

Single-to-mix Modality Alignment with Multimodal Large Language Model for Document Image Machine Translation

Yupu Liang^{1,2}, Yaping Zhang^{1,2}*, Zhiyang Zhang^{1,2},
Yang Zhao^{1,2}, Lu Xiang^{1,2}, Chengqing Zong^{1,2}, Yu Zhou^{1,3}

¹ State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS),
Institute of Automation, Chinese Academy of Sciences, Beijing, China

² School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³ Fanyu AI Laboratory, Zhongke Fanyu Technology Co., Ltd, Beijing, China

{liangyupu2021, zhangzhiyang2020}@ia.ac.cn, {yaping.zhang, yang.zhao, lu.xiang, cqzong, yzhou}@nlpr.ia.ac.cn

Abstract

Document Image Machine Translation (DIMIT) aims to translate text within document images, facing generalization challenges due to limited training data and the complex interplay between visual and textual information. To address these challenges, we introduce M4Doc, a novel single-to-mix modality alignment framework leveraging Multimodal Large Language Models (MLLMs). M4Doc aligns an image-only encoder with the multimodal representations of an MLLM, pre-trained on large-scale document image datasets. This alignment enables a lightweight DIMIT model to learn crucial visual-textual correlations during training. During inference, M4Doc bypasses the MLLM, maintaining computational efficiency while benefiting from its multimodal knowledge. Comprehensive experiments demonstrate substantial improvements in translation quality, especially in cross-domain generalization and challenging document image scenarios.¹

1 Introduction

Document Image Machine Translation (DIMIT) aims to translate text within document images from one language to another while preserving the logical layout (Liang et al., 2024). With vast amounts of information stored in document images (e.g., academic papers, magazines, scanned documents, Figure 1), DIMIT has gained increasing attention as a critical sub-task of visual document understanding in the era of multimodal large language models (Ye et al., 2023; Zhang et al., 2023a; Hu et al., 2024; Yu et al., 2024).

Recent advancements in DIMIT can be categorized into two primary approaches: (1) Cascade systems (Hinami et al., 2021; Sable et al., 2023; Yao, 2023; Zhang et al., 2023c), which employ multiple models sequentially and encounter issues such

* Corresponding author.

¹Our code is available at: <https://github.com/liangyupu/M4Doc>

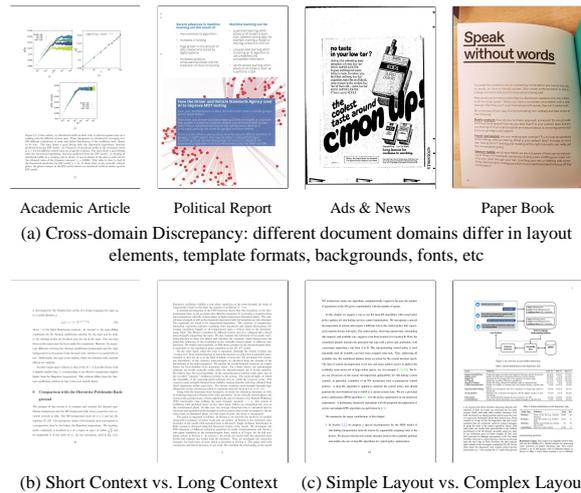


Figure 1: Different test scenarios of end-to-end DIMIT.

		BLEU	BLEU-PT
1	DoTA Test Set	38.68	42.34
2	Cross-domain	12.64	15.03
3	Long Context	34.85	33.73
4	Complex Layout	30.30	35.16

Table 1: Scores of the end-to-end DIMIT model on different test scenarios. The model is trained on the DoTA dataset. *Cross-domain*, *Long Context*, and *Complex Layout* mean testing on the DITrans political report subset, the DoTA long context subset, and the complex layout subset, separately.

as structural redundancy, error propagation, and high latency. (2) End-to-end methods (Jain et al., 2021; Ma et al., 2022; Zhu et al., 2023; Zhang et al., 2023b; Liang et al., 2024; Zhang et al., 2025b,a), which streamline the process by optimizing a unified training objective, thus enhancing structural efficiency.

However, both cascade and end-to-end approaches are hindered by the lack of large and diverse DIMIT datasets, which limits their ability to generalize to new types of documents. This limited generalization is evident in performance

drops across several key scenarios, as shown in Table 1: (1) **cross-domain generalization**: the model, trained on the DoTA dataset (Liang et al., 2024), achieves a BLEU score of 38.68 on the original test set, but only 12.64 on the DITrans Political Report test set (Zhang et al., 2023b) in a cross-domain zero-shot scenario, which includes document images with varying layouts, fonts, and background. (2) **long context generalization**: there is a decrease of 3.83 BLEU score when testing on the same dataset but with a long context subset, containing document images with more than 750 English words. (3) **complex layout generalization**: the model’s performance drops by 8.38 BLEU when the test set includes more images with complex layouts.²

Recent Multimodal Large Language Models (MLLMs), pre-trained on extensive datasets of images and text, have demonstrated impressive generalization across various domains, contexts, and layouts (Ye et al., 2023; Zhang et al., 2023a; Hu et al., 2024; Yu et al., 2024). While MLLMs hold great promise for DIMT (Liu et al., 2024a,b), their large size and computational demands make them difficult to use directly, especially in resource-constrained environments.³

To address these limitations, we propose **M4Doc** (single-to-mix Modality alignment with Multimodal large language Model for Document image Machine translation), a novel framework that leverages the strong generalization capabilities of MLLMs to enhance the performance and efficiency of smaller DIMT models through a single-to-mix modality alignment strategy. This strategy aligns an image-only representation with the MLLM’s rich multimodal representations, effectively transferring knowledge from the MLLM. Specifically, a novel single-to-mix modality alignment encoder is designed as a bridge to connect the MLLM and the DIMT model. This encoder with **only image** input learns to align with the mix-modality representation of MLLM, using both **image and text** as inputs. The alignment encoder can serve as an alternative to the MLLM and provide mix-modality information to the DIMT model in the inference stage. A major advantage of this approach is that it requires aligning the DIMT model with the MLLM only during training, allowing the use of a smaller model during inference, which achieves the trade-off between performance and

inference speed.

Our contributions are summarized as follows:

- A novel method, M4Doc, has been proposed, which uses the pre-trained knowledge of the MLLM to assist a small DIMT model in the training stage and achieves the trade-off between translation quality and inference speed.
- A new approach, single-to-mix alignment, has been developed, which only takes images as input and aligns with the mix-modality representation of the MLLM.
- Extensive experiments demonstrate the effectiveness of the proposed method, and the DIMT model’s performances on cross-domain, long context, and complex layout scenarios are also improved.

2 M4Doc Method

In this section, we introduce M4Doc, a novel single-to-mix modality alignment framework designed to enhance DIMT by leveraging the MLLMs. The model architecture of M4Doc is illustrated in Figure 2. The key idea of M4Doc is to align the representations of an image-only encoder with the rich multimodal representations of an MLLM during training, enabling a lightweight DIMT model to effectively capture the interplay between textual and visual features. The whole model contains an MLLM, an alignment encoder, an image encoder, and a translation decoder. In the training stage, the alignment encoder simultaneously learns to align with the MLLM and provides mix-modality information to the translation decoder. In the inference stage, the alignment encoder serves as an alternative to the MLLM and continues providing mix-modality information with only image input.

2.1 Mix-modality Representation Extraction

The MLLM acts as a guide for the alignment encoder to provide mix-modality information with image and text inputs. We input the image $I \in \mathbb{R}^{H \times W \times 3}$ and corresponding ground truth source language text $\mathbf{X} = \{x_1, x_2, \dots, x_m\}$ into the MLLM. The input format is <System Prompt> <Image Token> <User Prompt> <Source Text>, which is the same as the format used in the MLLM pre-training. <System Prompt> and <User Prompt> are also the same as those used by the MLLM in the OCR task. <Source Text> is the

²The criteria for complex layouts is in Appendix A.1.

³More details on fine-tuning MLLMs for DIMT are in Appendix B.3.

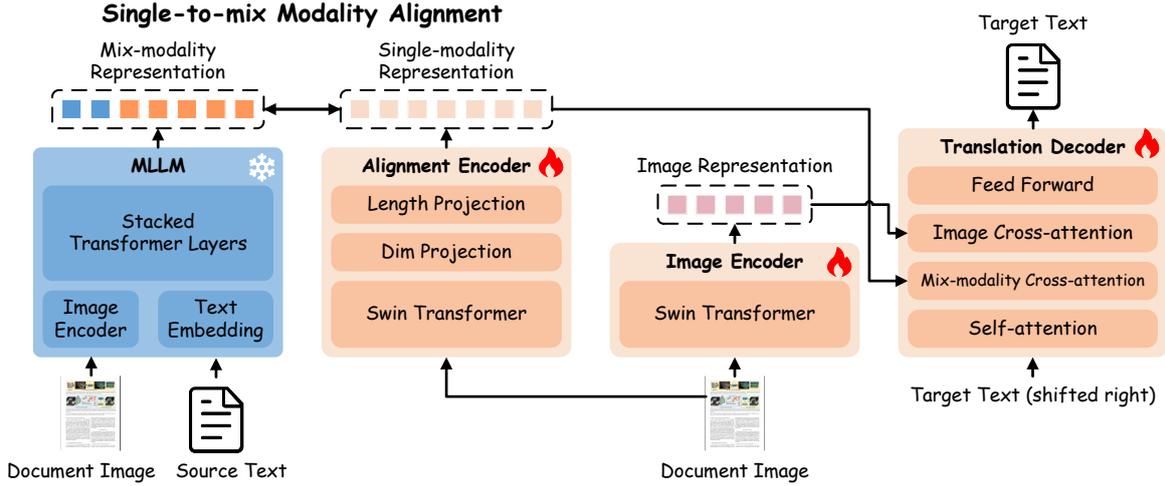


Figure 2: The diagram of the proposed M4Doc. During training, the alignment encoder learns to align with the MLLM’s mix-modality representation with single-modality input. The MLLM is frozen, while the other modules remain trainable. During inference, the MLLM is discarded for faster inference speed, while the alignment encoder provides aligned mix-modality information to guide the translation decoder.

ground truth OCR text of the corresponding image. We can get the mix-modality representation⁴ $\mathbf{H}_{\text{MLLM}} \in \mathbb{R}^{l_{\text{MLLM}} \times d_{\text{MLLM}}}$, which can be formulated as:

$$\mathbf{H}_{\text{MLLM}} = \text{MLLM}(\mathbf{I}, \mathbf{X}) \quad (1)$$

where l_{MLLM} and d_{MLLM} are the sequence length and dimension of the MLLM.

