DAVE: A VLM VISION ENCODER FOR DOCUMENT UNDERSTANDING AND WEB AGENTS

Anonymous authors

000

001

002003004

010 011

012

013

014

016

017

018

019

021

025

026

027

028

031

033

034

037

038

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

While Vision-language models (VLMs) have demonstrated remarkable performance across multi-modal tasks, their choice of vision encoders presents a fundamental weakness: their low-level features lack the robust structural and spatial information essential for document understanding and web agents. To bridge this gap, we introduce DAVE, a vision encoder purpose-built for VLMs and tailored for these tasks. Our training pipeline is designed to leverage abundant unlabeled data to bypass the need for costly large-scale annotations for document and web images. We begin with a self-supervised pretraining stage on unlabeled images, followed by a supervised autoregressive pretraining stage, where the model learns tasks like parsing and localization from limited, high-quality data. Within the supervised stage, we adopt two strategies to improve our encoder's alignment with both general visual knowledge and diverse document and web agentic tasks: (i) We introduce a novel model-merging scheme, combining encoders trained with different text decoders to ensure broad compatibility with different web agentic architectures. (ii) We use ensemble training to fuse features from pretrained generalist encoders (e.g., SigLIP2) with our own document and web-specific representations. Extensive experiments on classic document tasks, VQAs, web localization, and agent-based benchmarks validate the effectiveness of our approach, establishing DAVE as a strong vision encoder for document and web applications.

1 Introduction

Vision-Language Models (VLMs) (Liu et al., 2024a; Bai et al., 2025; Achiam et al., 2023) have shown remarkable capabilities in multi-modal reasoning and understanding, enabling a wide range of applications from image captioning to interactive web agents (Wu et al., 2024; Qin et al., 2025; Liu et al., 2023; Li et al., 2022b). A central component of VLMs is the vision encoder (Yin et al., 2024), which converts images into visual tokens for the language backbone to process (Li et al., 2023b; Liu et al., 2023). As such, the design of these encoders has become a key focus in VLM research (Tong et al., 2024; Shi et al., 2024).

However, the vision encoders predominantly used by VLMs suffer from a fundamental weakness: their low-level features lack the robust structural and spatial information essential for document understanding and web agents (Tschannen et al., 2025; Radford et al., 2021). Conversely, DINO-style models (Oquab et al., 2023) offer low-level features yet are tuned to natural images (Oquab et al., 2023; Siméoni et al., 2025) and transfer poorly to documents, UIs, and diagrams (Tong et al., 2024). To bridge this gap, we introduce DAVE (*Document and web Agents Vision Encoder*), a vision encoder purpose-built for VLMs and tailored for these tasks. See Figure 1 for an overview.

A key challenge in training on document and web images is the scarcity of high-quality annotated data, since current annotations rely on OCR models that bottleneck both scalability and quality (Kim et al., 2021; Lee et al., 2023). We address this with a two-stage training process: self-supervised pretraining on large-scale unlabeled data, followed by autoregressive pretraining on limited annotated structural and grounding data. Yet, the supervised stage introduces a further challenge: specialized vision encoders trained with a single text decoder tend to overfit to that decoder, leading to misalignment when paired with other decoders. We mitigate this issue via multi-decoder pretraining and weight-merging ("model soup") (Wortsman et al., 2022), which produces a largely decoder-agnostic encoder. To preserve high-level semantics while leveraging our structural and spatial features, we




Figure 1: **DAVE Overview**. Stage 1 trains the vision encoder with a decoder using MAE, learning strong structural and spatial priors from unlabeled data. Stage 2 performs autoregressive pretraining on diverse tasks with different text decoders and fuses the high-level semantic features from SigLIP 2. After that, different encoders are combined into a single one by learning a merge coefficient using unsupervised representation alignment, while keeping the encoders frozen.

fuse a generalist encoder with our encoder during pretraining via ensembling. The resulting specialized encoder integrates seamlessly with diverse decoders and generalist VLM stacks. Finally, our purpose-built vision encoder serves as the core component of a VLM architecture specifically designed for document and web agent tasks. We evaluate DAVE across diverse settings, including traditional and vision—language tasks. On traditional benchmarks, DAVE surpasses all SOTA models in document recognition and segmentation, while achieving competitive results with SigLIP 2 on screenshot classification. For vision—language evaluation, we test on document understanding, web localization, and web agent. Notably, **DAVE improves performance by an average of 10.5% compared to SigLIP 2**, highlighting its ability to provide specialized representations. Finally, we evaluate web agent settings on the Mind2Web benchmark (Deng et al., 2023), where DAVE improves web agent performance by 5% over the strongest baseline vision encoder.

We summarize our main contributions as follows: (i) We introduce DAVE, a vision encoder purpose-built for VLMs and tailored for document and web agents tasks; (ii) We propose a two-stage pretraining framework that combines a self-supervised pretraining stage on unlabeled images, followed by a supervised autoregressive pretraining stage; (iii) We introduce model weight merging and ensemble training strategies to enhance the alignment of structural and spatial representation for documents and web images with diverse VLM and agentic frameworks; (iv) We perform extensive experiments to evaluate the impact of visual representations on document and web understanding, as well as downstream web-agent tasks.

2 Related Work

Pretrained Vision Encoders. Image–text contrastive learning (Radford et al., 2021; Zhai et al., 2023; Jia et al., 2021; Yu et al., 2022) and self-supervised representation learning (He et al., 2022; Zhou et al., 2021; Oquab et al., 2023; Chen et al., 2020; Caron et al., 2021) are the two main paradigms for large-scale visual pretraining. Recent work has sought to combine the strengths of the two approaches through joint training or post-hoc alignment (Tschannen et al., 2025; Fini et al., 2025; Maninis et al., 2024; Naeem et al., 2024), yielding strong encoders for both vision and multimodal tasks. The works most closely related to ours are Eagle (Shi et al., 2024) and Perception Encoder (Bolya et al., 2025), which align pretrained vision encoders to a pretrained LLM with large-scale general datasets. In contrast, we pretrain the vision encoder from scratch with self-supervised

learning, use multiple pretrained LLMs as a backbone, and use domain-specific data to create a specialized foundation model.

Document Understanding with LLMs. Document processing and understanding have transitioned from OCR-based (Xu et al., 2020b;a) methods to vision-language-based. One line of work trains encoder-decoder models with various objectives, like extracting structures (Feng et al., 2025; Kim et al., 2022; Lv et al., 2023; Lee et al., 2023), and image reconstruction conditioned on text (Huang et al., 2022; Tang et al., 2023). Recent advancements in VLMs led to another line of work that uses document data to finetune a VLM (Nassar et al., 2025; Hu et al., 2024; Liu et al., 2024g) or incorporate it as part of the VLM training (Team et al., 2025; Bai et al., 2025; Chen et al., 2024). Our work seeks to bridge these two lines of work by introducing a specialized vision encoder within a general VLM framework.

