
Vision Encoders in Vision-Language Models: A Survey

Han Xiao
 Jina AI *by* Elastic
 han.xiao@jina.ai

Abstract

Vision encoders have remained comparatively small while language models scaled from billions to hundreds of billions of parameters. This survey analyzes vision encoders across 70+ vision-language models from 2023–2025 and finds that training methodology matters more than encoder size: improvements in loss functions, data curation, and feature objectives yield larger gains than scaling by an order of magnitude. Native resolution handling improves document understanding, and multi-encoder fusion captures complementary features no single encoder provides. We organize encoders into contrastive, self-supervised, and LLM-aligned families, providing a taxonomy and practical selection guidance for encoder design and deployment.

1 Introduction

Vision-language models have achieved strong performance, yet an asymmetry defines their architecture. While language models scaled from billions to hundreds of billions of parameters between 2020 and 2024, vision encoders remained largely frozen in time. The same 300–600 million parameter CLIP variants that powered early VLMs still dominate production systems today. Chen et al. [15] observed that “the progress in vision and vision-language foundation models has not kept pace with LLMs.” This asymmetry raises a fundamental question: *does the vision encoder matter?*

The answer, as this survey reveals, is nuanced. Encoder choice significantly impacts performance on vision-centric tasks such as document understanding, spatial reasoning, and fine-grained recognition, while mattering less for tasks where language reasoning dominates. More surprisingly, *how* an encoder is trained matters more than *how large* it is: a 400M-parameter SigLIP₂ encoder outperforms a 5.9B-parameter InternViT-6B on most VLM benchmarks. Understanding these trade-offs is essential for practitioners selecting encoders and researchers designing the next generation of vision-language systems.

The field originated with OpenAI’s CLIP [58] in 2021, which established contrastive image-text pretraining as the dominant paradigm. LLaVA [46] in 2023 demonstrated that connecting a frozen CLIP encoder to a language model through a simple projection could yield effective multimodal capabilities, as illustrated in Figure 1. For the next two years, most VLMs reused the same CLIP ViT-L/14 encoder while scaling only the language component.

Recent work has begun addressing this gap. Google’s SigLIP [88] replaced CLIP’s softmax-based contrastive loss with a sigmoid formulation that improved scaling and zero-shot performance. Shanghai AI Lab’s InternViT-6B [15] scaled the vision encoder to six billion parameters. By 2025, SigLIP₂ [70] added multilingual capabilities and dense features, becoming the encoder of choice for Qwen3-VL and Gemma 3. Cambrian-1 [69] showed that combining multiple encoders captures complementary visual information, while EVE [21] demonstrated that encoder-free architectures achieve competitive performance.

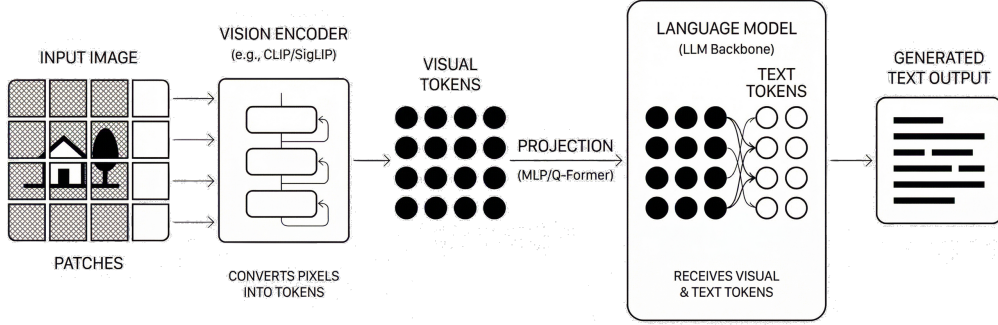


Figure 1: The canonical VLM architecture. An input image is divided into patches and processed by a vision encoder (e.g., CLIP, SigLIP) that converts pixels into visual tokens. A projection module (MLP or Q-Former) maps these tokens into the language model’s embedding space, where they are concatenated with text tokens. The language model generates text output conditioned on both visual and textual inputs.