2.2 Single-to-mix Modality Alignment

The alignment encoder bridges the gap between the single-modality (image-only) input and the mix-modality (image and text) representations of the MLLM. Using a pre-trained Swin Transformer (Blecher et al., 2024), the alignment encoder extracts visual features $\mathbf{H}_{\text{Swin}} \in \mathbb{R}^{l_{\text{Swin}} \times d_{\text{Swin}}}$ from the image \mathbf{I} :

$$\mathbf{H}_{\text{Swin}} = \text{Swin}(\mathbf{I}) \quad (2)$$

To match the dimensions of the MLLM output, two Feed Forward Networks (FFNs) are used to project \mathbf{H}_{Swin} to $\mathbf{H}_{\text{Align}} \in \mathbb{R}^{l_{\text{MLLM}} \times d_{\text{MLLM}}}$:

$$\mathbf{H}_{\text{Align}} = \text{FFN}_{\text{length}}(\text{FFN}_{\text{dim}}(\mathbf{H}_{\text{Swin}})^T)^T \quad (3)$$

After this process, an alignment loss guides the alignment encoder to mimic the mix-modality representation $\mathbf{H}_{\text{Align}} \in \mathbb{R}^{l_{\text{MLLM}} \times d_{\text{MLLM}}}$ ⁵:

$$\mathcal{L}_{\text{align}} = 1 - \text{Cos}(\mathbf{H}_{\text{MLLM}}, \mathbf{H}_{\text{Align}}) \quad (4)$$

where Cos is the cosine similarity of two tensors.

⁴The last layer’s output hidden states of the MLLM.

⁵The effect of different alignment loss functions can be found in Appendix B.1.

2.3 Aligned Mix-modality Guided Translation

The image encoder encodes the input image \mathbf{I} to its semantic representation $\mathbf{H}_{\text{Image}} \in \mathbb{R}^{l_{\text{Image}} \times d_{\text{Image}}}$. We also use a pre-trained Swin Transformer (Blecher et al., 2024) to construct the image encoder. $\mathbf{H}_{\text{Image}}$ is calculated as follows:

$$\mathbf{H}_{\text{Image}} = \text{Encoder}_{\text{Image}}(\mathbf{I}) \quad (5)$$

where l_{Image} is the number of output vectors and d_{Image} is the vectors’ dimension.

The translation decoder is aimed to generate target language text under the guidance of the alignment encoder and image encoder. We modify the vanilla Transformer’s decoder (Vaswani et al., 2017) by incorporating a mix-modality cross-attention module and an image cross-attention module in each layer to receive representations from the alignment encoder and the image encoder. At each decoding timestep t , the translation decoder takes $\mathbf{H}_{\text{Align}}$, $\mathbf{H}_{\text{Image}}$ and generated target tokens $y_{<t} = \{y_1, y_2, \dots, y_{t-1}\}$ as input and outputs the probability distribution of next target token y_t . This process can be defined as:

$$p(y_t | y_{<t}, \mathbf{I}, \mathbf{X}) = \text{Decoder}(y_{<t}, \mathbf{H}_{\text{Align}}, \mathbf{H}_{\text{Image}}) \quad (6)$$

where $\mathbf{H}_{\text{Align}}$ and $\mathbf{H}_{\text{Image}}$ both need to be converted to the same dimension of the decoder through two FFNs which are not shown in Figure 2 for simplicity.

The translation loss is as follows:

$$\mathcal{L}_{\text{trans}} = - \sum_{t=1}^n \log p(y_t | y_{<t}, \mathbf{I}, \mathbf{X}; \theta) \quad (7)$$

where θ is the parameters of the alignment encoder, image encoder, and translation decoder.

2.4 Training & Inference Strategy

In the training stage, the Swin Transformer modules of the alignment encoder and image encoder are initialized from the pre-trained OCR model’s encoder. The translation decoder’s FFN, image cross-attention, and self-attention modules are initialized from the pre-trained text translation model’s decoder. The other parts are randomly initialized. The parameters of the MLLM are frozen.

The total loss of M4Doc is as follows:

$$\mathcal{L} = \alpha \times \mathcal{L}_{\text{align}} + \mathcal{L}_{\text{trans}} \quad (8)$$

where α is a hyperparameter.⁶

In the inference stage, as shown in Figure 2, only the alignment encoder, image encoder, and translation decoder are involved, which contain much fewer parameters compared with the MLLM. Furthermore, due to the introduction of the alignment encoder aligning with the MLLM during the training stage, the entire model maintains high translation quality while achieving fast inference speed.

3 Experiments

3.1 Dataset & Metrics

Our models are comprehensively evaluated on two public benchmarks DoTA (Liang et al., 2024) and DITrans (Zhang et al., 2023c), under academic article and political report scenarios. Detailed dataset setting can be seen in Appendix A.1.

We thoroughly evaluate the models’ capabilities in three aspects: (1) **full-text translation**, which means the translation quality of all the text in the image - BLEU and COMET (Rei et al., 2020). (2) **plain-text translation**, which means the translation quality of the text after removing formulas and tables - BLEU-PT. (3) **structure preserving**, which means the model’s ability to restore the layout structure of the document images - STEDS (Structure Tree-Edit-Distance-based Similarity).

We calculate BLEU, BLEU-PT, and STEDS the same as Liang et al. (2024). For COMET calculation, due to the original COMET’s inability to process long texts, we first used Trankit (Nguyen et al., 2021) to segment the source and translated texts into sentences, then used Sentalign (Stein-grimsson et al., 2023) for sentence-level alignment,

⁶The effect of different hyperparameters can be found in Appendix B.2.

and finally calculated the average of COMET score in reference-free mode.⁷

3.2 Settings

Pre-trained Models Selection For the MLLM, we select four MLLMs with different numbers of parameters and training data: Vary-toy (Wei et al., 2024), Vary-base (Wei et al., 2023), Llava-next (Liu et al., 2024a) and Textmonkey (Liu et al., 2024b). The Swin Transformers of alignment encoder and image encoder are initialized from the encoder of pre-trained OCR model Nougat (Blecher et al., 2024). We follow the vanilla Transformer-base (Vaswani et al., 2017) setting, pre-train an English-Chinese translation model on UN Corpus En-Zh (Ziemski et al., 2016), and use the pre-trained decoder to initialize the translation decoder in M4Doc.

Other Settings The hyperparameter α is set to 1.0. During training, we use the Adam optimizer and employ a linear decay learning rate schedule with a learning rate of 5e-5. The maximum number of training steps is 15K and the batch size is 64. More detailed settings are in Appendix A.2.

3.3 Baselines

We evaluate our method against diverse baselines, including text-only, cascade, end-to-end, and knowledge distillation methods, to comprehensively assess its performance and validate its effectiveness.

Text-only MT (Vaswani et al., 2017) We use the DoTA dataset to fine-tune the Transformer-base model pre-trained on UN Corpus En-Zh (Ziemski et al., 2016).

3.3.1 Cascade Baselines

LARDIT (Zhang et al., 2023c) This cascade system employs a layout analysis model (Yao, 2023), an **OCR tool**, and a text-only MT, sequentially.

Nougat-trans We utilize the Nougat model (Blecher et al., 2024) for combined layout analysis and OCR and the text-only MT is employed for translation.

3.3.2 End-to-End Baselines

We evaluate the existing end-to-end methods under two distinct settings: **Document-level** and **Text-line-level**. The specific end-to-end models evaluated are:

⁷The COMET model we used is [wmt22-comet-da](#).