Vision-Language Models. Inspired by the success of recent large language models (LLMs) (Brown et al., 2020; Dubey et al., 2024), Vision-language models (Radford et al., 2021; Liu et al., 2024a; Hua et al., 2024a; Ye et al., 2023; Tang et al., 2024; Bai et al., 2025; Zhu et al., 2025; Hua et al., 2025; Tong et al., 2024) aim to achieve multimodal intelligence by jointly understanding and generating visual and language information. Flamingo (Alayrac et al., 2022) and BLIP-2 (Li et al., 2023b) are two of the early works that explore integrating LLMs as part of VLM. Beginning with LLaVA (Liu et al., 2024a), researchers have used instruction-following chat data in VQA format for instruction tuning, achieving significantly improved results (Liu et al., 2024e; Li et al., 2023a; Hua et al., 2024b; Tang et al., 2025; Lu et al., 2023).

Multimodal Web Agents. Vision-based autonomous web agents have recently attracted attention for their simplicity and stronger generalization compared to LLM-based agents (Deng et al., 2023; Kim et al., 2023). Early efforts such as WebGUM (Furuta et al., 2023) and CogAgent (Hong et al., 2024) pretrained vision—language models (VLMs) on web and GUI data to enhance agentic capabilities. Follow-up work (Wu et al., 2024; Xu et al., 2024; Liu et al., 2024b; Qin et al., 2025) expanded training with large-scale grounding and interaction datasets, further improving VLM performance on web-based tasks. While these works are promising steps toward a general-purpose visual web agent, the role of the vision encoder in these systems remains underexplored.

3 THE DAVE MODEL

3.1 Preliminarilies

Masked Autoencoder. Masked Autoencoder (MAE) (He et al., 2022) adopts an asymmetric encoder–decoder architecture. The vision encoder processes only a subset of visible patches, while a lightweight transformer decoder reconstructs the full image from the encoded features and mask tokens. After pretraining, only the encoder is retained for downstream tasks.

By default, MAE applies per-patch normalization before computing the reconstruction loss to learn better representations. The reconstruction loss is then defined as

$$\mathcal{L}_{\text{MAE}} = \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \left\| f_{\theta}(\tilde{x})_i - \frac{x_i - \mu(x_i)}{\sqrt{\sigma^2(x_i) + \epsilon}} \right\|_2^2, \quad \epsilon = 10^{-6}$$
 (1)

where \tilde{x} denotes the input image with masked patches, $f_{\theta}(\tilde{x})_i$ is the reconstructed output of patch i predicted by the decoder with parameters θ , x_i is the ground-truth pixel values of patch i, $\mu(x_i)$ and $\sigma^2(x_i)$ are the mean and variance of pixels within patch i, \mathcal{M} is the set of masked patches.

Vision-Language Model. A vision-language model (VLM) consists of a vision encoder ϕ , an MLP projector, and a text decoder Θ . The vision encoder takes an input image $x \in \mathbb{R}^{H \times W \times 3}$ and produces a sequence of patch-level features $\{v_i\}_{i=1}^N$. These are projected into the text embedding space via the MLP projector and used along with the text input.

In what follows, we describe the process to get our vision encoder.

3.2 STAGE 1: SSL ON DOCUMENT AND WEB IMAGES

In Stage 1, we conduct self-supervised training with MAE to learn rich structural and spatial features from document and web images. While MAE demonstrates the strongest performance on OCR

tasks compared to other methods, the training is prone to instability when trained at scale (Fan et al., 2025). We find that this instability is particularly acute for document and web images. Our analysis (Section 5.4) attributes this to their characteristically low inter-patch variance, which destabilizes the target normalization in the standard MAE objective (Equation 1). To address this, we modify the objective to reconstruct raw pixel values directly:

$$\mathcal{L}_{\text{MAE-pixel}} = \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \|f_{\theta}(\tilde{x})_i - x_i\|_2^2, \tag{2}$$

This change stabilizes training and enables scaling the training sample to 20 million images without additional hyperparameter tuning. With this robust self-supervised procedure, we then proceed to the second stage of our training process.

3.3 Stage 2: Supervised Multi-Task Pretraining

To further enhance the encoder's structural and spatial understanding for document and web images, we perform a supervised, autoregressive pretraining stage. This stage utilizes a limited set of high-quality labeled data for tasks like OCR, layout extraction, and web localizations. We use the VLM architecture as described in Section 3.1, and DAVE as the vision in the VLM architecture.

Weight-space Merging. A key limitation of the architecture discussed above is that the pretrained vision encoder becomes tightly coupled to its specific text decoder. This coupling significantly degrades performance when the encoder is integrated with a different decoder. To address this, we use model merging to create a vision encoder that is agnostic to the choice of text decoder.

Formally, given a set of n pretrained text decoders $\{\Theta_1, \dots, \Theta_n\}$, we train n corresponding instances of our vision encoder, denoted $\{\phi_1, \dots, \phi_n\}$. Each instance is aligned with a different text decoder but otherwise shares the same architecture, training data, and hyperparameters.

To merge these different encoders, we propose a distillation-based merging scheme that learns a small set of coefficients to combine pretrained weights while keeping the original parameters frozen. Considering each encoder ϕ_i as a set of m weights $\{\theta_i^{(j)}\}_{j=1}^m$, we form each merged weight $\theta_{\text{merge}}^{(j)}$ by learning a set of coefficients $\{\alpha_i^{(j)}\}_{i=1}^n$ to compute the weighted sum of the corresponding weight for each of the n encoders:

$$\theta_{\text{merge}}^{(j)} = \sum_{i=1}^{n} \alpha_i^{(j)} \theta_i^{(j)} \qquad \alpha_i^{(j)} \in [0, 1].$$

The merged encoder is thus composed of m merged weights: $\phi_{\text{merge}} = \{\theta_{\text{merge}}^{(j)}\}_{i=1}^m$.

The resulting encoder produces patch-level features $\mathbf{z} = \phi_{\text{merge}}(\mathcal{I})$. To ensure that ϕ_{merge} preserves the features from all n teacher encoders, we define a distillation loss, $\mathcal{L}_{\text{distill}}$. This objective minimizes the average Mean Squared Error between the merged features \mathbf{z} and the features from each teacher encoder $\mathbf{z}_i = \phi_i(\mathcal{I})$:

$$\mathcal{L}_{ ext{distill}} = rac{1}{n} \sum_{i=1}^{n} \left\| \hat{\mathbf{z}}_i - \mathbf{z}_i
ight\|_2^2.$$

During the distillation process, all encoder parameters remain frozen; only the newly introduced combination coefficients are optimized. The final vision encoder is given by the merged encoder:

$$\phi_{\text{DAVE}}^{\text{final}} = \phi_{\text{merge}}(\{\alpha_i^{\star}\}_{i=1}^n),$$

where $\{\alpha_i^*\}_{i=1}^n$ are the optimized coefficients obtained from this training.

The next section details our method for fusing specialized document and web features with general visual representations.

Ensemble Training. Pretraining the encoder exclusively on document and web data provides strong structural and spatial features, but it also limits its grasp of general visual representation. This is a crucial shortcoming, as the high-level semantic features learned from diverse, large-scale datasets are equally important for robust performance.