Research Questions

1. Which training paradigm (contrastive, self-supervised, LLM-aligned) yields the best VLM performance?
2. When does encoder scale matter, and when does training methodology dominate?
3. How should variable-resolution images be handled at the encoder level?
4. When do multi-encoder approaches outperform single encoders?
5. What is the trajectory of vision encoding: specialized encoders or encoder-free unification?

Section 2 addresses questions 1–3 through paradigm analysis. Section 3 provides context for questions 4–5 through adoption patterns. Section 4 answers all five questions with empirical evidence, and Section 5 synthesizes the answers into practical guidance.

Contributions. This survey makes four contributions: (1) We present the first systematic taxonomy of vision encoders for VLMs, organizing 70+ models by training paradigm, integration architecture, and resolution strategy. (2) We provide quantitative evidence that training methodology improvements outperform parameter scaling: SigLIP 2 at 400M parameters exceeds InternViT-6B at 5.9B on most VLM tasks. (3) We compile reference tables (Appendix A) enabling encoder selection based on application requirements. (4) We identify the emerging tension between specialized encoders and encoder-free architectures as the field’s central design question.

The remainder of this survey is organized as follows. Section 2 establishes architectural foundations and training paradigms. Section 3 documents encoder adoption across VLM families. Section 4 answers the central question of what works through quantitative comparison. Section 5 synthesizes findings and identifies future directions. Table 1 provides a reference of VLMs organized by encoder family.

Table 1: VLMs Organized by Vision Encoder Family. Covers models through late 2025; includes peer-reviewed publications, arXiv preprints with released code/weights, and official technical reports. Unreleased models marked as preview. *Encoder trained from scratch by the model developers.

Model	Vision Encoder	Enc. Params	Date	Organization
OpenAI CLIP Family [58]				
LLaVA [46]	CLIP ViT-L/14	304M	2023/04	UW-Madison
LLaVA-1.5 [46]	CLIP ViT-L/14-336	304M	2023/10	UW-Madison
LLaVA-NeXT [45]	CLIP ViT-L/14	304M	2024/01	ByteDance

Continued on next page

Table 1 – continued

Model	Vision Encoder	Enc. Params	Date	Organization
Phi-3.5-Vision [1]	CLIP ViT-L/14	304M	2024/04	Microsoft
Phi-4-Vision [1]	CLIP ViT-L/14+	304M	2024/12	Microsoft
MM1 [52]	CLIP ViT-H/14	632M	2024/03	Apple
MM1.5 [52]	CLIP ViT-H/14	632M	2024/09	Apple
Ferret [84]	CLIP ViT-L/14	304M	2023/10	Apple
Ferret-v2 [90]	CLIP ViT-L/14	304M	2024/04	Apple
Yi-VL [85]	CLIP ViT-H/14	632M	2024/03	01.AI
Molmo [19]	CLIP ViT-L/14	304M	2024/09	AI2
VL-Mamba [57]	CLIP ViT-L/14	304M	2024/03	CASIA/Adelaide
Google SigLIP Family [88, 70]				
LLaVA-OneVision [39]	SigLIP-S0400M	400M	2024/08	ByteDance
DeepSeek-VL [47]	SigLIP-L + SAM-B	393M	2024/03	DeepSeek
DeepSeek-VL2 [79]	SigLIP-S0400M	400M	2024/12	DeepSeek
PaliGemma [9]	SigLIP-S0400M	400M	2024/07	Google
PaliGemma 2 [9]	SigLIP-S0400M	400M	2024/12	Google
Gemma 3 [25]	SigLIP-S0400M	400M	2025/03	Google
Qwen3-VL [6]	SigLIP 2 S0400M	400M	2025/11	Alibaba
Idefics2 [36]	SigLIP-S0400M	400M	2024/04	HuggingFace
Idefics3 [35]	SigLIP-S0400M	400M	2024/08	HuggingFace
SmolVLM [51]	SigLIP-S0400M	400M	2025/04	HuggingFace
Nemotron Nano V2 [20]	SigLIP-S0400M	400M	2025/11	NVIDIA
MiniCPM-V 2.6 [82]	SigLIP-S0400M	400M	2024/08	Tsinghua
Baichuan-Omni [42]	SigLIP-S0400M-384	400M	2024/10	Baichuan
VideoLLaMA 3 [89]	SigLIP/DFN	400M	2025/01	DAMO Academy
MiniCPM-V 4.5 [82]	SigLIP + 3D-Resampler	400M	2025/09	Tsinghua
Jina-VLM [34]	SigLIP 2 S0400M	400M	2025/12	Jina AI
Janus-Pro [13]	SigLIP-L	303M	2025/01	DeepSeek
EVA-CLIP Family [66]				
CogVLM [76]	EVA2-CLIP-E	4.4B	2023/11	Tsinghua
CogVLM2 [29]	EVA2-CLIP-E	4.4B	2024/05	Tsinghua
Emu [67]	EVA-CLIP-g	1B	2023/07	BAAI
Emu2 [64]	EVA-CLIP-E	4.4B	2023/12	BAAI
InternViT Family [15]				
InternVL 1.0 [15]	InternViT-6B	6B	2023/12	Shanghai AI Lab
InternVL 1.5 [15]	InternViT-6B	6B	2024/04	Shanghai AI Lab
InternVL 2.0 [14]	InternViT-300M/6B	300M–6B	2024/07	Shanghai AI Lab
InternVL 2.5 [14]	InternViT-300M/6B	300M–6B	2024/12	Shanghai AI Lab
InternVL 3.5 [75]	InternViT + ViR	300M–6B	2025/08	Shanghai AI Lab
NVLM-D/X/H [17]	InternViT-6B	6B	2024/09	NVIDIA