		DoTA				DITrans				# Params	Time
		B	C	BP	S	B	C	BP	S	(M)	(s/page)
1	Text-only MT	47.61	67.51	54.16	92.89	21.50	48.76	22.55	86.96	99.5	8.81
Cascade Baselines											
2	LARDIT	35.58	54.48	41.75	75.83	14.66	30.16	16.58	57.77	$99.5 + \theta_1$	12.46
3	Nougat-trans	43.37	65.25	50.79	88.16	18.39	35.80	19.21	52.12	346.9	17.03
End-to-end TIMT Baselines (Document-level)											
4	ItNet	3.84	21.94	2.27	48.46	1.64	21.52	1.71	41.63	97.5	8.43
5	E2ETIT	1.51	20.80	1.69	32.90	2.71	23.45	2.83	40.53	122.0	8.19
6	PEIT	5.81	24.98	4.52	55.79	4.13	21.98	4.21	41.59	135.1	2.57
End-to-end TIMT Baselines (Text-line-level)											
7	ItNet	21.75	43.29	23.52	75.83	6.16	28.82	8.77	57.77	$97.5 + \theta_2$	7.20
8	E2ETIT	17.42	38.25	17.74	75.83	6.72	28.55	7.81	57.77	$122.0 + \theta_2$	7.59
9	PEIT	27.43	44.08	31.29	75.83	9.08	26.18	9.38	57.77	$135.1 + \theta_2$	2.42
Knowledge Distillation Baselines											
10	Seq-KD	34.42	53.54	36.63	82.51	10.58	25.89	11.38	56.92	212.4	9.76
11	MTKD	37.32	60.32	39.96	82.28	13.24	29.33	15.33	59.58	212.4	9.56
12	RD (Original)	5.13	23.86	3.85	53.06	0.53	24.37	0.56	40.07	212.4	8.38
13	RD (Trans)	31.05	48.16	32.00	77.62	9.31	22.69	9.72	58.24	212.4	9.86
End-to-end DIMT (Document-level)											
14	Base	37.60	61.52	40.85	83.08	11.91	30.59	14.00	52.89	127.6	9.16
15	DIMTDA	38.68	61.30	42.34	84.44	12.64	32.30	15.03	60.86	242.6	9.82
16	M4Doc (Vary-toy)	39.95	62.78	42.33	83.97	14.79	32.03	18.67	53.73	212.4	9.61
17	M4Doc (Vary-base)	41.22	63.10	42.09	86.06	14.52	30.53	16.55	55.89	215.6	9.43
18	M4Doc (Llava-next)	34.36	57.88	37.60	82.67	11.03	30.79	12.58	57.81	216.8	9.96
19	M4Doc (Textmonkey)	42.98	65.41	44.92	86.69	18.18	35.27	19.82	59.98	215.6	9.52

Table 2: Results on DoTA and DITrans English-Chinese test set. The models are trained on DoTA, and tested on DoTA and DITrans. **B**, **C**, **BP**, and **S** represent BLEU, COMET, BLEU-PT, and STEDS, respectively. **# Params** is the number of parameters of the model during inference. **Time** is the average inference time on a single NVIDIA V100 GPU. θ_1 denotes the parameters of the layout analysis model and OCR model. θ_2 denotes the parameters of the parameters of the layout analysis model and sentence splitting model. The **bold numbers** represent the best performance of the end-to-end DIMT.

Base This baseline end-to-end DIMT model uses the same image encoder and translation decoder architecture as M4Doc, without incorporating an alignment encoder or multimodal knowledge transfer.

DIMTDA (Liang et al., 2024) This end-to-end DIMT model uses a model assembler to integrate multiple pre-trained models to enhance the understanding of layout and translation capabilities.

ItNet (Jain et al., 2021) This end-to-end Text Image Machine Translation (TIMT) system first pre-trains a vanilla Transformer on a text parallel dataset. The combination of the image encoder and pre-trained decoder is fine-tuned.

E2ETIT (Ma et al., 2022) This end-to-end TIMT model uses a TPSNet and a ResNet as an image encoder combined with a Transformer decoder and utilizes text translation as an auxiliary task.

PEIT (Zhu et al., 2023) This end-to-end TIMT system employs a vision-text representation aligner

and a cross-model regularize to bridge the modality gap between visual inputs and textual inputs.

3.3.3 Knowledge Distillation Baselines

We conduct experiments to compare our method with three different knowledge distillation methods.

Seq-KD (Kim and Rush, 2016) This is the vanilla sequence-level knowledge distillation method for machine translation.

MTKD (Ma et al., 2023c) This method employs a pre-trained OCR model and a pre-trained machine translation model as teacher models, with a TIMT model serving as the student model.

RD (Zhu et al., 2024) This approach leverages an LLM, based on the OCR results of document images, to generate rationales, subsequently employing these rationales to train a document understanding model. As the original RD method performs poorly, we mix the generated rationales with the translation data from the DoTA dataset during training, resulting in the RD (Trans) method.

4 Results & Analysis

4.1 Main Results

Table 2 reports the performance of all methods. It can be observed that M4Doc outperforms the cascade methods LARDIT (line 2 vs. 19) by 7.40 BLEU, 10.93 COMET and 10.86 STEDS scores on the DoTA test set. Besides, M4Doc also achieves comparable performance with Nougat-trans (line 3 vs. 19) on both DoTA and DITrans test sets, while the number of parameters of M4Doc is reduced by 37.8% compared to Nougat-trans and inference time decreases by 44.1%.

Moreover, our method outperforms all the end-to-end TIMT baselines on both document-level and text-line-level settings. As the TIMT models are designed for text-line-level images, the performances of all TIMT models under the text-line-level settings are better than document-level settings. However, M4Doc still surpasses the highest-performing TIMT model (line 9 vs. 19) by a margin of 15.55 BLEU on the DoTA test set and 9.10 BLEU on the DITrans test set.

By comparing line 11 and line 19, our method is superior to the end-to-end DIMT baseline in in-domain and cross-domain zero-shot settings. In the in-domain setting, there is an increase of 4.30 BLEU, 4.11 COMET, and 2.58 BLEU-PT scores. In the cross-domain zero-shot setting, our method outperforms DIMTDA by 5.54 BLEU, 2.97 COMET, and 4.79 BLEU-PT scores, which confirms introducing MLLMs as auxiliaries during training can enhance the model’s generalization abilities.

From the results presented in lines 10–13, M4Doc demonstrates superior performance compared to all knowledge distillation baselines. Furthermore, M4Doc surpasses the highest-performing baseline (lines 11 vs. 16) by 2.63 BLEU, 2.46 COMET, and 2.37 BLEU-PT scores on the DoTA test set.

From the results of line 16-19, as the number of parameters in the MLLM models increases, the DIMT model’s translation quality also generally improves. However, due to the difference in pre-training data between Llava-next and other MLLMs, MLLMs pre-trained on document images are more suitable for assisting in the training of the DIMT model.

		Ads & News		Political Report	
		BLEU	STEDS	BLEU	STEDS
1	DIMTDA	14.21	77.12	26.71	89.33
2	M4Doc (Vary-toy)	17.07	78.18	27.62	88.83
3	M4Doc (Vary-base)	21.30	80.67	31.71	90.76
4	M4Doc (Textmonkey)	24.28	82.05	34.26	91.06

Table 3: Results on DITrans English-Chinese test set after finetuning.

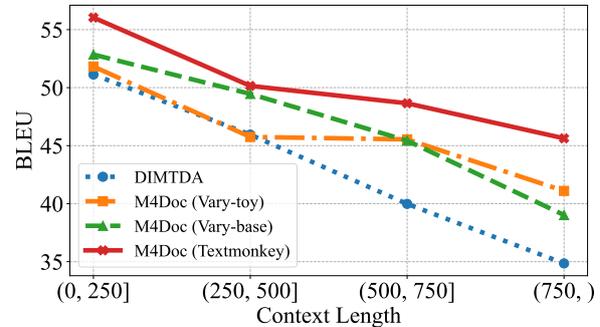


Figure 3: BLEU scores of M4Doc models testing on different context length valid sets. Detailed data can be seen in Appendix C.

4.2 Generalization Ability towards Difficult DIMT Scenarios

We pay special attention to challenging DIMT scenarios, where our model exhibits advantages through single-to-mix modality alignment. Consequently, we conduct three sets of experiments.

4.2.1 Cross-domain

We fine-tune end-to-end DIMT models on two subsets of the DITrans dataset separately after training on the DoTA dataset. Detailed dataset setting can be seen in Appendix A.1. The results are shown in Table 3. With the help of MLLM during training, all three variants of our method achieve better performance than DIMTDA. This could be because the MLLM is pre-trained on a large amount of data, allowing the alignment encoder to learn similar representations from the MLLM, thus enhancing the generalization capability of the DIMT model.

4.2.2 Long Context

We select samples from the valid set within different context lengths.⁸ Detailed settings can be seen in Appendix A.1. Results are shown in Figure 3. Our models outperform the baseline across all context length scenarios. The performance of all models decreases as the context length increases, but the decline is less pronounced with our mod-

⁸Context length refers to the number of English words in the image.

			BLEU	BLEU-PT	STEDS
1	Simple	DIMTDA	55.24	55.26	90.54
2		M4Doc	56.88	56.72	92.25
3	Complex	DIMTDA	30.30	35.16	84.57
4		M4Doc	35.88	41.24	83.76

Table 4: Results of different layout complexity on DoTA English-Chinese valid set.

			BLEU	BLEU-PT	STEDS
1	M4Doc (Vary-toy)		40.05	42.58	83.93
2	w/o \mathcal{L}_{align}		36.58	40.06	83.54
3	w/o Alignment Encoder		36.62	40.09	83.65
4	w MLLM Output		42.56	46.93	89.48

Table 5: Ablation study results on DoTA English-Chinese valid set.

els compared to DIMTDA, which indicates that the introducing of MLLM can improve the DIMT models’ ability to handle images with long context.

4.2.3 Complex Layout

We select two subsets (images with simple layout and complex layout) from the valid set of the DoTA dataset. Detailed settings can be seen in Appendix A.1. As shown in Table 4, our methods perform similarly to DIMTDA on images with simple layouts, but on images with complex layouts, our methods can achieve up to 5.58 BLEU and 6.08 BLEU-PT scores higher than DIMTDA. It suggests that the assistance of the MLLM during training can improve the DIMT models’ ability to understand complex layout structures and further improve the translation ability.