To obtain both types of features, we design an ensemble pretraining paradigm that combines a frozen pretrained generalist encoder ϕ_{gen} with our document and web specialist encoder ϕ_{spec} from the previous stage: $\phi_{\text{DAVE}}(x) = \text{Concat}(\phi_{\text{gen}}(x), \phi_{\text{spec}}(x))$.

This design provides two main benefits: (i) it encourages $\phi_{\rm spec}(x)$ to focus on learning low-level structural and spatial representations, as high-level semantics are captured by $\phi_{\rm gen}(x)$; and (ii) it enables early fusion of structural and spatial features with high-level semantic features.

4 EVALUATION

4.1 IMPLEMENTATION DETAILS

Self-supervised Stage. For the SSL stage, we follow the original MAE implementation (He et al., 2022), with ViT-L16-384 (Dosovitskiy et al., 2020) as our vision encoder. We train the MAE with a batch size of 4096 for 120K steps. More detail can be found in Appendix A.1

Supervised Training Stage. We adopt the VLM architecture discussed in Section 3.1, with an image tilting size of 384. During the ensemble training, we employ frozen SigLIP-2 as the generalist vision encoder, while the encoder from the Self-supervised stage serves as the domain-specialized component. This forms the full DAVE encoder. We experiment with multiple LLMs of varying scales and architectures, including QWen2.5-0.5B-Instruct (Bai et al., 2025), Phi-4-mini-Instruct (Abouelenin et al., 2025), and Granite-3.1-3B-Instruct (Team et al., 2025). After a single epoch of full-parameter training, we retain only the DAVE encoder for weight-merging and downstream tasks. During the weight-merging distillation, we train the merge coefficient for 20 epochs on unlabeled documents and web images. Refer to Appendix X for implementation details.

4.2 EVALUATION SETTING

In this section, we first describe our encoder baselines, followed by two types of evaluations for the vision encoders: (i) finetuning vision encoders for classic document tasks, (ii) performing instruction tuning to build a VLM with the vision encoder for vision-language tasks and agentic tasks.

Baselines. We compare our vision encoder with both generalist and specialist encoders. For generalist encoders, we utilize DinoV2 (Oquab et al., 2023) and Web-SSL MAE (Fan et al., 2025), a MAE variant trained on 2 billion images. We also include SigLIP2 (Tschannen et al., 2025) and AIMv2 (Fini et al., 2025), which are SOTA encoders trained with both contrastive and reconstruction. For specialist vision encoders, we compare against DiT (Li et al., 2022a), Pix2Struct (Lee et al., 2023), and Dolphin (Feng et al., 2025). Since both Pix2struct and Dolphin are encoder-decoder models, we use their encoders for comparison.

Classic Document Tasks. For each classic document task, we finetune the vision encoder and the suitable prediction heads. In DocBank, we use attention pooling to pool the visual feature, followed by an MLP head to predict the bounding box. In Doclaynet, we train an MLP head to predict the semantic segmentation based on the feature patches. For RICO-SCA, we use attention pooling followed by an MLP to predict the class. Refer to Appendix C for more details.

Vision-Language Model as Evaluator. To evaluate our vision encoders on more complicated and realistic tasks, such as VQA and instruction-based localization, we use VLMs as evaluators. Specifically, we follow the standard LLaVA architecture as discussed in Section 3.1. We use a tilting size of 336 for AIMv2 and 384 for all other vision encoders. For vision encoders that produce a different number of visual tokens than the tilting resolution, we use bilinear interpolation as in LLaVA-Onevision (Li et al., 2024) to interpolate the visual tokens. Note that we omit the projector alignment phase to improve training efficiency, as prior work has shown that this stage yields only limited gains Tong et al. (2024); Karamcheti et al. (2024). We use Llama-3.2-3B-Instruct (Dubey et al., 2024) and Qwen2.5-7B-Instruc (Bai et al., 2025) as the LLM backbone. We train for one epoch with the vision encoder frozen. For Mind2web, we finetune the VLMs on the training set before performing offline evaluation on the test set. More detail can be found in Appendix D.

	Document			Ge	neral V	QA		Web			
Model / Variant	\$2P	O Regard	702	00000	Po deputo de la constante de l	Many	A CONTROLLER	A CONTRACTOR	None West	Spania	Wedsh.
			Llan	na-3.2-3	B-Instr	uct					
MAE-Web	53.0	27.6	43.8	26.2	43.8	29.9	48.6	35.1	36.9	51.1	46.8
DinoV2	53.2	2.6	13.7	20.4	14.1	35.0	49.4	13.7	13.2	16.5	17.6
DiT	49.9	2.1	11.3	19.2	12.6	28.8	43.9	10.0	9.2	2.5	14.9
Pix2Struct	51.0	6.9	22.8	21.4	21.5	33.3	45.0	16.4	22.6	32.6	26.5
Dophine	50.8	44.7	<u>74.8</u>	37.4	60.3	33.3	44.4	42.8	48.8	<u>56.3</u>	<u>76.0</u>
AIMv2	<u>56.1</u>	48.3	56.6	35.6	48.7	36.3	51.4	<u>58.7</u>	46.7	32.0	49.4
SigLIP 2	58.0	<u>51.5</u>	72.1	<u>40.6</u>	51.8	36.9	<u>53.5</u>	64.4	<u>54.7</u>	40.7	67.8
DAVE (ours)	59.6	62.2	82.1	50.2	63.1	<u>36.6</u>	55.6	69.2	59.2	64.5	82.6
			Qwe	en-2.5-7	B-Instr	uct					
MAE-Web	63.8	33.2	55.1	28.6	50.3	38.4	45.4	40.6	59.5	70.1	60.7
Pix2struct	58.7	10.0	29.8	20.0	28.3	37.9	39.1	18.2	52.3	47.8	30.3
Dophine	63.3	50.9	83.3	42.0	66.3	40.2	44.6	49.4	60.7	75.2	85.6
AIMv2	69.7	69.3	89.3	57.8	77.4	41.3	56.2	73.1	64.9	79.3	86.1
SigLIP 2	73.6	<u>67.4</u>	90.2	<u>55.2</u>	<u>75.9</u>	<u>43.9</u>	58.7	74.3	<u>65.4</u>	81.4	88.3
DAVE (ours)	74.0	67.5	90.9	60.2	82.5	45.8	<u>55.2</u>	<u>73.7</u>	67.3	82.9	88.6

Table 1: The DAVE's performance on Document understanding, general VQA, and Web understanding benchmarks with two VLM architectures using different LLMs. The best result per row is highlighted in **bold** and the second best with <u>underline</u>. Higher values represent better performance.

4.3 BENCHMARKS

In this section, we describe our evaluation datasets for several downstream benchmarks.

Classic Document Understanding Tasks. To directly evaluate the structural and spatial visual representation from the vision encoders, we consider several document parsing tasks that can be performed with only images. We employ DocBank (Li et al., 2020a) to comprehensively evaluate document recognition on different categories, including tables, charts, and paragraphs. We also use Doclaynet (Pfitzmann et al., 2022) to test document segmentation. For web tasks, we use RICO-SCA (Li et al., 2020b) to perform web UI classification.