Continued on next page

Table 1 – continued

Model	Vision Encoder	Enc. Params	Date	Organization
Custom/Proprietary Encoders				
Qwen-VL [4]	ViT-bigG	1.9B	2023/08	Alibaba
Qwen2-VL [74]	NaViT	675M	2024/10	Alibaba
GLM-4.1V-Thinking [26]	AIMv2-Huge	600M	2025/07	Zhipu AI
GLM-4.5V [26]	AIMv2-Huge	600M	2025/08	Zhipu AI
MiniMax-VL-01 [38]	ViT*	303M	2025/01	MiniMax
Qwen2.5-VL [5]	NaViT	675M	2025/02	Alibaba
Llama 3.2-Vision [27]	ViT-H/14	632M	2024/09	Meta
Llama 4	MetaCLIP	–	2025/04	Meta
Kimi-VL [33]	MoonViT	400M	2025/04	Moonshot AI
DeepSeek-OCR [78]	DeepEncoder	380M	2025/10	DeepSeek
FastVLM [72]	FastViTHD	125M	2024/12	Apple
Pixtral 12B [2]	ViT*	400M	2024/10	Mistral
Ovis2.5 [48]	NaViT	300M	2025/08	Alibaba
MiMo-VL [87]	NaViT	675M	2025/06	Xiaomi
Step-3 [63]	EVA-CLIP 5B	5B	2025/07	StepFun
QVQ-72B-Preview [5]	NaViT	675M	2024/12	Alibaba
Multi-Encoder Approaches				
Cambrian-1 [69]	CLIP, SigLIP, ConvNeXt, DINOv2	1.9B	2024/06	NYU
Eagle [61]	CLIP, ConvNeXt, Pix2Struct, EVA-02	1.8B	2024/08	NVIDIA
Cobra [92]	DINOv2+SigLIP	704M	2024/03	Westlake U.
Encoder-Free / Native Multimodal				
Fuyu-8B [8]	Linear proj.	–	2023/10	Adept
EVE [21]	PEL	–	2024/06	BAAI
EVEv2 [22]	PEL	–	2025/02	BAAI
ELVA [41]	Token merge	–	2025/03	CAS
Chameleon [10]	VQ-VAE	–	2024/05	Meta
Emu3 [77]/3.5	VQ-VAE	–	2024/09	BAAI

2 Training Paradigms and Architectural Foundations

This section establishes the foundations for understanding vision encoders in VLMs. We first clarify scope and definitions, then cover the Vision Transformer architecture, three training paradigms (contrastive, self-supervised, LLM-aligned), connector designs, resolution strategies, and alternative architectures.

2.1 Scope and Definitions

Definition: Vision Encoder for VLMs

A *vision encoder* is a neural network component that transforms pixel inputs into dense vector representations consumed by a language model. It operates before the vision-language connector and produces a sequence of visual tokens $\mathbf{Z}_v \in \mathbb{R}^{N \times D}$, where N is the number of spatial tokens and D is the embedding dimension.

