precise binary masks that distinguish text lines from the background, with no overlap between line instances. Our methodology is designed to adhere to the official evaluation protocol described in [31,2].

4.2 DIVA-HisDB

The DIVA-HisDB dataset was introduced for the ICDAR 2017 HisDoc-Layout-Comp [22] competition. It is recognized for its complex page layouts, which include marginalia, decorated initials, and interlinear glosses.

Technical Details. The DIVA-HisDB consists of 150 pages from three manuscripts (CB55, CSG18, CSG863), provided as 600 dpi RGB scans with ground truth in PAGE XML format. For our experiments, we downsample all images by a factor

of three to make the image dimensions between the two datasets coherent. The dataset supports multiple layout analysis tasks, including full text line segmentation (Task 3), equal to the FEST 2025 competition. However, the polygon annotations in PAGE XML format for this task feature overlapping regions between adjacent lines, as illustrated in Fig. 3. This characteristic is incompatible with models designed to produce a single, non-overlapping binary segmentation mask, making a pixel-level evaluation ambiguous.

To address this, we opt for the task of baseline detection (Task 2), which requires the prediction of start and end-points of baselines instead of full polygons, providing a suitable proxy for text line localization without the issue of overlapping regions. To maintain consistency with the few-shot evaluation setting of U-DIADS-TL, we do not use the full training partition of 20 pages per manuscript, but instead, for each experiment, we randomly sample three pages for training, allowing us to assess our model’s few-shot generalization.

5 Evaluation

In the following section, we evaluate our proposed methodology. We begin by introducing the relevant evaluation metrics and experimental setup, followed by an ablation study of each component using our primary dataset, U-DIADS-TL. Finally, we assess the full approach on both benchmark datasets to underscore its effectiveness and generalizability.

5.1 Metrics

For the U-DIADS-TL dataset, we report Pixel IoU and Line IoU as well as Detection Rate (DR), Recognition Accuracy (RA), and F-Measure (FM) from matched line segments as proposed by Zottin et al. [31,2]. Pixel IoU and Line IoU are metrics based on the IoU, which is defined as:

$$\text{IoU} = \frac{TP}{TP + FP + FN}, \quad (4)$$

where TP is the number of true positives, FP false positives, and FN false negatives of the predicted binary mask. Subsequently, Pixel IoU evaluates segmentation quality at the pixel level, whereas Line IoU evaluates detection accuracy at the line level, where connected components are matched if pixel-level precision and recall exceed a threshold of 75%.

Concerning DR, RA, and FM, these metrics are based on a matching score, MatchScore. For predicted line R_i and ground-truth line G_j , the MatchScore is defined as:

$$\text{MatchScore}(i, j) = \frac{T(G_j \cap R_i)}{T(G_j \cup R_i)}, \quad (5)$$

where a pair (i, j) is considered a one-to-one match if $\text{MatchScore}(i, j) \geq 0.75$. Let M be the number of such matches, N_1 the number of ground-truth lines,

and N_2 the number of detected lines, then the match-based metrics are defined as follows:

$$\text{DR} = \frac{M}{N_1}, \text{RA} = \frac{M}{N_2}, \text{FM} = \frac{2 \cdot \text{DR} \cdot \text{RA}}{\text{DR} + \text{RA}}. \quad (6)$$

Based on overlapping text line annotations and resulting problems when evaluating text line segmentation on pixel-level (cf. Fig. 3), for the DIVA-HisDB dataset, we opt for the simplified task of text baseline detection. With this comes the necessity of employing suitable evaluation metrics for which we follow closely the evaluation protocol denoted in the work by Simistira et al. [22], which is based on the cBAD competition [9]. Here, an F-Measure metric is employed, which assesses baseline extraction by calculating recall and precision for polygonal chains representing baselines [9].

5.2 Evaluation Protocol

For model training, an AdamW optimizer is employed with a learning rate of $3e - 4$, and models are trained or fine-tuned for 100 epochs using a batch size of eight. An early stopping criterion is implemented to prevent overfitting, halting training after 10 iterations and fine-tuning after 30 iterations without improvement in validation performance. Model selection for final evaluation is based on the highest validation mean IoU achieved during each run.

Regarding evaluation across datasets, we report the average performance obtained from training with at least three random seeds for U-DIADS-TL. Conversely, for the DIVA-HisDB dataset, performance is reported from the model that achieves the highest validation performance, consistent with its competition setting. Furthermore, specific to the DIVA-HisDB dataset, a post-processing pipeline is applied to facilitate comparison with ICDAR competition results from 2017 [22], which include filtering out segments shorter than 50 pixels and subsequently merging lines that are both geometrically close (within a 50-pixel distance threshold) and similarly oriented (within a 15-degree angle threshold).