Document and General VQA. For document understanding and question answering, we evaluate on AI2D (Kembhavi et al., 2016), OCRBench (Liu et al., 2024f), DocVQA (Mathew et al., 2021), InfoVQA (Mathew et al., 2022), and ChartQA (Masry et al., 2022). For general vision-language understanding, we report results on MMMU (Yue et al., 2024), RealWorldQA (xAI, 2024), and TextVQA (Singh et al., 2019).

Web UI and Agent Tasks. For web UI localization, we use Screenspot-V2 (Wu et al., 2024), which consists of text instructions (e.g., "click the button in cooridnate (x, y)"), and the corresponding bounding boxes. We report the center accuracy: the label-box center lies within the predicted box. We also evaluate on WebSRC (Chen et al., 2021) and VisualWebBench (Liu et al., 2024d) for web UI question answering. For the agentic benchmark, we evaluate on Mind2Web (Deng et al., 2023), where the model receives a user goal and a webpage screenshot, and must predict an action like clicking.

4.4 DATASETS

In this section, we provide a description of our training datasets for the different stages.

Self-supervised Learning Data. We use 20 million images for our self-supervised training. For document data, we sample 10 million PDF images from DocFM (Team et al., 2025), which contains 85 million document pages extracted from Common Crawl, Wikipedia, and ESG (Environmental,

			Mino	d2Web		
Model	Cross-Task		Cross-	-Website	Cross-Domain	
_	Step SR	Element Acc.	Step SR	Element Acc.	Step SR	Element Acc.
		Llar	na-3.2-3B-Ins	struct		
MAE	21.0	16.8	16.5	11.9	15.5	11.6
DinoV2	13.5	11.4	10.2	6.6	9.5	7.2
DiT	5.8	4.1	2.7	1.2	2.1	1.1
Pix2Struct	15.8	12.3	12.1	8.2	10.8	7.5
Dolphin	24.0	<u>19.6</u>	20.9	<u>13.6</u>	19.3	14.5
AIMv2	$\overline{14.7}$	11.6	9.2	5.9	$\frac{19.3}{8.8}$	6.1
SigLIP 2	20.6	17.3	11.8	8.7	12.8	9.7
DAVE (ours)	30.8	26.1	24.2	18.0	23.9	19.1

Table 2: **Results on Web Agent.** Performance on Mind2Web with three splits (Cross-Task, Cross-Website, Cross-Domain). We report the stepwise accuracy (correct grounding) and the element accuracy (correct grounding and action) for each task.

Social, and Governance) reports. The PDFs are filtered to include English only. For web screenshot, we sample 10 million from Common-Web (com, 2023) without any language filtering.

Supervised Learning Data. For the autoregressive supervised training, we leverage data from diverse sources, including PlotQA (Methani et al., 2020), ChartQA (Masry et al., 2022), fintabnet (Zheng et al., 2021), Datikz (Belouadi et al., 2023), Pubtables (Smock et al., 2022), and DocFM (Team et al., 2025). Importantly, we use the version from Granite Vision (Team et al., 2025) where all the problems are reformulated into content extraction. The data for this stage is a mixture of curated benchmarks and a large-scale, self-processed dataset. The curated benchmarks cover tasks such as chart-to-markdown, table-to-caption, and web UI grounding with UGround Gou et al. (2024). To expand on this, we processed 500K PDF samples from arXiv with an OCR model (Cui et al., 2025), formulating them into additional recognition and grounding tasks. Altogether, this training stage comprises approximately 2 million samples.

Instruction Tuning Data. Following Cambrian-1 (Tong et al., 2024), We use the LLaVA-1.5-mix 665K (Liu et al., 2023), DocVQA (Mathew et al., 2021), ChartQA (Masry et al., 2022), and AI2D (Kembhavi et al., 2016) as our instruction tuning data. We also integrate Pixmo-Doc (Yang et al., 2025) into the data mix to improve document understanding. Additionally, we add MultiUI (Liu et al., 2024c) to improve the web and agentic capability. This combined to a total of 2.5 million training images.


Model	DocLayNet	DocBank	RICO-SCA
DinoV2	68.4	38.3	85.6
MAE-Web	64.6	44.5	88.3
Dolphin	53.8	50.5	88.4
Pix2Struct	56.7	47.2	90.1
AIMv2	70.5	42.4	91.3
SipLIP 2	70.8	51.7	93.3
DAVE	74.1	56.9	<u>92.8</u>

Table 3: Performance comparison on **classic document tasks**. DocLayNet and DocBank use mAP, while RICO-SCA uses classification accuracy.

5 RESULTS

5.1 VISION-LANGUAGE TASKS

Table 1 shows the evaluation results for document, general, and web benchmarks. Overall, DAVE consistently outperforms the strongest baseline, SigLIP 2 in the Llama-3.2-3B-Instruct setup on 8 document and web benchmarks by an average of 10.5%. This is achieved without losing the general VQA capabilities, like MMMU and RealWorldQA, which implies that the ensemble training successfully merges the structural and spatial features with the general visual features. In addition, when using Qwen-2.5-7B-Instruct as the VLM decoder, we show that the improvements in benchmarks hold a similar trend. This suggests that our merging scheme effectively aligns with different text decoders. Finally, one interesting observation is that specialized models like Pix2Struct and






Figure 2: Each row indicates an additional modification to the training strategy.

Figure 3: Inter-patch standard deviation across different data sources.

Dolphin perform suboptimally compared to SigLIP 2 and AIMv2, which shows that both general visual features and document/web-specific features are critical to the performance. This hypothesis is further supported by DAVE's strong results.

5.2 Web Agentic Tasks

The results for multimodal-mind2web are outlined in Table 2. DAVE outperforms the best baseline encoder, Dolphin, by an average of 5%. Surprisingly, self-supervised vision encoders like MAE and DinoV2 perform competitively with SigLIP 2 and AIMv2, while document-specialized encoder Dolphine achieves the best accuracy among all baselines. This suggests that structural and spatial aspects may be more important in web navigation compared to broad general features.

5.3 CLASSIC DOCUMENT TASKS

As shown in Table 3, we evaluate DAVE on dense document tasks. DAVE outperforms both specialized and generalist encoders on DocBank and Doylaynet. This demonstrates the strong structural and spatial visual features learned by our 2-stage pretraining. Furthermore, DAVE performs competitively with SigLIP 2 on screen classification, a task that emphasizes semantic understanding and for which DAVE is not designed to tackle.

5.4 ABLATIONS

In this section, we present detailed ablations on the design choice of our training method. In all of the ablations, we report the average performance on document and web vision-language benchmarks, which we denote as Doc and Web in the Table 4.

Training Method. To assess the contribution of each design choice in the stage 2 supervised pretraining, we incrementally add each choice and plot the performance trajectory on document visionlanguage tasks tasks in Figure 2. Beginning with the a scratch text decoder and subsequently incorporating a pretrained LLM as the text decoder, ensemble training, and finally weight merging. The consistently improving trend highlights the positive impact of our proposed design choices.