4.3 Ablation Study

4.3.1 Effect of Different Module

To investigate the effectiveness of the proposed modules, we conduct ablation experiments. The results are shown in Table 5.

w/o \mathcal{L}_{align} We remove the MLLM during training, keep the alignment encoder, and only use \mathcal{L}_{trans} to guide the model. By comparing line 1 and line 2, a decline of 3.47 BLEU and 2.52 BLEU-PT scores can be observed, which demonstrates the effectiveness of the MLLM during training.

w/o Alignment Encoder We remove the Swin Transformer in the alignment encoder and use the output of the image encoder to do alignment with MLLM and image encoding simultaneously. It can be seen from the comparison between line 1 and line 3 that simultaneously achieving alignment and

			BLEU	BLEU-PT	STEDS
1	M4Doc (Vary-toy)		40.05	42.58	83.93
2	w/o MLLM Image Input		38.77	42.00	84.99
3	w/o MLLM Text Input		37.06	38.69	84.98

Table 6: Results on DoTA English-Chinese valid set with different modalities input.

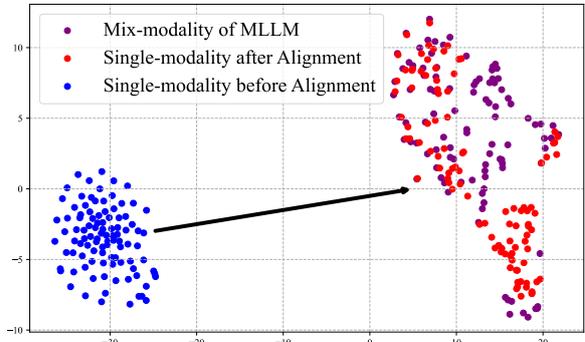


Figure 4: T-SNE visualization of different representations for MLLM and alignment encoder.

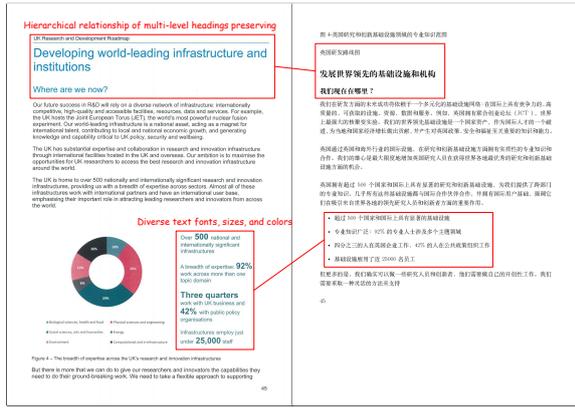
image encoding is challenging for a single encoder and causes a decrease in translation quality.

w MLLM Output The output hidden states of MLLM are directly sent to the translation decoder without the alignment encoder as an intermediary. By comparing line 1 and line 4, there is an increase of 2.51 BLEU and 5.55 STEDS scores. However, this approach significantly increases the parameters of the model ($\times 11.53$) and inference time ($\times 1.26$). Our method strikes a balance between translation quality and inference speed.

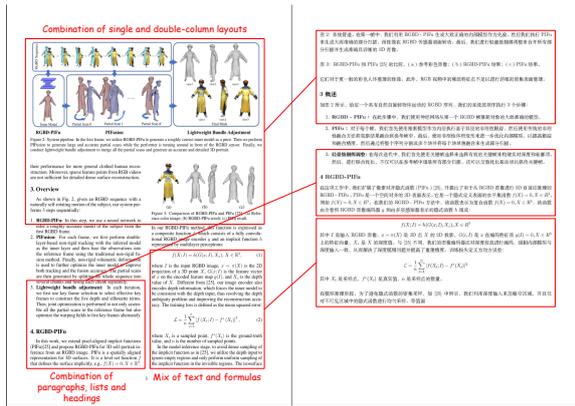
4.3.2 Effect of Mix-modality Input

To explore the impact of mix-modality input, we only send English text or the corresponding image to the MLLM during training. The input formats are <System Prompt> <Image Token> and <System Prompt> <Source Text>. As the results of Table 6 show, the performance degradation of the model is greater when text input is removed compared to when image input is removed. This may be because the source text contains more translation-relevant textual information. When MLLM has only image input, the alignment from image modality to image modality does not introduce additional information.

We provide a visualization of the representations of 100 samples in Figure 4. The single-modality representation output by the alignment encoder, after training, largely overlaps with the distribution of the mix-modality representation output by



(a) Cross-domain: Political Report



(b) Complex Layout

Figure 5: The output samples of M4Doc. For each image pair, the left is the input document image, and the right is the output translations in markdown format after rendering. It is better to zoom in for a clearer view. More samples can be seen in Appendix D.

the MLLM. This demonstrates that our proposed single-to-mix modality alignment allows the alignment encoder to effectively learn the MLLM outputs, providing additional information to guide the translation decoder in generating translations.

4.4 Case Study

We provide the output samples of M4Doc cross-domain and complex layout scenarios in Figure 5. More samples can be seen in Appendix D.

Figure 5 (a) is an image from the political report subset of the DITrans dataset. The fonts, sizes, and colors of the texts are diverse, which is quite different from the DoTA dataset used for training. After fine-tuning, our model can still perform translations and obtain the hierarchical relationship of multi-level headings and lists.

Figure 5 (b) comes from the DoTA dataset. The image contains a mix of text, figures, and formulas, with headings and lists, and a combination of

single and double-column layouts, resulting in a highly complex layout structure. Our model can still output the translation results in a logical order, formatted in Markdown.

5 Related Work

Text Image Machine Translation (TIMT) refers to translating texts from one language to another within images, as explored by Lan et al. (2023). In recent years, various end-to-end TIMT methods (Ma et al., 2023a,b,c, 2024a; Tian et al., 2023; Ma et al., 2024b; Lan et al., 2024; Qian et al., 2024; Guan et al., 2025) have been proposed. Jain et al. (2021) follows the encoder-decoder paradigm and uses a convolutional encoder and an autoregressive Transformer decoder to build the model. Ma et al. (2022) proposes a text translation enhanced text image translation method, which trains the end-to-end TIMT model with text translation as an auxiliary task. Zhu et al. (2023) introduces an end-to-end TIMT framework that bridges the modality gap with pre-trained models. While these end-to-end methods have demonstrated satisfactory performance, their effectiveness is limited to images with short context and simple layout structure, different from document images.

Recent advancements in MLLMs have significantly improved the processing and understanding of text-rich images (Ye et al., 2023; Wei et al., 2024; Hu et al., 2024; Yu et al., 2024). Wei et al. (2023) explores adding fine-grained vision perception for document images to the MLLM without affecting its existing natural image understanding capabilities. Zhang et al. (2023a) and Liu et al. (2024a) utilize GPT-4 (Yang et al., 2023) to construct a visual instruction tuning dataset and improve LLaVA’s (Liu et al., 2023) ability to comprehend textual detail within images. Liu et al. (2024b) proposes shifted window attention to achieve cross-window connectivity at higher input resolutions and token resampler to filter out significant tokens. As MLLMs take both images and texts as inputs during the pre-training stage, the integration of visual and linguistic information provides a better understanding of document images, which inspires us to leverage the MLLM for the DIMT task.

6 Conclusion

In this paper, we propose a novel method, single-to-mix modality alignment with multimodal large language model for document image machine transla-

tion (M4Doc), which has three advantages. Firstly, single-to-mix modality alignment allows the alignment encoder to infer more textual information from the image input. Secondly, the alignment with MLLM enhances the generalization capability towards three difficult DIMT scenarios. Finally, the introduction of the alignment encoder achieves the SOTA translation quality while preserving high inference efficiency. Extensive experiments demonstrate the effectiveness of M4Doc and highlight its advantage in enhancing the performance of the DIMT model in cross-domain and complex document image scenarios.

Limitations

Although M4Doc achieves notable results on the DIMT task, current end-to-end models generate the entire translated text of the document image in a single output. In the future, we plan to explore integrating user prompts to translate text in specific regions of the image, thereby making the translation more aligned with user preferences.

Acknowledgements

We thank anonymous reviewers for helpful suggestions. This work is supported by the National Natural Science Foundation of China (No. 62336008 and No. 62476275).

References

- Lukas Blecher, Guillem Cucurull, Thomas Scialom, and Robert Stojnic. 2024. *Nougat: Neural optical understanding for academic documents*. In *The Twelfth International Conference on Learning Representations*.
- Boyu Guan, Yining Zhang, Yang Zhao, and Chengqing Zong. 2025. *TriFine: A large-scale dataset of vision-audio-subtitle for tri-modal machine translation and benchmark with fine-grained annotated tags*. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 8215–8231, Abu Dhabi, UAE. Association for Computational Linguistics.
- Ryota Hinami, Shonosuke Ishiwatari, Kazuhiko Yasuda, and Yusuke Matsui. 2021. Towards fully automated manga translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12998–13008.
- Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang Zhang, Bo Zhang, Chen Li, Ji Zhang, Qin Jin, Fei Huang, et al. 2024. *mplug-docowl 1.5: Unified structure learning for ocr-free document understanding*. *arXiv preprint arXiv:2403.12895*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. *Lora: Low-rank adaptation of large language models*. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. *Gpt-4o system card*. *arXiv preprint arXiv:2410.21276*.
- Puneet Jain, Orhan Firat, Qi Ge, and Sihang Liang. 2021. *Image translation network*.
- Yoon Kim and Alexander M. Rush. 2016. *Sequence-level knowledge distillation*. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Zhibin Lan, Liqiang Niu, Fandong Meng, Jie Zhou, Min Zhang, and Jinsong Su. 2024. *Translatotron-v (ison): An end-to-end model for in-image machine translation*. *arXiv preprint arXiv:2407.02894*.
- Zhibin Lan, Jiawei Yu, Xiang Li, Wen Zhang, Jian Luan, Bin Wang, Degen Huang, and Jinsong Su. 2023. *Exploring better text image translation with multimodal codebook*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3479–3491, Toronto, Canada. Association for Computational Linguistics.
- Yupu Liang, Yaping Zhang, Cong Ma, Zhiyang Zhang, Yang Zhao, Lu Xiang, Chengqing Zong, and Yu Zhou. 2024. *Document image machine translation with dynamic multi-pre-trained models assembling*. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7084–7095, Mexico City, Mexico. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024a. *Llava-next: Improved reasoning, ocr, and world knowledge*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. *Visual instruction tuning*. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 34892–34916.
- Yuliang Liu, Biao Yang, Qiang Liu, Zhang Li, Zhiyin Ma, Shuo Zhang, and Xiang Bai. 2024b. *Textmonkey: An ocr-free large multimodal model for understanding document*. *arXiv preprint arXiv:2403.04473*.
- Cong Ma, Xu Han, Linghui Wu, Yaping Zhang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023a. *Modal contrastive learning based end-to-end text image machine translation*. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.