Name	Pixel IoU	Line IoU	DR	RA	FM
Baseline [31]	49.0	44.3	29.1	17.7	21.7
DeepLabV3+	52.44	51.90	48.32	30.00	36.44
+ UNet++	64.50	70.06	66.97	41.83	50.77
+ Lighter Encoder	65.29	71.24	68.30	42.39	51.94
+ Connectivity Loss	67.83	75.17	72.07	46.36	55.66
Final Model (Tuned)	70.64	77.38	73.90	52.96	61.40

Table 1: Ablation Study of the Proposed Method.

5.3 Ablation

In the following section, we will ground our proposed method by conducting an ablation study on the individual components. Each component yields significant performance improvements, which are denoted in Table 1, which summarizes the results of the performed study.

Model Architectures First, we ablate different model architectures. For architecture selection, we considered the same as Zottin et al. [31] plus a transformer-based one; however, this time, training on patches instead of the whole image this time. Inspecting the results in Table 2, we observe that simpler architectures, e.g., the UNet [19] or derived variants, i.e., the UNet++ [28], yield drastically better performance over all metrics. In line with the theory that less bias reduces the risk of overfitting [11], we select the UNet++ given our context of data-scarce training regimes in the few-shot setting.

Name	Pixel IoU	Line IoU	DR	RA	FM
PSPNet[27]	42.22	44.08	37.31	33.33	34.18
SegFormer [26]	50.10	49.27	46.42	29.66	35.50
UNet++ [28]	64.50	70.06	66.97	41.83	50.77
DeepLabV3+ [4]	52.44	51.90	48.32	30.00	36.44
UNet [19]	64.34	69.56	66.71	39.92	49.31

Table 2: Ablating Architectures.

Encoder Parameter Size Continuing with the previous made insight, we ablate encoder depth, evaluating ResNets [13] pre-trained on ImageNet [5] with varying amounts of network layers. We illustrate our findings in Fig. 4, which indicates a similar trend: performance tends to increase with decreasing depth. However, this trend is not stable as there is some fluctuation. For instance, a ResNet with a depth of 18 performs clearly worse than one with 34 layers, while being better than larger variants with up to 152 layers.






Fig. 4: Ablation on Network Depth.

Connectivity Loss An essential component of our proposed methodology for few-shot text line segmentation is the tuned connectivity-aware loss function proposed by Grim et al. [12] and leveraged within our methodology. As the hyperparameters α and β influence the trade-off between pixel-level and topological loss and between penalizing false merge and false split errors, we perform an ablation study on these parameters to select them optimally.

We illustrate our findings in Fig. 5 and observe in Fig. 5a that performance is best when considering only the topology of the mask (i.e., $\alpha = 1$), neglecting pixel-level performance, which is expected due to our additional incorporation of the Dice loss. In contrast to initial assumptions that both false merges and false splits play a crucial role in performance, results of Fig. 5b show that only false merges are problematic (i.e., $\beta = 0$). Given these results, we performed the same ablation for the DIVA-HisDB baseline detection task and found $\beta = 0$ to also work best.



(a) Varying α , while keeping β constant at 0.5.



(b) Varying β , while keeping α constant at 1.

Fig. 5: Ablating the Hyper-Parameters of the Connectivity Loss.

5.4 Results

Our comprehensive evaluation on two distinct historical document datasets, DIVA-HisDB and U-DIADS-TL, unequivocally demonstrates the competitive performance and generalizability of our proposed few-shot methodology in extremely scarce data settings of only three available annotated pages. Summarized in Table 3 and Table 4, our results highlight the impact of the connectivity-aware

loss, which we propose for text line segmentation and baseline detection, in addition to our use of smaller and simpler architectures compared to previous SOTA methods.

Name	Samples	CB55	CSG18	CSG863	Avg. F1
System-8 [22]		98.96	98.53	97.16	98.22
System-2 [22]	20	95.97	98.79	98.30	97.68
UNet++		99.66	94.31	92.24	95.40
UNet++		93.17	83.17	86.62	87.65
+Connectivity Loss	<u>3</u>	96.07	87.27	86.02	89.79
+Post-Processing		<u>99.00</u>	<u>94.13</u>	<u>93.62</u>	<u>95.58</u>

Table 3: Final Results DIVA-HisDB.