Image Inter-Patch Variance. To show the instability of the MAE training on document and web image discussed in Section 3.3, we plot the distribution density of inter-patch variance across different data sources in Figure 3. Compared to ImageNet (Deng et al., 2009), document and web images exhibit much lower variance. This distributional gap underscores the need for a specialized vision encoder capable of handling such data. Additional implementation details are provided in Appendix E.

Different Merge Setting. As described in Section 3.3, we use weight merging to create a decoder-agnostic vision encoder. Table 4a examines the effect of different merge strategies. Our end-to-end approach for learning merge coefficients outperforms both simple averaging and the heuristic-based Fisher Merge (Matena & Raffel, 2022) method, which weights parameters by their estimated

Method	Doc	Web
No merge	55.6	53.0
Average	62.8	67.7
Fisher Merge	60.3	67.0
Learned Coef	63.4	68.2

Method	Doc	Web
From Scratch	52.2	54.7
Granite	55.6	53.0
Granite+Qwen	62.1	64.8
Granite+Qwen+Phi	63.4	68.2

⁽a) Comparison of Different Merge Methods.

(b)	Com (parison	of	Different	N.	lerging	LLMs.
-----	-------	---------	----	-----------	----	---------	-------

Method	Doc	Web
SIgLIP 2 + DiT	50.3	47.6
SIgLIP 2 + PS	48.9	42.2
SIgLIP 2 + DP	49.7	44.0
DAVE	63.4	68.2

Method	Doc	Web
SigLIP 2	49.1	54.4
Finetune SigLIP 2	58.2	65.2
DAVE	63.4	68.2

⁽c) Comparison of Different Specialized Encoder Fusion. Key: PS-Pix2struct, DP-Dolphin

(d) Comparison with Finetuning.

Table 4: **Ablation study**: (a) merge methods, (b) LLM merging, and additional analyses (c, d).

importance using the Fisher information. Table 4b evaluates merging DAVE from different LLM backbones, where performance improves consistently with the number of merged encoders.

Multi-encoder Comparison. Because DAVE combines with both generalist and specialist features, we ask how DAVE performs compared to multi-encoder settings. Specifically, we concatenate different document and web specialized models with SigLIP 2 during the instruction tuning of building the VLM. In Table 4c, we present the result of different multi-encoder settings. We highlight that without any feature fusion, all the specialized encoders fail to improve performance on document and web benchmarks.

Comparison with Finetuning. To test whether our pretraining approach is more effective than directly finetuning an existing vision encoder on the supervised data used in our pretraining, we finetune SigLIP 2 with our supervised pretraining data. In Table 4d, we show that DAVE outperforms finetuned SigLIP 2, demonstrating the effectiveness of our approach.

6 CONCLUSION

We introduced DAVE, a foundational vision encoder for document and web understanding. Our approach combines self-supervised learning on large-scale unlabeled data with supervision from limited vision—language annotations, enabling data-efficient training in specialized domains. By leveraging pretrained generalist encoders, DAVE can focus on learning domain-specific features. Extensive experiments demonstrate that strong visual representations are critical for document, web, and agentic tasks, where our encoder achieves significant improvements over prior models. Ablation studies further confirm the method's flexibility on different VLM and agentic frameworks.

7 LIMITATIONS AND FUTURE WORK

Despite the significant progress achieved in document and web understanding by DAVE, several limitations remain. The model operates at a fixed resolution, depends on an external pretrained generalist encoder, and requires substantial compute for self-supervised training. In agentic settings, its vision encoder is also not conditioned on prior actions. These limitations highlight opportunities for developing vision encoders that are more compute- and data-efficient, can process inputs of any resolution or aspect ratio, and incorporate prior agentic actions to produce more tailored representations. In addition, this paradigm offers a promising path for domains with scarce vision—language data, such as medical imaging. While we do not anticipate specific negative impacts, as with any machine learning method, caution should be exercised during deployment.

REFERENCES

- Common screens. https://registry.opendata.aws/comonscreens, 2023.
 - Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, et al. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*, 2025.
 - Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.
 - Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 2022.
 - Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
 - Jonas Belouadi, Anne Lauscher, and Steffen Eger. Automatikz: Text-guided synthesis of scientific vector graphics with tikz. *arXiv preprint arXiv:2310.00367*, 2023.
 - Daniel Bolya, Po-Yao Huang, Peize Sun, Jang Hyun Cho, Andrea Madotto, Chen Wei, Tengyu Ma, Jiale Zhi, Jathushan Rajasegaran, Hanoona Rasheed, et al. Perception encoder: The best visual embeddings are not at the output of the network. *arXiv preprint arXiv:2504.13181*, 2025.
 - Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
 - Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
 - Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PmLR, 2020.
 - Xingyu Chen, Zihan Zhao, Lu Chen, Danyang Zhang, Jiabao Ji, Ao Luo, Yuxuan Xiong, and Kai Yu. Websrc: A dataset for web-based structural reading comprehension. *arXiv* preprint *arXiv*:2101.09465, 2021.
 - Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024.
 - Cheng Cui, Ting Sun, Manhui Lin, Tingquan Gao, Yubo Zhang, Jiaxuan Liu, Xueqing Wang, Zelun Zhang, Changda Zhou, Hongen Liu, et al. Paddleocr 3.0 technical report. *arXiv preprint arXiv:2507.05595*, 2025.
 - Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pp. 248–255. Ieee, 2009.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su.
 Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36:28091–28114, 2023.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. arXiv e-prints, pp. arXiv–2407, 2024.
 - David Fan, Shengbang Tong, Jiachen Zhu, Koustuv Sinha, Zhuang Liu, Xinlei Chen, Michael Rabbat, Nicolas Ballas, Yann LeCun, Amir Bar, and Saining Xie. Scaling language-free visual representation learning. 2025.
 - Hao Feng, Shu Wei, Xiang Fei, Wei Shi, Yingdong Han, Lei Liao, Jinghui Lu, Binghong Wu, Qi Liu, Chunhui Lin, et al. Dolphin: Document image parsing via heterogeneous anchor prompting. *arXiv* preprint arXiv:2505.14059, 2025.
 - Enrico Fini, Mustafa Shukor, Xiujun Li, Philipp Dufter, Michal Klein, David Haldimann, Sai Aitharaju, Victor G Turrisi da Costa, Louis Béthune, Zhe Gan, et al. Multimodal autoregressive pre-training of large vision encoders. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 9641–9654, 2025.
 - Hiroki Furuta, Kuang-Huei Lee, Ofir Nachum, Yutaka Matsuo, Aleksandra Faust, Shixiang Shane Gu, and Izzeddin Gur. Multimodal web navigation with instruction-finetuned foundation models. *arXiv preprint arXiv:2305.11854*, 2023.
 - Boyu Gou, Ruohan Wang, Boyuan Zheng, Yanan Xie, Cheng Chang, Yiheng Shu, Huan Sun, and Yu Su. Navigating the digital world as humans do: Universal visual grounding for gui agents. *arXiv preprint arXiv:2410.05243*, 2024.
 - Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022.
 - Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. Cogagent: A visual language model for gui agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14281–14290, 2024.
 - Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang Zhang, Bo Zhang, Chen Li, Ji Zhang, Qin Jin, Fei Huang, et al. mplug-docowl 1.5: Unified structure learning for ocr-free document understanding. arXiv preprint arXiv:2403.12895, 2024.
 - Hang Hua, Yunlong Tang, Chenliang Xu, and Jiebo Luo. V2xum-llm: Cross-modal video summarization with temporal prompt instruction tuning. *arXiv preprint arXiv:2404.12353*, 2024a.
 - Hang Hua, Yunlong Tang, Ziyun Zeng, Liangliang Cao, Zhengyuan Yang, Hangfeng He, Chenliang Xu, and Jiebo Luo. Mmcomposition: Revisiting the compositionality of pre-trained vision-language models. *arXiv preprint arXiv:2410.09733*, 2024b.
 - Hang Hua, Qing Liu, Lingzhi Zhang, Jing Shi, Soo Ye Kim, Zhifei Zhang, Yilin Wang, Jianming Zhang, Zhe Lin, and Jiebo Luo. Finecaption: Compositional image captioning focusing on wherever you want at any granularity. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 24763–24773, 2025.
 - Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. Layoutlmv3: Pre-training for document ai with unified text and image masking. In *Proceedings of the 30th ACM international conference on multimedia*, pp. 4083–4091, 2022.
 - Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35 (1):73–101, 1964.
 - Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.

- Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa
 Sadigh. Prismatic vlms: Investigating the design space of visually-conditioned language models.
 In Forty-first International Conference on Machine Learning, 2024.
 - Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *Computer Vision ECCV 2016*, pp. 235–251, Cham, 2016. Springer. doi: 10.1007/978-3-319-46493-0_15.
 - Geewook Kim, Teakgyu Hong, Moonbin Yim, Jinyoung Park, Jinyoong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. Donut: Document understanding transformer without ocr. *arXiv preprint arXiv:2111.15664*, 7(15):2, 2021.
 - Geewook Kim, Teakgyu Hong, Moonbin Yim, JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. Ocr-free document understanding transformer. In *European Conference on Computer Vision*, pp. 498–517. Springer, 2022.
 - Geunwoo Kim, Pierre Baldi, and Stephen McAleer. Language models can solve computer tasks. *Advances in Neural Information Processing Systems*, 36:39648–39677, 2023.
 - Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. Pix2struct: Screenshot parsing as pretraining for visual language understanding. In *International Conference on Machine Learning*, pp. 18893–18912. PMLR, 2023.
 - Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.
 - Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023a.
 - Junlong Li, Yiheng Xu, Tengchao Lv, Lei Cui, Cha Zhang, and Furu Wei. Dit: Self-supervised pre-training for document image transformer. In *Proceedings of the 30th ACM international conference on multimedia*, pp. 3530–3539, 2022a.
 - Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pretraining for unified vision-language understanding and generation. In *International conference on machine learning*, pp. 12888–12900. PMLR, 2022b.
 - Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023b.
 - Minghao Li, Yiheng Xu, Lei Cui, Shaohan Huang, Furu Wei, Zhoujun Li, and Ming Zhou. Docbank: A benchmark dataset for document layout analysis. *arXiv preprint arXiv:2006.01038*, 2020a.
 - Yang Li, Jiacong He, Xin Zhou, Yuan Zhang, and Jason Baldridge. Mapping natural language instructions to mobile ui action sequences. *arXiv* preprint arXiv:2005.03776, 2020b.
 - Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023.
 - Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024a. URL https://llava-vl.github.io/blog/2024-01-30-llava-next/.
 - Junpeng Liu, Tianyue Ou, Yifan Song, Yuxiao Qu, Wai Lam, Chenyan Xiong, Wenhu Chen, Graham Neubig, and Xiang Yue. Harnessing webpage uis for text-rich visual understanding. *arXiv* preprint arXiv:2410.13824, 2024b.
 - Junpeng Liu, Tianyue Ou, Yifan Song, Yuxiao Qu, Wai Lam, Chenyan Xiong, Wenhu Chen, Graham Neubig, and Xiang Yue. Harnessing webpage uis for text-rich visual understanding, 2024c. URL https://arxiv.org/abs/2410.13824.

- Junpeng Liu, Yifan Song, Bill Yuchen Lin, Wai Lam, Graham Neubig, Yuanzhi Li, and Xiang Yue. Visualwebbench: How far have multimodal llms evolved in web page understanding and grounding? *arXiv preprint arXiv:2404.05955*, 2024d.
 - Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pp. 216–233. Springer, 2024e.
 - Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xucheng Yin, Cheng-lin Liu, Lianwen Jin, and Xiang Bai. Ocrbench: on the hidden mystery of ocr in large multimodal models. *Science China Information Sciences*, 67(12):220102, December 2024f. ISSN 1869-1919. doi: 10.1007/s11432-024-4235-6. URL http://dx.doi.org/10.1007/s11432-024-4235-6.
 - Yuliang Liu, Biao Yang, Qiang Liu, Zhang Li, Zhiyin Ma, Shuo Zhang, and Xiang Bai. Textmonkey: An ocr-free large multimodal model for understanding document. *arXiv* preprint *arXiv*:2403.04473, 2024g.
 - Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
 - Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv* preprint arXiv:2310.02255, 2023.
 - Tengchao Lv, Yupan Huang, Jingye Chen, Yuzhong Zhao, Yilin Jia, Lei Cui, Shuming Ma, Yaoyao Chang, Shaohan Huang, Wenhui Wang, et al. Kosmos-2.5: A multimodal literate model. *arXiv* preprint arXiv:2309.11419, 2023.
 - Kevis-Kokitsi Maninis, Kaifeng Chen, Soham Ghosh, Arjun Karpur, Koert Chen, Ye Xia, Bingyi Cao, Daniel Salz, Guangxing Han, Jan Dlabal, et al. Tips: Text-image pretraining with spatial awareness. *arXiv preprint arXiv:2410.16512*, 2024.
 - Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv* preprint *arXiv*:2203.10244, 2022.
 - Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022.
 - Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 2199–2208, 2021. doi: 10.1109/WACV48630.2021.00225.
 - Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and C.V. Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1697–1706, January 2022.
 - Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. Plotqa: Reasoning over scientific plots. In *Proceedings of the ieee/cvf winter conference on applications of computer vision*, pp. 1527–1536, 2020.
 - Muhammad Ferjad Naeem, Yongqin Xian, Xiaohua Zhai, Lukas Hoyer, Luc Van Gool, and Federico Tombari. Silc: Improving vision language pretraining with self-distillation. In *European Conference on Computer Vision*, pp. 38–55. Springer, 2024.
 - Ahmed Nassar, Andres Marafioti, Matteo Omenetti, Maksym Lysak, Nikolaos Livathinos, Christoph Auer, Lucas Morin, Rafael Teixeira de Lima, Yusik Kim, A Said Gurbuz, et al. Smoldocling: An ultra-compact vision-language model for end-to-end multi-modal document conversion. *arXiv* preprint arXiv:2503.11576, 2025.
 - Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv* preprint arXiv:2304.07193, 2023.