- Cong Ma, Yaping Zhang, Mei Tu, Xu Han, Linghui Wu, Yang Zhao, and Yu Zhou. 2022. Improving end-to-end text image translation from the auxiliary text translation task. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 1664–1670. IEEE.
- Cong Ma, Yaping Zhang, Mei Tu, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023b. E2timt: Efficient and effective modal adapter for text image machine translation. In *The 17th International Conference on Document Analysis and Recognition (ICDAR)*, pages 70–88.
- Cong Ma, Yaping Zhang, Mei Tu, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023c. Multi-teacher knowledge distillation for text image machine translation. In *The 17th International Conference on Document Analysis and Recognition (ICDAR)*, pages 484–501.
- Cong Ma, Yaping Zhang, Zhiyang Zhang, Yupu Liang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2024a. Born a BabyNet with hierarchical parental supervision for end-to-end text image machine translation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2468–2479, Torino, Italia. ELRA and ICCL.
- Cong Ma, Yaping Zhang, Zhiyang Zhang, Yupu Liang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2024b. Born a BabyNet with hierarchical parental supervision for end-to-end text image machine translation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2468–2479, Torino, Italia. ELRA and ICCL.
- Minh Van Nguyen, Viet Dac Lai, Amir Poursan Ben Veyseh, and Thien Huu Nguyen. 2021. *Trankit: A lightweight transformer-based toolkit for multilingual natural language processing*. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90, Online. Association for Computational Linguistics.
- Zhipeng Qian, Pei Zhang, Baosong Yang, Kai Fan, Yiwei Ma, Derek F Wong, Xiaoshuai Sun, and Rongrong Ji. 2024. Anytrans: Translate anytext in the image with large scale models. *arXiv preprint arXiv:2406.11432*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. *COMET: A neural framework for MT evaluation*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Nilesh P Sable, Priya Shelke, Ninad Deogaonkar, Nachiket Joshi, Rudra Khabadi, and Tushar Joshi. 2023. Doc-handler: Document scanner, manipulator, and translator based on image and natural language processing. In *2023 International Conference on Emerging Smart Computing and Informatics (ESCI)*, pages 1–6. IEEE.
- Steinthor Steingrímsson, Hrafn Loftsson, and Andy Way. 2023. *SentAlign: Accurate and scalable sentence alignment*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 256–263, Singapore. Association for Computational Linguistics.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Yanzhi Tian, Xiang Li, Zeming Liu, Yuhang Guo, and Bin Wang. 2023. *In-image neural machine translation with segmented pixel sequence-to-sequence model*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15046–15057, Singapore. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, Jinrong Yang, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2023. Vary: Scaling up the vision vocabulary for large vision-language models. *arXiv preprint arXiv:2312.06109*.
- Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, En Yu, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2024. Small language model meets with reinforced vision vocabulary. *arXiv preprint arXiv:2401.12503*.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. The dawn of Imms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9(1):1.
- Cong Yao. 2023. Docxchain: A powerful open-source toolchain for document parsing and beyond. *arXiv preprint arXiv:2310.12430*.
- Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Yuhao Dan, Chenlin Zhao, Guohai Xu, Chenliang Li, Junfeng Tian, et al. 2023. mplug-docowl: Modularized multimodal large language model for document understanding. *arXiv preprint arXiv:2307.02499*.
- Ya-Qi Yu, Minghui Liao, Jihao Wu, Yongxin Liao, Xiaoyu Zheng, and Wei Zeng. 2024. Texthawk: Exploring efficient fine-grained perception of multimodal large language models. *arXiv preprint arXiv:2404.09204*.

Yanzhe Zhang, Ruiyi Zhang, Jiuxiang Gu, Yufan Zhou, Nedim Lipka, Diyi Yang, and Tong Sun. 2023a. LlavAr: Enhanced visual instruction tuning for text-rich image understanding. *arXiv preprint arXiv:2306.17107*.

Zhiyang Zhang, Yaping Zhang, Yupu Liang, Cong Ma, Lu Xiang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2025a. Understand layout and translate text: Unified feature-conductive end-to-end document image translation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(5):3358–3376.

Zhiyang Zhang, Yaping Zhang, Yupu Liang, Lu Xiang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023b. LayoutDIT: Layout-aware end-to-end document image translation with multi-step conductive decoder. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10043–10053, Singapore. Association for Computational Linguistics.

Zhiyang Zhang, Yaping Zhang, Yupu Liang, Lu Xiang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2025b. From chaotic OCR words to coherent document: A fine-to-coarse zoom-out network for complex-layout document image translation. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10877–10890, Abu Dhabi, UAE. Association for Computational Linguistics.

Zhiyang Zhang, Yaping Zhang, Lu Xiang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023c. A novel dataset and benchmark analysis on document image translation. In *China Conference on Machine Translation*, pages 103–115. Springer.

Shaolin Zhu, Shangjie Li, Yikun Lei, and Deyi Xiong. 2023. PEIT: Bridging the modality gap with pre-trained models for end-to-end image translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13433–13447, Toronto, Canada. Association for Computational Linguistics.

Wang Zhu, Alekh Agarwal, Mandar Joshi, Robin Jia, Jesse Thomason, and Kristina Toutanova. 2024. Efficient end-to-end visual document understanding with rationale distillation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8401–8424, Mexico City, Mexico. Association for Computational Linguistics.

Michał Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. The United Nations parallel corpus v1.0. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3530–3534, Portorož, Slovenia. European Language Resources Association (ELRA).

Appendix

A Setting Details

A.1 Dataset Settings

In the DITrans dataset, the sample sizes for the advertisement, news, and political report subdomains are 485, 610, and 1397, respectively. Due to the small number of images in the advertisement and news domains and their similar layout structures as scanned document images, we merge these two domains. We then randomly select 100 images as the test set and another 100 images as the valid set. For the political report domain, we also randomly select 100 images as the test set and another 100 images as the valid set.

In the experiment of varying context length, to shield the impact of layout difference, we select images with a single column and without formulas, tables, or figures. Context length refers to the number of English words in the image.

In the experiment of varying layout complexity, we transform samples from the valid set into trees, selecting the 100 trees with the fewest nodes as a simple layout set and the 100 trees with the most nodes as a complex layout set.

A.2 Main Experiment Settings

We segment the Chinese texts with `jieba` and apply WordPiece to segment both English and Chinese texts and the vocabulary size of both English and Chinese is 52K. We use the pre-trained OCR model Nougat’s encoder (Blecher et al., 2024) to initialize the Swin Transformer of alignment encoder and image encoder. The layer numbers and window size are 2, 2, 14, 2 and 7. The hidden size of each layer is 1024 and the patch size is 4. The input image size is 896 × 672. We follow the vanilla Transformer-base (Vaswani et al., 2017) setting and pre-train an English-Chinese translation model on the UN Corpus. We set the decoder’s max length and max position embeddings to 1536 to cover most input texts. For the MLLM, we select four MLLMs with different numbers of parameters and training data: Vary-toy (Wei et al., 2024), Vary-base (Wei et al., 2023), Llava-next (Liu et al., 2024a) and Textmonkey (Liu et al., 2024b), as shown in Table 7. The hyperparameter α is set to 1.0 and the sequence lengths for all MLLMs, except Llava-next, are set to 2048 to cover the long context of document images. The sequence length for Llava-next is 4096 due to the different image encoders and prompts

	MLLM	# Params (M)	Training Data
1	Vary-toy	2237.4	Document images
2	Vary-base	8123.7	Document images
3	Llava-next	8354.8	Natural images
4	Textmonkey	9715.8	Document images

Table 7: The number of parameters and training data of different MLLMs.

		BLEU	BLEU-PT	STEDS
1	Cross-entropy	31.67	34.60	81.60
2	MSE	33.00	35.67	82.68
3	Cosine-similarity	40.05	42.58	83.93

Table 8: Results on DoTA English-Chinese valid set with different alignment loss functions.

used by the MLLMs.

During translation model pre-training, the maximum training step is 100K and the maximum token per batch is 4096. A linear decay learning rate schedule with a learning rate of $7e-4$ and a warmup ratio of 0.05 is used. During the training stage of M4Doc, the maximum number of training steps is 15K and the batch size is 64. We use a linear decay learning rate schedule with a learning rate of $5e-5$ and the number of warmup steps is 1000. We use Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-8$ for both training stages. We used two NVIDIA A100 GPUs and spent 28 hours to complete the training of M4Doc (Vary-toy), and 68 hours to complete the training of M4Doc (Textmonkey). For inference, we use beam search with 4 beams.

A.3 Prompts for Each MLLM

The <System Prompt> and <User Prompt> used in the main experiment are listed as follows.

Prompts for Vary-toy/base

<System Prompt>
None

<User Prompt>
Convert the image to markdown/latex format.