This survey focuses on vision encoders within vision-language model pipelines. We distinguish vision encoders from related but distinct components:

Vision Encoder vs. Image Tokenizer. Discrete tokenizers like VQ-VAE [71] map images to code-book indices for autoregressive generation (Chameleon, Emu3). We cover these in Section 3.7 as encoder-free alternatives but focus primarily on continuous-output encoders.

Vision Encoder vs. Raw Pixel Embedding. Encoder-free architectures (Fuyu-8B, EVE) project patches directly into language model space without pretrained vision components. These bypass the encoder entirely; we discuss them in Section 3.7.

Vision Encoder vs. Generation Encoder. Image generation models (Stable Diffusion, DALL-E) use encoders for reconstruction objectives. Our scope is encoders optimized for *understanding* within VLM pipelines, not generation.

2.2 Vision Transformer Architecture

The vision transformer (ViT) architecture underlies nearly all modern vision encoders for VLMs. A ViT first partitions an input image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ into a sequence of non-overlapping patches $\{\mathbf{x}_p^i\}_{i=1}^N$, where $N = HW/P^2$ for patch size P . These patches are linearly projected and combined with positional embeddings:

$$\mathbf{z}_0 = [\mathbf{x}_{\text{cls}}; \mathbf{E}\mathbf{x}_p^1; \mathbf{E}\mathbf{x}_p^2; \dots; \mathbf{E}\mathbf{x}_p^N] + \mathbf{E}_{\text{pos}} \quad (1)$$

where $\mathbf{E} \in \mathbb{R}^{D \times (P^2 \cdot C)}$ is the patch embedding projection, \mathbf{x}_{cls} is a learnable class token, and $\mathbf{E}_{\text{pos}} \in \mathbb{R}^{(N+1) \times D}$ provides positional information. The resulting sequence is processed through L transformer layers, each applying multi-head self-attention followed by a feed-forward network.

This architecture produces N spatial tokens plus one class token that capture visual information at different image locations, as shown in Figure 2. For VLM applications, these tokens are projected into the language model’s embedding space, enabling the language model to attend to visual features. The patch size P and image resolution together determine the number of visual tokens: a 336×336 image with $P = 14$ produces 576 tokens, while higher resolutions or smaller patches yield proportionally more tokens with associated computational costs.

2.3 Contrastive Learning

Contrastive training, pioneered by CLIP [58], trains vision and text encoders jointly on image-text pairs, learning to align matching pairs in a shared embedding space. Given a batch $\mathcal{B} = \{(\mathbf{I}_i, \mathbf{T}_i)\}_{i=1}^{|\mathcal{B}|}$ of image-text pairs, let $\mathbf{x}_i = f(\mathbf{I}_i)/\|f(\mathbf{I}_i)\|$ and $\mathbf{y}_i = g(\mathbf{T}_i)/\|g(\mathbf{T}_i)\|$ denote the normalized embeddings from the image encoder f and text encoder g . The CLIP objective minimizes:

$$\mathcal{L}_{\text{CLIP}} = -\frac{1}{2|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} \left[\log \frac{e^{\tau \mathbf{x}_i^\top \mathbf{y}_i}}{\sum_{j=1}^{|\mathcal{B}|} e^{\tau \mathbf{x}_i^\top \mathbf{y}_j}} + \log \frac{e^{\tau \mathbf{x}_i^\top \mathbf{y}_i}}{\sum_{j=1}^{|\mathcal{B}|} e^{\tau \mathbf{x}_j^\top \mathbf{y}_i}} \right] \quad (2)$$

where τ is a learnable temperature. This produces encoders that excel at semantic understanding and zero-shot recognition, as representations are directly grounded in natural language.

OpenAI’s CLIP ViT-L/14, with 304 million parameters trained on 400 million image-text pairs, became the standard choice for early VLMs due to its zero-shot transfer capabilities. CLIP’s contrastive objective optimizes for image-level semantics rather than fine-grained spatial understanding, a limitation that later encoders address.