In more detail, on the DIVA-HisDB dataset with the task of baseline detection (Table 3), our suggested lightweight UNet-based methodology, trained on patches with only three annotated pages, achieves an average F1 score of 95.58%. In comparison to the UNet++ baseline trained on a significantly larger dataset (20 samples versus 3, 95.40% F1) and approaching or even outperforming the F1 scores of top competition systems (System-8: 98.22%, System-2: 97.68%) that leveraged considerably more training data, the efficacy of our methodology in the few-shot setting is underscored.

Metric	Latin 14396	Latin 2	Syriaque 341	Avg. Ours	Avg. [31]	Rel. Gain to [31]
Pixel IoU	73.71	77.83	60.36	70.64	49.0	+44.16
Line IoU	79.78	86.28	66.09	77.38	44.3	+74.67
DR	79.45	84.37	57.89	73.90	29.1	+153.95
RA	51.20	65.40	42.27	52.96	17.7	+199.21
F-Measure	61.96	73.43	48.83	61.40	21.7	+182.95

Table 4: Final Results U-DIADS-TL.

On the more challenging task of text line segmentation of the U-DIADS-TL dataset (Table 4), the impact of our methodology is further exemplified, where our final, tuned model sets a new SOTA for few-shot text line segmentation. Across all metrics, our approach achieves relative gains ranging from 75% to 200%. Specifically, we observe a nearly 200% gain in RA and an approximate 185% gain in FM. Complementary improvements include a plus of around 75% in Line IoU (reaching 77.38%), an over 150% relative gain in DR (attaining 73.90%), and a nearly 45% relative improvement in Pixel IoU, with an absolute performance of 70.64%. An excerpt of qualitative results is depicted in Fig. 1.

6 Conclusion

In this work, we challenged the prevailing trend of increasing model complexity for few-shot text line segmentation in historical documents. We demonstrated that a simpler, lightweight UNet++ architecture, when paired with a connectivity-aware loss function, yielded superior performance in a low-data regime. Our approach directly addressed the critical issue of maintaining topological correctness by penalizing common structural errors such as line fragmentation and inadvertent merges.

Our experimental results substantiated the effectiveness of this methodology. On the U-DIADS-TL benchmark, our model established a new state-of-the-art, improving the Line IoU by over 75% and the RA by nearly 200% relative to the baseline. Furthermore, it showed strong generalization capabilities on the DIVA-HisDB dataset, producing results competitive with top-performing systems while utilizing 85% less training data. These findings confirmed our central hypothesis: architectural simplicity, when combined with a topologically informed loss function, provides a more potent and efficient solution for few-shot segmentation than more complex, data-hungry models.

References

1. Alberti, M., Vögtlin, L., Pondenkandath, V., Seuret, M., Ingold, R., Liwicki, M.: Labeling, cutting, grouping: an efficient text line segmentation method for medieval manuscripts. In: 2019 International Conference on Document Analysis and Recognition (ICDAR). pp. 1200–1206. IEEE (2019)
2. AVML Lab Fest Organizers: FEST @ ICDAR 2025: Competition on few-shot text line segmentation of ancient handwritten documents (2025), accessed: 2025-06-19
3. Barakat, B., Droby, A., Kassis, M., El-Sana, J.: Text line segmentation for challenging handwritten document images using fully convolutional network. In: 2018 16th International Conference on Frontiers in Handwriting Recognition (ICFHR). pp. 374–379. IEEE (2018)
4. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(4), 834–848 (Apr 2018). <https://doi.org/10.1109/tpami.2017.2699184>
5. De Nardin, A., Zottin, S., Colombi, E., Piciarelli, C., Foresti, G.L.: Is imagenet always the best option? An overview on transfer learning strategies for document layout analysis. In: International Conference on Image Analysis and Processing. pp. 489–499. Springer (2023)
6. De Nardin, A., Zottin, S., Paier, M., Foresti, G.L., Colombi, E., Piciarelli, C.: Efficient few-shot learning for pixel-precise handwritten document layout analysis. In: 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). IEEE (Jan 2023)
7. De Nardin, A., Zottin, S., Piciarelli, C., Colombi, E., Foresti, G.L.: Few-shot pixel-precise document layout segmentation via dynamic instance generation and local thresholding. *International Journal of Neural Systems* **33**(10) (Aug 2023)