Prompts for Llava-next

<System Prompt>
You are a helpful language and vision assistant. You are able to understand the visual content that the user provides, and assist the user with a variety of tasks using natural language.

<User Prompt>

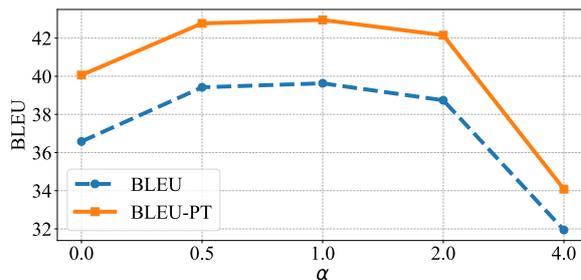


Figure 6: BLEU and BLEU-PT scores of M4Doc (Vary-toy) trained with different α values on the valid set. Detailed data can be seen in Appendix C.

OCR this image.

Prompts for Textmonkey

<System Prompt>
You are a helpful assistant.

<User Prompt>
Read all the text in the image.

B Detailed Analysis

B.1 Effect of Different Alignment Loss Functions

To explore the impact of different alignment loss functions, we use cross-entropy loss, mean square error (MSE) loss, and cosine-similarity loss as the alignment loss function and conduct experiments with the same setting as the main experiment M4Doc (Vary-toy). The results are shown in Table 8.

As shown in the table, using cosine similarity as the alignment loss function yields the best results. We think this may be because the loss values calculated by cosine-similarity range between $[-1, 1]$, allowing the model to strike a balance between learning the alignment task and the translation task. Therefore, we choose cosine-similarity loss for the main experiment.

B.2 Hyperparameter Sensitivity Analysis

To explore the impact of α in the loss function, we vary α and get results in Figure 6. As shown in the figure, the model’s performance initially increases and then decreases with the increase in α , achieving the best performance when $\alpha = 1.0$. This could potentially be attributed to the fact that a small α diminishes the influence of MLLM, while a large α introduces too much noise. So, we set $\alpha = 1.0$ in the main experiment.

		DoTA			DITrans			# Params	Time
		BLEU	BLEU-PT	STEDS	BLEU	BLEU-PT	STEDS	(M)	(s/page)
1	Vary-toy (Original)	10.64	4.92	66.23	2.07	2.10	45.12	2237.4	62.58
2	Vary-toy (Fine-tuned)	22.67	22.75	73.99	5.87	6.14	49.59	2253.9	57.60
3	M4Doc (Vary-toy)	39.95	42.33	83.97	14.79	18.67	53.73	212.4	9.61
4	Vary-base (Original)	13.45	5.79	76.26	2.84	2.79	56.21	8123.7	68.84
5	Vary-base (Fine-tuned)	38.60	38.53	82.95	11.61	11.72	54.59	8137.9	69.67
6	M4Doc (Vary-base)	41.22	42.09	86.06	14.52	16.55	55.89	215.6	9.43

Table 9: Results of directly fine-tuning MLLMs on the DoTA dataset. **# Params** is the number of parameters of the model during inference. **Time** is the average inference time on a single NVIDIA V100 GPU.

		B	C	BP	S
1	GPT-4o	29.17	60.32	31.95	59.45
2	Gemini	30.31	59.67	31.69	63.32
3	DIMTDA	38.73	61.33	42.37	84.98
4	M4Doc (Vary-toy)	39.45	62.42	42.59	83.50
5	M4Doc (Vary-base)	41.11	63.57	42.00	86.62
6	M4Doc (Textmonkey)	42.27	65.40	44.17	86.93

Table 10: Results on comparison with commercial MLLMs. **B**, **C**, **BP**, and **S** represent BLEU, COMET, BLEU-PT, and STEDS, respectively.

B.3 Comparison with Fine-tuning MLLM

We conduct comparative experiments to evaluate the DIMT capabilities of MLLMs by directly applying MLLMs to the DIMT task and fine-tuning them specifically for this task. We fine-tune Vary-toy and Vary-base using LoRA (Hu et al., 2022) with a *lora_rank* of 32, while keeping other settings consistent with the main experiment. The results are presented in Table 9.

As shown in the table, directly using MLLMs for the DIMT task yields very poor performance (line 1 and 4), with almost no translation capability on the political report domain. After fine-tuning on the DoTA dataset, the DIMT capability of MLLMs improves significantly but still falls short of the performance achieved by our proposed M4Doc method. By comparing line 5 and 6, our method outperforms the best-performing Vary-base (Fine-tuned) model by 2.62 BLEU scores and achieves a greater improvement of 2.91 BLEU scores in zero-shot cross-domain scenarios. This highlights the potential of our method for efficiently leveraging MLLMs across various downstream tasks.

B.4 Comparison with Commercial MLLMs

With the rapid development of MLLMs, some commercial MLLMs (Hurst et al., 2024; Team et al., 2024) demonstrate the capability of understanding text-rich document images. To assess their ability

to accomplish the DIMT task, we randomly choose 200 samples from the test set of the DoTA dataset, then prompt GPT-4o and Gemini with three different prompts to complete the document image machine translation task. The prompts we used are as follows.

Prompts for GPT-4o and Gemini to complete DIMT task

<Prompt 1>
Output the Chinese translations of this image in markdown format.

<Prompt 2>
Please extract and provide the Chinese translations of the text contained within this image, ensuring that the translations are accurately represented, and format them using markdown for clear presentation.

<Prompt 3>
Please translate the all texts in this image into English and adhere to the following translation standards:
Accuracy: Ensure that the translation fully captures the meaning of all the texts in the image without adding or omitting any information.
Fluency: The translation should read naturally and smoothly, reflecting the conventions of the target language and the translation should follow the reading order of the image.
Format: The translation should be presented in markdown format.

We average the metric values of the translation results obtained from different prompts to determine the final results. As the output format of MLLMs may be unstable, we filter the English parts of the output text and only keep the Chinese parts.

Table 10 reveals that both GPT-4o and Gemini can accomplish the DIMT task directly, but exhibit inferior performance compared to M4Doc (line 2 vs. line 6). This may be because commercial MLLMs are not trained on the DoTA dataset, their output formats differ from the reference. This leads to commercial MLLMs performing signifi-

		En-Fr		En-De	
		BLEU	STEDS	BLEU	STEDS
1	Text-only MT	59.68	95.93	49.25	96.04
Cascade Baselines					
2	LARDIT	42.79	75.59	32.65	75.59
3	Nougat-trans	55.82	90.77	43.73	89.92
End-to-end DIMT					
4	DIMTDA	45.82	84.84	37.83	85.92
5	M4Doc (Vary-toy)	48.88	85.04	41.47	86.72
6	M4Doc (Vary-base)	49.18	86.83	42.61	86.64
7	M4Doc (Textmonkey)	54.64	89.85	46.70	89.58

Table 11: Results on DoTA English-French and English-German test set.

cantly worse than M4Doc on metrics like BLEU and STEDS. However, semantic-based evaluation metrics, such as COMET, can more accurately reflect the model’s translation performance, which shows that the DIMT ability of existing commercial MLLMs is comparable to that of M4Doc.

B.5 Evaluation on Other Languages

To verify our method’s effectiveness in other languages, we conduct English-French and English-German DIMT experiments. The text machine translation models are pre-trained on the UN Corpus En-Fr and WMT14 En-De dataset. We use the En-Fr and En-De subsets of the DoTA dataset to train our models. The rest of the settings remain the same as the main experiment. Table 11 demonstrates the effectiveness of M4Doc on other languages’ DIMT tasks.

C Detailed Data

Table 12 presents the detailed data corresponding to the BLEU scores of M4Doc models tested on validation sets with different context lengths, as shown in Figure 3. Table 13 provides the detailed data corresponding to the BLEU and BLEU-PT scores of M4Doc (Vary-toy) trained with different α values on the validation set, as illustrated in Figure 6.

D Output Samples

We provide the output samples of M4Doc in cross-domain and long context scenarios in Figure 7. Figure 7 (a) is an image from the ads & news subset of the DITrans dataset. The scanned document image contains a lot of noise, and the font size varies significantly, which makes the image difficult to handle. After fine-tuning our model on the subset, it can translate the text in the image, even if some of the text appears blurry.

	(0,250]	(250,500]
DIMTDA	51.13	45.95
M4Doc (Vary-toy)	51.82	45.75
M4Doc (Vary-base)	52.85	49.46
M4Doc (Textmonkey)	56.04	50.15
	(500,750]	(750,)
DIMTDA	39.98	34.85
M4Doc (Vary-toy)	45.54	41.10
M4Doc (Vary-base)	45.43	39.02
M4Doc (Textmonkey)	48.65	45.63

Table 12: Detailed data of Figure 3.