- Birgit Pfitzmann, Christoph Auer, Michele Dolfi, Ahmed S Nassar, and Peter Staar. Doclaynet: A large human-annotated dataset for document-layout segmentation. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 3743–3751, 2022.
- Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, et al. Ui-tars: Pioneering automated gui interaction with native agents. *arXiv preprint arXiv:2501.12326*, 2025.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.
- Claude E Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27 (3):379–423, 1948.
- Min Shi, Fuxiao Liu, Shihao Wang, Shijia Liao, Subhashree Radhakrishnan, Yilin Zhao, De-An Huang, Hongxu Yin, Karan Sapra, Yaser Yacoob, et al. Eagle: Exploring the design space for multimodal llms with mixture of encoders. *arXiv preprint arXiv:2408.15998*, 2024.
- Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv* preprint arXiv:2508.10104, 2025.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Brandon Smock, Rohith Pesala, and Robin Abraham. Pubtables-1m: Towards comprehensive table extraction from unstructured documents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4634–4642, 2022.
- Yunlong Tang, Daiki Shimada, Jing Bi, and Chenliang Xu. Avicuna: Audio-visual llm with interleaver and context-boundary alignment for temporal referential dialogue. *arXiv* preprint arXiv:2403.16276, 2024.
- Yunlong Tang, Junjia Guo, Hang Hua, Susan Liang, Mingqian Feng, Xinyang Li, Rui Mao, Chao Huang, Jing Bi, Zeliang Zhang, et al. Vidcomposition: Can mllms analyze compositions in compiled videos? In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 8490–8500, 2025.
- Zineng Tang, Ziyi Yang, Guoxin Wang, Yuwei Fang, Yang Liu, Chenguang Zhu, Michael Zeng, Cha Zhang, and Mohit Bansal. Unifying vision, text, and layout for universal document processing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19254–19264, 2023.
- Granite Vision Team, Leonid Karlinsky, Assaf Arbelle, Abraham Daniels, Ahmed Nassar, Amit Alfassi, Bo Wu, Eli Schwartz, Dhiraj Joshi, Jovana Kondic, et al. Granite vision: a lightweight, open-source multimodal model for enterprise intelligence. *arXiv preprint arXiv:2502.09927*, 2025.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*, 2024.
- Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pp. 23965–23998. PMLR, 2022.

- Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yian Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, et al. Os-atlas: A foundation action model for generalist gui agents. *arXiv* preprint arXiv:2410.23218, 2024.
 - xAI. Realworldqa: A benchmark for evaluating real-world spatial understanding in multimodal ai models. https://huggingface.co/datasets/xai-org/RealworldQA, 2024. Accessed: 2025-09-17.
 - Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, et al. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding. *arXiv preprint arXiv:2012.14740*, 2020a.
 - Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pretraining of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1192–1200, 2020b.
 - Yiheng Xu, Zekun Wang, Junli Wang, Dunjie Lu, Tianbao Xie, Amrita Saha, Doyen Sahoo, Tao Yu, and Caiming Xiong. Aguvis: Unified pure vision agents for autonomous gui interaction. *arXiv* preprint arXiv:2412.04454, 2024.
 - Yue Yang, Ajay Patel, Matt Deitke, Tanmay Gupta, Luca Weihs, Andrew Head, Mark Yatskar, Chris Callison-Burch, Ranjay Krishna, Aniruddha Kembhavi, et al. Scaling text-rich image understanding via code-guided synthetic multimodal data generation. *arXiv preprint arXiv:2502.14846*, 2025.
 - Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
 - Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *National Science Review*, 11(12):nwae403, 2024.
 - Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
 - Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multi-modal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9556–9567, 2024.
 - Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 11975–11986, 2023.
 - Xinyi Zheng, Douglas Burdick, Lucian Popa, Xu Zhong, and Nancy Xin Ru Wang. Global table extractor (gte): A framework for joint table identification and cell structure recognition using visual context. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 697–706, 2021.
 - Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.
 - Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv* preprint arXiv:2504.10479, 2025.

Hyperparameter	Value
Batch Size	4096
Learning Rate	1e-5
Epochs	25
Warmup Epochs	5
Weight Decay	0.05
Mask Ratio	0.75
LR Scheduler	Cosine

Table 5: Training hyperparameters for MAE training.

A STAGE 1: SELF-SUPERVISED TRAINING

Here, we present the implementation detail and dataset used for the self-supervised training.

A.1 IMPLEMENTATION DETAILS

We use the official implementation of masked autoencoder (MAE) for our training, and we use 32 H200 GPUs for the training. Our vision encoder is a ViT-L-384 initialized from scratch, whereas our decoder is a transformer model with a depth of 4 and 16 attention heads. The detail training hyperparameter can be found in Table 5

A.2 DATASETS

Our self-supervised pretraining dataset consists of 10 million document images from DocFM and 10 million web screenshots from Common Screen

DocFM. DocFM is a large scale document data collected by IBM consisting of 85 million document pages extracted from unique PDF documents sourced from Common Crawl, Wikipedia, and ESG (Environmental, Social, and Governance) reports. From this dataset, we first filter and keep only the English pdfs, then we randomly sample 10 million pdfs and store them as images.

Common Screen. Common Screen is a large scale web screenshot data consisting of 70 million screenshot images based on the Common Crawl data. We randomly sample 10 million images as our training data, without any language filtering.

A.3 ADDITIONAL RESULTS ON MAE

Model	DocBank	Doclaynet	RICO-SCA
MAE-Web	64.6	44.5	88.3
MAE-Doc	72.8	52.5	90.7

Table 6: Evaluation results on classic document tasks.

In this section, we demonstrates the effectiveness of pretraining on document and web images by comparing our stage-1 encoder, which we denote as MAE-Doc, with the pretrained MAE-Web introduced in 4.2 in classic document tasks. In Table 6, the superior performance of MAE-Doc over MAE-Web highlights importance of the domain-specific pretraining for document and web tasks.

B STAGE 2: SUPERVISED PRETRAINING

In this section, we present the detail implementation and the dataset composition of our supervised pretraining after the self-supervised pretraining.

Hyperparameter	Value	
Max sequence length	20000	
Learning Rate	3e-5	
Epochs	1	
Warmup ratio	0.03	
LR Scheduler	Cosine	

Table 7: Training hyperparameters for supervised pretraining.

B.1 IMPLEMENTATION DETAILS

We use the official LLaVA-Next (Li et al., 2024) repo with some customization to support different LLM and encoder settings, and we use 32 H200 GPUs for the training. To handle high resolution images, we use AnyRes image tilting with size of 384x384. The hyperparameters are listed in Table 7.

B.2 DATASETS

Dataset	#Image	Description	
FM4D	250k	Doc/Chart/Table Extraction	
PlotQA	146k	Plot Extraction	
Fintabnet	88k	Table Extraction	
ChartQA	17k	Chart Extraction	
Datikz	94k	Image to Latex	
Pubtables	480k	Table to HTML	
UGround	756k	Web UI Grounding	
Self-curated Data	250k	Doc Grounding/structure Recognition	

Table 8: Dataset description for supervised pretraining.