The BAAI research group addressed scaling challenges with EVA-CLIP [66], which combined the EVA vision transformer architecture with CLIP-style contrastive training. Their largest model, EVA-02-CLIP-E/14+ with five billion parameters, achieved 82.0% zero-shot top-1 accuracy on ImageNet. The subsequent EVA-CLIP-18B [65] pushed the scale further to eighteen billion parameters, though empirical studies found diminishing returns in VLM contexts.

Google’s SigLIP [88] reconsidered the training loss rather than simply scaling. CLIP’s softmax-based contrastive loss creates artificial competition within each batch. SigLIP replaced this with a

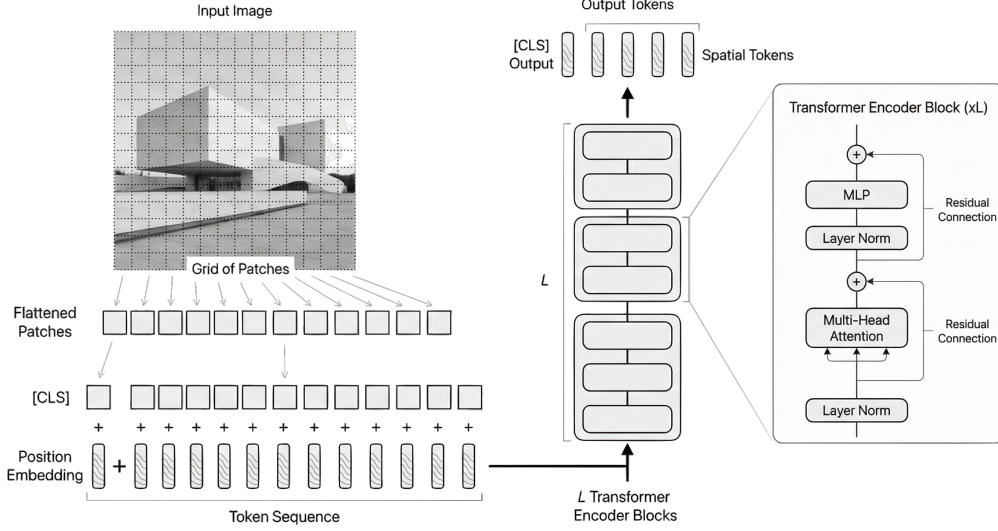


Figure 2: Vision Transformer (ViT) architecture. An input image is divided into a grid of patches, which are flattened and linearly projected into embeddings. A learnable [CLS] token is prepended, and positional embeddings are added to form the token sequence. This sequence passes through L transformer encoder blocks, each containing multi-head self-attention and MLP layers with residual connections and layer normalization. The output consists of spatial tokens plus the [CLS] token for downstream tasks.

sigmoid loss that treats each image-text pair independently:

$$\mathcal{L}_{\text{SigLIP}} = -\frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} \sum_{j=1}^{|\mathcal{B}|} \log \sigma(z_{ij}(\tau \mathbf{x}_i^\top \mathbf{y}_j + b)) \quad (3)$$

where $z_{ij} = 1$ for positive pairs and -1 otherwise. This enables stable training at larger batch sizes and simplified distributed training. The SigLIP-S0400M variant became a common choice for VLMs including LLaVA-OneVision and DeepSeek-VL2.

SigLIP 2 [70], released in February 2025, unifies multiple training objectives into a staged recipe. The first stage combines the sigmoid loss with LocCa [73], a decoder-based objective for captioning and referring expression comprehension. A lightweight transformer decoder with cross-attention to visual features is trained for three tasks: image captioning, referring expression prediction (predicting bounding boxes for region descriptions), and grounded captioning (predicting region-specific captions given coordinates). The second stage adds self-supervised losses from SILC [54] and TIPS [50]: local-to-global self-distillation where partial image views match the full-image teacher representation, and masked prediction where the student reconstructs teacher features at masked patch locations. The complete SigLIP 2 objective during the second stage is:

$$\mathcal{L}_{\text{SigLIP2}} = \mathcal{L}_{\text{sig}} + \mathcal{L}_{\text{LocCa}} + \alpha(\mathcal{L}_{\text{distill}} + \mathcal{L}_{\text{mask}}) \quad (4)$$

where α is a weighting factor that varies by model size. SigLIP 2 yields improved dense features for segmentation and depth estimation while maintaining strong zero-shot classification. It trains on multilingual data spanning over one hundred languages and has been adopted by Qwen3-VL, Gemma 3, and many 2025 frontier models.