8. De Nardin, A., Zottin, S., Piciarelli, C., Colombi, E., Foresti, G.L.: A one-shot learning approach to document layout segmentation of ancient arabic manuscripts. In: 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). IEEE (Jan 2024)
9. Diem, M., Kleber, F., Fiel, S., Gruning, T., Gatos, B.: cBAD: ICDAR2017 Competition on Baseline Detection. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR). p. 1355–1360. IEEE (Nov 2017). <https://doi.org/10.1109/icdar.2017.222>
10. Fizaine, F.C., Bard, P., Paindavoine, M., Robin, C., Bouyé, E., Lefèvre, R., Vinter, A.: Historical text line segmentation using deep learning algorithms: Mask-RCNN against U-Net networks. *Journal of Imaging* **10**(3), 65 (2024)
11. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press (2016)
12. Grim, A., Chandrashekar, J., Sümbül, U.: Efficient connectivity-preserving instance segmentation with supervoxel-based loss function. *Proceedings of the AAAI Conference on Artificial Intelligence* **39**(3), 3167–3175 (Apr 2025). <https://doi.org/10.1609/aaai.v39i3.32326>
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). p. 770–778. IEEE (Jun 2016)
14. Likforman-Sulem, L., Hanimyan, A., Faure, C.: A Hough-based algorithm for extracting text lines in handwritten documents. In: *Proceedings of 3rd international conference on document analysis and recognition*. vol. 2, pp. 774–777. IEEE (1995)
15. Likforman-Sulem, L., Zahour, A., Taconet, B.: Text line segmentation of historical documents: a survey. *International Journal of Document Analysis and Recognition (IJDAR)* **9**, 123–138 (2007)
16. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3431–3440 (2015)
17. Mosinska, A., Marquez-Neila, P., Koziński, M., Fua, P.: Beyond the pixel-wise loss for topology-aware delineation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3136–3145 (2018)
18. O’Gorman, L.: The document spectrum for page layout analysis. *IEEE Transactions on pattern analysis and machine intelligence* **15**(11), 1162–1173 (1993)
19. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 234–241. Springer International Publishing, Cham (2015)
20. Salehi, S.S.M., Erdogmus, D., Gholipour, A.: Tversky loss function for image segmentation using 3d fully convolutional deep networks. In: *International workshop on machine learning in medical imaging*. pp. 379–387. Springer (2017)
21. Shapiro, V., Gluhchev, G., Sgurev, V.: Handwritten document image segmentation and analysis. *Pattern Recognition Letters* **14**(1), 71–78 (1993)
22. Simistira, F., Bouillon, M., Seuret, M., Würsch, M., Alberti, M., Ingold, R., Liwicki, M.: ICDAR2017 Competition on Layout Analysis for Challenging Medieval Manuscripts. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR). vol. 01, pp. 1361–1370 (2017). <https://doi.org/10.1109/ICDAR.2017.223>
23. Sterzinger, R., Brenner, S., Sablatnig, R.: Drawing the Line: Deep Segmentation for Extracting Art from Ancient Etruscan Mirrors, p. 39–56. Springer Nature Switzerland (2024). https://doi.org/10.1007/978-3-031-70543-4_3

24. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations, p. 240–248. Springer International Publishing (2017). https://doi.org/10.1007/978-3-319-67558-9_28
25. Vadlamudi, N., Krishna, R., Sarvadevabhatla, R.K.: Seamformer: High precision text line segmentation for handwritten documents. In: International Conference on Document Analysis and Recognition. pp. 313–331. Springer (2023)
26. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems. vol. 34, pp. 12077–12090. Curran Associates, Inc. (2021)
27. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (Jul 2017). <https://doi.org/10.1109/cvpr.2017.660>
28. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: UNet++: A Nested U-Net Architecture for Medical Image Segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, pp. 3–11. Springer (2018)
29. Zottin, S., De Nardin, A., Colombi, E., Piciarelli, C., Pavan, F., Foresti, G.L.: U-DIADS-Bib: a full and few-shot pixel-precise dataset for document layout analysis of ancient manuscripts. *Neural Computing and Applications* **36**(20), 11777–11789 (Jan 2024). <https://doi.org/10.1007/s00521-023-09356-5>
30. Zottin, S., De Nardin, A., Foresti, G.L., Colombi, E., Piciarelli, C.: ICDAR 2024 Competition on Few-Shot and Many-Shot Layout Segmentation of Ancient Manuscripts (SAM). In: International Conference on Document Analysis and Recognition. pp. 315–331. Springer (2024)
31. Zottin, S., Nardin, A.D., Branca, G., Colombi, E., Piciarelli, C., Shujat, H., Foresti, G.L.: Exploring few-shot text line segmentation approaches in challenging ancient manuscripts. In: Cornia, M., Nunzio, G.M.D., Firmani, D., Mizzaro, S., Serra, G., Tonelli, S., Tremamunno, A. (eds.) Proceedings of the 21st Conference on Information and Research science Connecting to Digital and Library science, Udine, Italy, February 20-21, 2025. CEUR Workshop Proceedings, vol. 3937. CEUR-WS.org (2025)