We use a wide range of publicly available data for our supervised pretraining. Notably, all the data excluding UGround are reformulated into layout and information extraction. For more detail on the data curation, refer to Granite Vision (Team et al., 2025). For the Self-curated Data, we use PaddleOCR to first parse the document pdf from arxiv, then we use the bounding box and ocr text to formulate grounding and recognition data. The data composition is shown in Table 8.

B.3 Weight-Merging

For the weight-merging, we randomly sample 10k document and web images from self-supervised pretraining as our training data. We use the AdamW Optimizer (Loshchilov & Hutter, 2017), a learning rate of 1e-4, and a batch size of 32 to train for 20 epoches. The only trainable component is the newly introduced merge coefficient, while all the vision encoders remain frozen.

C CLASSIC DOCUMENT AND WEB TASKS

C.1 DOCBANK

DocBank is a large-scale benchmark for document layout analysis, providing token-level annotations that capture both textual and structural information. It is widely used to evaluate models on tasks such as text detection, segmentation, and understanding of document formatting. In our work, we formulate DocBank as Document element recognition, where the task is the predict the bounding box of a specific element given a document image.

Implementation. DocBank consists of 13 categories. For each category, we finetune the vision encoder with an attention pooling network on the corresponding training set, and perform evaluation on the validation set. Before training, each image is processed and resized according to the default preprocessor for the specific vision encoder. For each category, we train with 5 epoch, a learning

Model	Abstract	Author	Caption	Equation	Figure	Footer
DinoV2	86.4	9.0	45.2	14.9	90.5	3.6
MAE	84.1	17.9	20.7	33.4	90.8	23.8
Dolphin	73.8	21.6	48.1	25.7	91.5	32.9
Pix2Struct	80.6	26.3	21.6	15.3	83.1	27.7
AIMv2	85.8	5.5	27.7	15.5	90.4	19.4
SigLIP 2	85.9	23.1	52.6	34.7	91.6	31.6
DAVE	87.4	35.3	54.1	53	91.5	17.9
Model	List	Paragraph	Reference	Section	Table	Title
DinoV2	42.2	10.3	86.1	3.2	68.5	0.0
MAE	40.8	10.2	88.8	9.2	88.1	26.5
Dolphin	68.7	11.5	88.0	9.7	84.3	50.2
Pix2Struct	33.9	13.7	87.3	5.4	62.2	64.4
AIMv2	45.8	9.8	86.2	2.8	80.3	39.0
SigLIP 2	51.8	13.2	88.7	7.5	90.9	48.8
DAVE	66.4	20.6	91.6	27	87.9	49.7

Table 9: Evaluation results on DocBank by category.

rate of 1e-5, weight decay of 1e-4, and a batch size of 32. We use Smooth L1 objective (Huber, 1964) for the training. The per-category result is shown in Table 9.

C.2 DOCLAYNET

DocLayNet is a large-scale dataset for document layout analysis that contains richly annotated page images collected from diverse sources such as scientific articles, reports, and business documents. It provides pixel-level segmentation masks for a broad set of layout elements, enabling fine-grained evaluation of models on tasks like detection, segmentation, and structural understanding. In our work, we use Doclaynet to perform semantic segmentation.

Implementation. Doclaynet consists of 11 distinct classes, and in our implementation, we introduce an additional class for the area with no class label. We finetune the vision encoder along with an MLP projector to project the visual feature patches into a segmentation mask. Then we use crossentropy (Shannon, 1948) between the prediction and the ground truth segmentation as objective and train on the training set. We train for 2 epoch, with a learning rate of 1e-5, weight decay of 1e-4, and a batch size of 32.

C.3 RICO-SCA

RICO-SCA is a benchmark derived from the RICO dataset, which contains large-scale mobile application user interfaces. The SCA (Semantic Component Annotation) extension enriches each screen with detailed semantic labels for UI elements such as class type, buttons, images, text fields, and navigation components. In our work, we use RICO-SCA to evaluate on web UI classification.

Implementation. We finetune the vision encoder together with an attention pooling network to classify the class type of a given web UI image, using the provided training set. We train for 3 epoch, a learning rate of 1e-5, a weight decay of 1e-4, and a batch size of 32. We then evaluate the finetuned model on the validation set.

D VLM AS EVALUATOR

Here we provide the detail implementation and data composition on using VLMs as evaluator for vision encoders.

Data Composition. We use a set of instruction tuning data with a focus on document and web understanding for the VLM training. We then use the Mind2Web training set to finetune our VLM on web agentic tasks. The full data composition can be found in Table 10

Dataset	#Image	Description
AI2D	2k	Chart understanding
DocVQA	10k	Document understanding
InfoVQA	2k	Infographic reasoning
ChartQA	17k	Chart understanding
Pixmo-Doc	250k	Document understanding
MultiUI	2m	Web localization and understanding
LLaVA 665K	340k	General visual instruction data

Table 10: Dataset description for VLM training.

Implementation. The training is conducted with 32 H200 GPUs. We use the same architecture as described in Section B. For image tilting, we use tilting to split the image into size of 336x336 for AIMv2, and 384x384 for other vision encoders. The instruction tuning hyperparameter and finetuning is the same as in 7, except that Mind2Web is finetuned for 2 epoch.

E ABLATIONS

 In this section, we provide the implementation and evaluation details of our ablation studies, along with additional ablations on DAVE.

E.1 INTER-PATCH VARIANCE

We randomly sample 1000 document and web images from our pretraining data to conduct the analysis. For each image, we divide it into non-overlapping 16×16 patches and compute the perpatch standard deviation of pixel intensities after normalization. This provides a distribution of inter-patch variance that reflects the inherent local variability in different domains. As discussed in Section 5.4, document-like images exhibit consistently lower variance across patches, while web-like and natural images demonstrate a broader spread. These differences highlight the difference in low-level representation, particularly structural and spatial, between natural and document/web images.

E.2 MULTI-ENCODER SETUP

Given a setup consisting of SigLIP 2 and a specialized encoder, the images are first processed separately by their respective preprocessor. The processed images are then passed into the corresponding vision encoders to obtain visual patch features. Finally, the patch features are concatenated channelwise. In scenarios where the specialized vision encoder produces a different number of patch features compared to SigLIP 2, we apply linear interpolation on the specialized features to match the patch count.

Model	Doc	Web
DAVE	63.4	68.2
DAVE-spec + AIMv2	64.5	64.8

Table 11: Evaluation results on document and web understanding.

E.3 GENERALIZATION TO UNSEEN GENERALIST ENCODER

We denote the specialized encoder within DAVE as DAVE-spec. To evaluate whether DAVE-spec can adopt to unseen generalist encoder, we pair it with with AIMv2 during instruction tuning. The result is show in Table 11. The comparable performance with DAVE highlights the generalizability of DAVE-spec to novel vision-language and agent frameworks.