MetaCLIP 2 [16], released in July 2025, provides the first recipe for training CLIP from scratch on worldwide web-scale image-text pairs. Through scalable substring matching and language-specific metadata curation that balances head and tail concepts across 300+ languages, MetaCLIP 2 ViT-H/14 surpasses English-only counterparts by 0.8% on ImageNet and achieves the highest reported scores on multilingual benchmarks including CVQA (57.4%), Babel-ImageNet (50.2%), and XM3600 (64.3% image-to-text retrieval), addressing the “curse of multilinguality” common in multilingual models.

TULIP [68] unifies contrastive and generative objectives, combining image-text contrastive learning with image-image contrastive learning and reconstruction regularization to address CLIP’s weakness

on vision-centric tasks requiring high-fidelity image understanding (counting, depth estimation, fine-grained recognition). TULIP scales to over 1B parameters and achieves up to $2\times$ improvement over SigLIP on RxRx1 in few-shot classification and $3\times$ higher scores on MMVP, showing that unified training objectives can better preserve visual details while maintaining semantic alignment.

Discussion. Contrastive training established vision-language alignment as the dominant paradigm, but its limitations motivated the alternatives discussed next: batch size requirements, English-centric data, and image-level rather than dense features. The evolution from CLIP to SigLIP to SigLIP 2 demonstrates that training objective improvements (sigmoid loss, multilingual data, dense supervision) yield gains that parameter scaling alone cannot match.

2.4 Self-Supervised Learning

Self-supervised methods train vision encoders without language supervision, typically using self-distillation or masked prediction. DINOv2 [55] employs a student-teacher framework where teacher parameters θ_t are an exponential moving average of student parameters θ_s :

$$\theta_t \leftarrow \lambda \theta_t + (1 - \lambda) \theta_s \quad (5)$$

The image-level self-distillation objective minimizes cross-entropy between student and teacher outputs over different augmented views:

$$\mathcal{L}_{\text{DINO}} = - \sum_{v \in \mathcal{V}_g} \mathbf{p}_t(v) \log \mathbf{p}_s(v) \quad (6)$$

where \mathcal{V}_g denotes global crop views, and $\mathbf{p}_t, \mathbf{p}_s$ are softmax-normalized prototype scores from teacher and student heads respectively. DINOv2 extends this with a patch-level masked prediction objective (iBOT), where random input patches are masked from the student but visible to the teacher:

$$\mathcal{L}_{\text{iBOT}} = - \sum_{i \in \mathcal{M}} \mathbf{p}_t^i \log \mathbf{p}_s^i \quad (7)$$

where \mathcal{M} is the set of masked patch indices. This patch-level objective encourages fine-grained local representations complementing the global image-level loss.

DINOv2 also employs the KoLeo regularizer [59], derived from the Kozachenko-Leonenko differential entropy estimator, which encourages features to span the embedding space uniformly:

$$\mathcal{L}_{\text{KoLeo}} = -\frac{1}{n} \sum_{i=1}^n \log(d_{n,i}), \quad d_{n,i} = \min_{j \neq i} \|\mathbf{x}_i - \mathbf{x}_j\| \quad (8)$$

where $d_{n,i}$ is the distance to the nearest neighbor in the batch. This regularizer prevents representation collapse and improves downstream transfer. The complete DINOv2 objective combines these terms: $\mathcal{L} = \mathcal{L}_{\text{DINO}} + \mathcal{L}_{\text{iBOT}} + \mathcal{L}_{\text{KoLeo}}$.

DINOv2 produces representations that transfer well to dense prediction tasks including segmentation, depth estimation, and fine-grained structure recognition, despite receiving no explicit supervision for these tasks. Language-supervised encoders like SigLIP excel at semantic understanding and OCR, while self-supervised encoders perform better on vision-centric tasks involving spatial relationships, object counting, and depth perception. This complementarity motivates multi-encoder approaches (Section 3.6).