α	BLEU	BLEU-PT
0.0	36.58	40.06
0.5	39.42	42.77
1.0	39.63	42.95
2.0	38.74	42.15
4.0	31.94	34.08

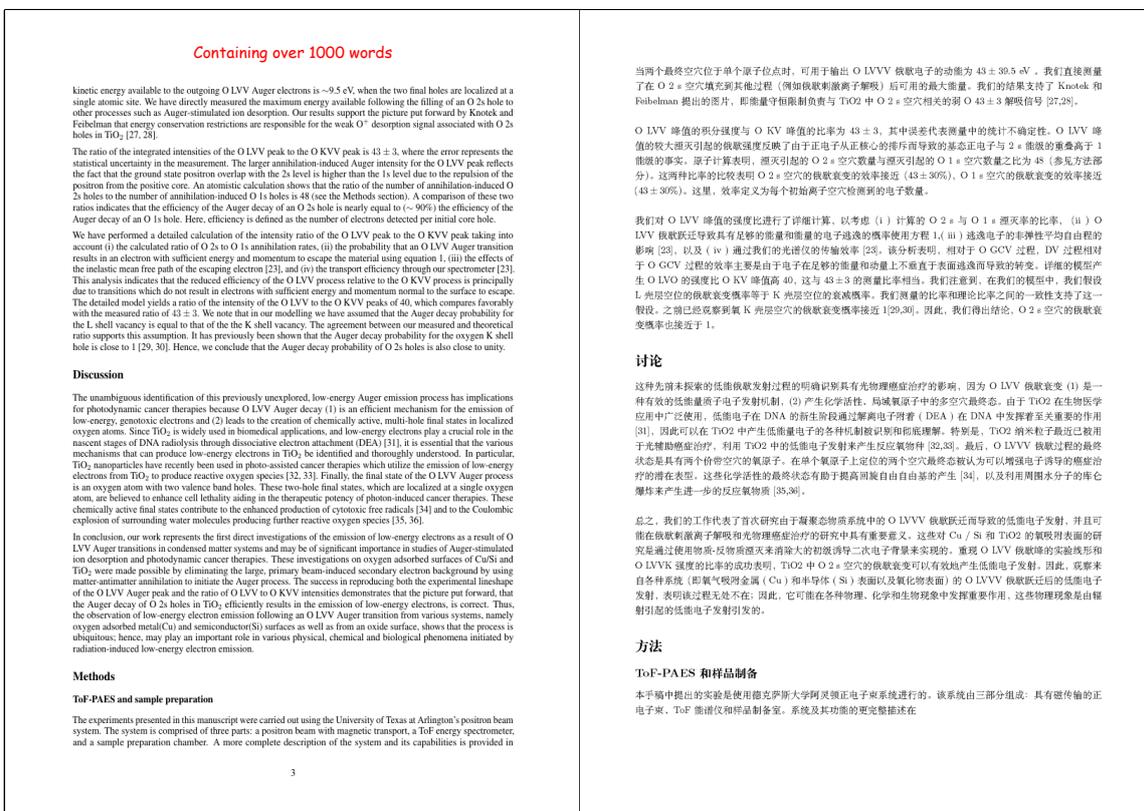
Table 13: Detailed data of Figure 6.

Figure 7 (b) is an image from the DoTA dataset, which contains more than 1000 English words. For images containing such long contexts, our model still achieves end-to-end DIMT without omissions.

We also list other output samples in Figure 8, Figure 9, and Figure 10.



(a) Cross-domain: Ads & News



(b) Long Context

Figure 7: The output samples of M4Doc. For each image pair, the left is the input document image, and the right is the output translations in markdown format after rendering. It is better to zoom in for a clearer view.

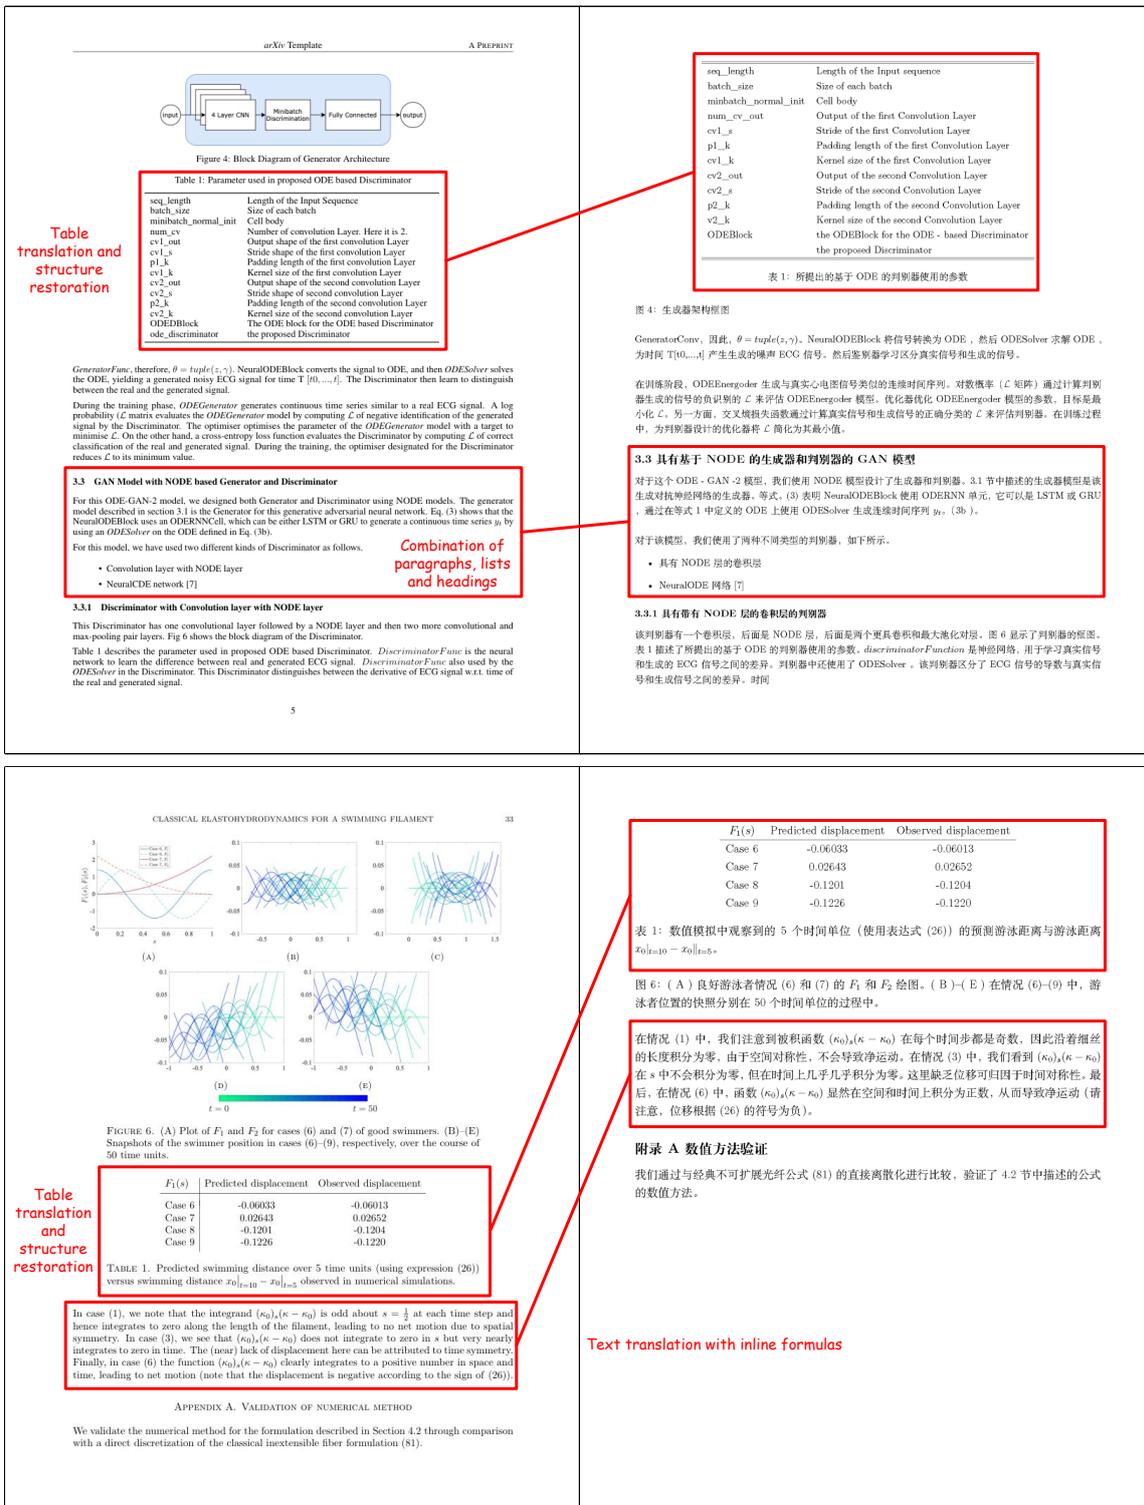


Figure 8: The output samples of M4Doc in the DoTA dataset. It is better to zoom in for a clearer view.