Integrating self-supervised encoders into VLMs requires bridging the gap between SSL representations and language model embeddings. Three strategies exist: frozen encoder with learned projection, contrastive alignment before VLM integration, and joint fine-tuning where the SSL encoder adapts during VLM training. The Web-SSL study [23] shows that with sufficient training data diversity, frozen SSL encoders with learned projections approach contrastive encoder performance.

DINOv3 [62], released in August 2025, scaled self-supervised vision transformers to seven billion parameters through Gram anchoring, which stabilizes training at large scales by anchoring representations to a fixed reference distribution.

Discussion. Self-supervised encoders like DINOv2 excel at spatial tasks but lack text alignment, requiring adaptation for VLM integration. Their strength in dense prediction complements contrastive encoders’ semantic alignment, motivating the multi-encoder approaches discussed in Section 3.6. The gap between self-supervised and contrastive encoders on VLM tasks suggests that text alignment during pretraining remains essential for strong vision-language performance.

2.5 LLM-Aligned Training

LLM-aligned training develops vision encoders specifically for VLM integration through progressive alignment, as illustrated in Figure 3. The training proceeds in two stages: first, contrastive pretraining aligns the vision encoder with a text encoder; then, generative fine-tuning jointly optimizes the vision encoder with a language model. The generative objective trains the model to predict text tokens conditioned on visual features:

$$\mathcal{L}_{\text{gen}} = - \sum_{t=1}^T \log P(w_t | w_{<t}, \mathbf{Z}_v; \theta_v, \theta_l) \quad (9)$$

where \mathbf{Z}_v denotes visual features from the vision encoder with parameters θ_v , and θ_l are the language model parameters. This joint optimization produces encoders architecturally compatible with language models.

The InternVL project [15] developed InternViT-6B, scaling the vision encoder to six billion parameters through progressive alignment across contrastive, generative, and supervised fine-tuning stages. Section 3.4 details this approach. They also released InternViT-300M, which retained much of the capability at lower computational cost.

SAILViT [83] addresses a challenge in LLM-aligned encoders: standard ViTs trained through contrastive learning or self-supervision struggle with connector-based co-training due to parameter initialization conflicts and modality semantic gaps. SAILViT uses gradual feature refinement through three stages: coarse-grained modality alignment, fine-grained feature refinement, and world knowledge infusion. This approach improves OpenCompass benchmark performance when integrated with existing MLLMs.

Google’s NaViT [18] addressed the fixed resolution constraint of standard vision transformers. NaViT introduced “Patch n’ Pack,” which packs variable-length patch sequences from multiple images into a single training example. For images with patch counts $\{n_1, n_2, \dots, n_k\}$, the packed sequence has length:

$$L_{\text{pack}} = \sum_{i=1}^k n_i, \quad \text{where } n_i = \left\lfloor \frac{H_i}{P} \right\rfloor \cdot \left\lfloor \frac{W_i}{P} \right\rfloor \quad (10)$$

This preserves native resolution and aspect ratio while maintaining computational efficiency through sequence packing.

Qwen2-VL [74] introduced Multimodal Rotary Position Embedding (M-RoPE) to handle positional information across modalities. Standard 1D-RoPE in language models encodes only sequential position, but M-RoPE decomposes the rotary embedding into three components: temporal, height, and width.

$$\text{M-RoPE}(\mathbf{x}, t, h, w) = \mathbf{x} \odot [\cos(\theta_t), \cos(\theta_h), \cos(\theta_w)] + \mathbf{x}' \odot [\sin(\theta_t), \sin(\theta_h), \sin(\theta_w)] \quad (11)$$

where $\theta_t, \theta_h, \theta_w$ are position-dependent rotation angles for temporal, height, and width dimensions respectively, and \mathbf{x}' is a rotated version of \mathbf{x} . For text, all three components share identical position IDs, reducing to standard 1D-RoPE. For images, temporal IDs remain constant while height and width IDs reflect spatial position. For video, temporal IDs increment per frame. This decomposition enables the model to extrapolate to longer sequences during inference by keeping position ID values smaller than equivalent 1D encodings.

Table 2 summarizes the key characteristics of foundation encoders. The choice depends on the target application: SigLIP 2 offers the best general-purpose performance, DINOv2 excels when spatial understanding is paramount, InternViT provides the deepest integration with language models, and NaViT-style architectures are preferred when preserving fine image details matters most.