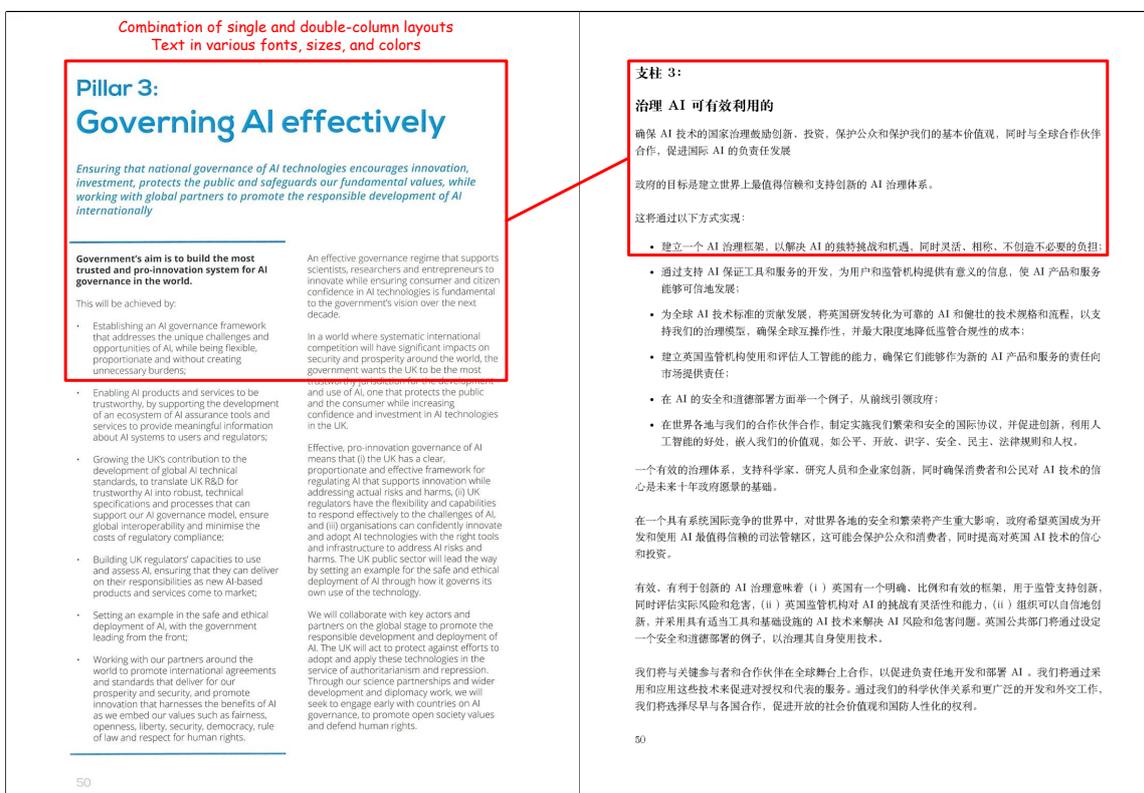


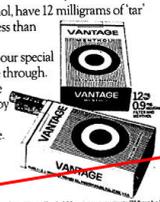
Figure 9: The output samples of M4Doc in the political report subset of the DITrans dataset. It is better to zoom in for a clearer view.

2

Rev. 11/4/01/1/2
Cip - 2/25/73

How many times have you decided to give up smoking?

Nobody these days is telling you not to give up smoking. But if you've given it up more times than you'd like to remember, the chances are you enjoy it too much to want to give it up at all. If you're like a lot of smokers these days, it probably isn't smoking that you want to give up. It's some of that 'tar' and nicotine you've been hearing about. So you tried cigarettes which were low in 'tar' and you found yourself checking every once in a while to see if they were still lit. Which drove you right back to your regular brand. But now, there is Vantage. Vantage cigarettes, either filter or menthol, have 12 milligrams of 'tar' and 0.9 milligrams of nicotine, considerably less than most cigarettes. And what really makes Vantage special is our special filter which allows the tobacco flavor to come through. Vantage isn't the lowest 'tar' and nicotine cigarette, but it sure is the lowest one you'll enjoy smoking. And that's what makes all the difference.



Warning: The Surgeon General Has Determined That Cigarette Smoking is Dangerous to Your Health.

Blurred text with significant noise

2

Reid 17-6/73 Cip -2/25/73

你决定戒烟多少次了？

现在没有人告诉你不要戒烟。

但如果你放弃的次数超过了你想记住的次数，那么很可能你太不想放弃。

如果你现在和很多吸烟者一样，你可能不会戒烟。这是你一直听说的一些“焦油”和尼古丁。

所以你尝试了低焦油含量的香烟，你发现你自己每隔一段时间就会检查一下，看看它们是否还亮着。这让你回到了你的常规品牌。

但现在，有了 Vantage。

Vantage 香烟，无论是过滤香烟还是薄荷醇香烟，都有 12 毫克“焦油”和 0.9 毫克的尼古丁毫克，大多数香烟。

真正让 Vantage 与众不同的特殊过滤器，它可以使烟草的味道通过。

Vantage 不是焦油和尼古丁含量最低的香烟，但它肯定是你喜欢的最低的一种。

这就是一切不同的原因。

警告：卫生部长已经确定吸烟对你的健康有害。

5

Long text images with substantial noise

Text in various fonts and sizes

x 5 Authors
A Smoking-Quitter
Method

A Psychosocial Model of Smoking Cessation and Maintenance of Cessation

JUDITH K. OCKENE,* RONALD NUTTALL,† ROBERT C. BENFARI,* IRVING PUWITZ,‡ AND IRA S. OCKENT†

* Department of Behavioral Sciences, Harvard School of Public Health, Boston, Massachusetts 02115; †Graduate School of Arts and Sciences, Boston College, Newton, Massachusetts 02459; ‡ Division of Cardiovascular Medicine, University of Massachusetts Medical School, Worcester, Massachusetts 01605

The long-term cessation rates obtained in smoking control programs have been disappointing. In order to provide a better understanding of the problems of smoking cessation and maintenance of cessation, and to facilitate the development of more effective smoking control programs, a model of smoking cessation is offered in which stress, or "life events," is a major determinant of the success of this process. The model assumes that chronic smoking is a maladaptive behavior, helping the smoker to deal with the discomfort generated by stress by providing a means he can use to maintain a "vital balance." A smoker's ability to handle stress without cigarettes is conditioned by the presence of personal and social resources or "psychosocial assets." Personal security, an individual's belief that he can control what happens to him, an ability to respond to stress with a low level of negative affect, and the availability of social supports provide a cushion or buffer against the effects of stress. It is therefore the combined effect of stress and the psychosocial variables which determine whether or not a smoker can successfully eliminate smoking from his repertoire of coping responses.

Despite antismoking campaigns, attempts at behavioral intervention, and acceptance by smokers of the dangers of smoking, many smokers are unable to stop or to maintain cessation after stopping. Much research has been done attempting to find a "cure" for smoking; none has been found. The long-term cessation rates obtained in large-scale smoking control programs have been disappointing. When reliable treatment effects have occurred, they have decreased rapidly during the first month, with relapse continuing to occur during the first 2 years of follow-up (11, 38, 59, 79).

Smoking, a complex psychosocial behavior pattern, goes through stages of initiation, transition, maintenance, and termination (21, 43). As indicated by the relapse curves, maintenance of cessation is yet another stage on the smoking behavior continuum. In each stage the individual has different motives and needs. Social forces or forces extrinsic to the individual are strongly associated with the initiation of smoking (1, 35, 55, 58, 80, 85). Initiation of smoking may also be related to psychological variables, such as self-esteem, status needs, and other personal needs (52, 82, 85, 87), although social forces are more critical at this stage (5). The sociological variables which are so important during the formation of the habit seem to play a minor role once smoking has become part of the life style of the smoker (35, 58). It therefore is a habit which initially has strong extrinsic motivations, but in time develops its own intrinsic motivation. As the habit continues, it becomes increasingly tied to psychological needs, and becomes an in-

633

0091-7435/81/05052-0000
Copyright © 1981 by Academic Press, Inc.
All rights of reproduction in any form reserved.

Small text with noise

TIMN 276028

5

预防医学 10/323-638 (1981) 9.

戒烟与戒烟的心理社会模式

朱迪斯·K·科克斯森* 雷诺塔杜尔* 罗伯特·C·本森*

Irwin Puwitz, 和 Ira S. Ockent

国家健康与预防中心，波士顿，马萨诸塞州；1996年10月11日；美国广播公司编辑，波士顿，波士顿，马萨诸塞州；1995年10月9日；美国广播公司编辑，密苏里州医学院心血管医学研究所，马萨诸塞州芝加哥斯德哥尔摩

戒烟控制计划获得的长期戒烟率已经下降。为了提供对戒烟控制问题的更好理解和维持的发展，并将其制定为更有效的戒烟控制计划，需要一个戒烟运动，以说明该过程的压力，或“修复活动”，这是成功的一个主要决定因素。该计划使戒烟成为一种令人绝望的行为，帮助吸烟者通过提供一种可以发挥“重要价值”的优势来消除与戒烟不同压力的风险。一名吸烟者注意到压力或“精神病患者”的行为。个人相信这种行为会与吸烟的行为相矛盾。个人相信这种行为会对他造成伤害，一种对低水平戒烟影响的能力，以及社会支持如何对压力产生影响的行为构成成瘾物质。因此，它对血压和心理变量的综合影响，这些变化决定了吸烟是否能够成功消除戒烟。

尽管戒烟运动，但在行为干预方面取得的尝试，以及吸烟者危险的节奏，许多吸烟者无法停止或维持戒烟。已经进行了大量研究来发现“诅咒”；没有发现。长期以来在大规模的戒烟控制计划中获得的戒烟率一直存在。当发生可靠的治疗效果时，他们已经在第一次2年内慢慢减少。累积坚持发生在第一次2年内 (11, 38, 59, 79)。

吸烟是一种复杂的心理行为模式，在最初的分期、分期、分期、维护和终止阶段 (21, 43)。正如拒绝治疗证据所表明的那样，戒烟维持是现阶段的一个阶段。在每个阶段，个人都有不同的动机和需求。社会力量或对个人的过度依赖与吸烟的发起强烈相关 (3, 35, 58, 80, 85, 80)。尽管社会力量需要，否则吸烟也可能与心理变量有关 (如自我外科医生、需求和其他个人) (5, 28, 82, 87, 90) 有关。但在这个阶段，社会力量也更重要。在吸烟的形成过程中，这种社会变化非常重要，很快就会成为吸烟者生命风格的一部分 (35, 8)。因此，这是一个习惯最初具有强烈的激励，但在时间上它确实具有固有的动机。随着习惯的继续，承载越来越符合心理需求，变得越来越棘手。

633

0091-7435/81/05052-0000

版权所有 P:1941, 咨询公司。他表示，在任何时候戒烟的自我方面都很重要。

Figure 10: The output samples of M4Doc in the ads & news subset of the DITrans dataset. It is better to zoom in for a clearer view.