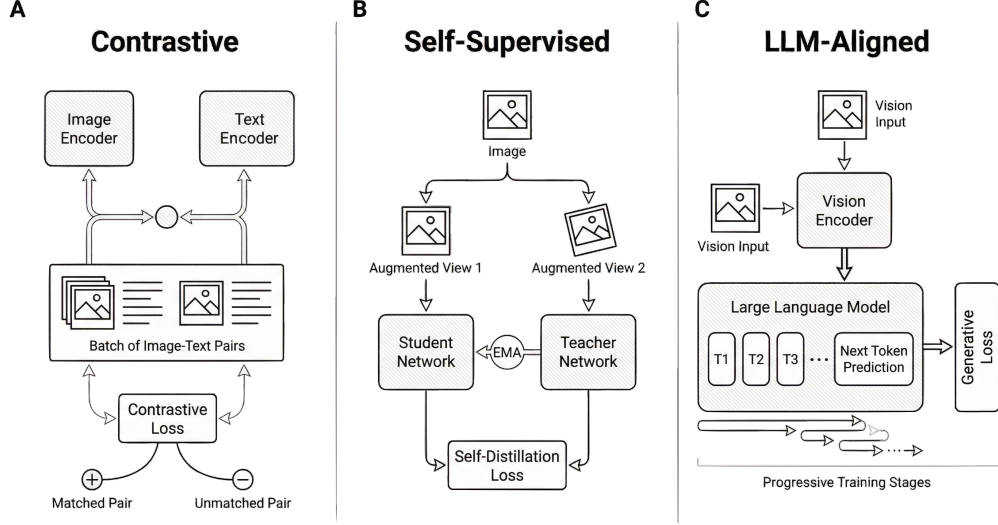


Figure 3: Vision encoder training paradigms. (A) **Contrastive**: image and text encoders are jointly trained on image-text pairs using a contrastive loss that aligns matching pairs (CLIP, SigLIP). (B) **Self-supervised**: a student network learns from a teacher network (updated via EMA) using augmented views of the same image, without text supervision (DINOv2). (C) **LLM-aligned**: the vision encoder is trained with a language model using generative objectives; progressive training may include an initial contrastive stage (InternViT).

Table 2: Foundation Vision Encoders for VLMs. Parameters indicate vision encoder size only. Resolution uses “/” for discrete checkpoint options (e.g., 224/336 = models available at either resolution) and “–” for supported ranges.

Model	Date	Org.	Params	Resolution	Key Innovation	Ref.
Contrastive (Language-Supervised)						
CLIP ViT-L/14	Jan 2021	OpenAI	304M	224/336	Contrastive pretraining	[58]
MetaCLIP	Sep 2023	Meta	304M–632M	224/336	Balanced data curation	[80]
EVA-CLIP	Mar 2023	BAAI	1B–5B	224/336	Scaled EVA + improved training	[66]
SigLIP-S0400M	Mar 2023	Google	400M	384	Sigmoid loss, batch flexibility	[88]
SigLIP 2	Feb 2025	Google	400M–1B	256/384/512	Multilingual, dense features	[70]
Self-Supervised						
DINOv2-L/g	Apr 2023	Meta	304M–1.1B	518	Self-distillation, dense features	[55]
DINOv3	Aug 2025	Meta	7B	518	Gram anchoring, stable scaling	[62]
Web-SSL	Apr 2025	Meta/NYU	up to 7B	224–518	SSL matching CLIP at scale	[23]
LLM-Aligned / Native Resolution						
InternViT-6B	Dec 2023	Shanghai AI	6B	448	LLM-aligned architecture	[15]
NaViT	Jul 2023	Google	Various	Native	Patch n’ Pack, any resolution	[18]
UniViTAR	Apr 2025	Meituan	300M–1B	Native	Resolution curriculum learning	[56]
Autoregressive						
AIMv2	Nov 2024	Apple	300M–2.7B	224/336/448	Multimodal autoregressive pretraining	[24]

Discussion. LLM-aligned encoders like InternViT represent a paradigm shift from adapting general-purpose encoders to building vision components specifically for language model integration. The progressive alignment strategy and LLM-supervised training yield strong VLM performance, but the $20\times$ parameter increase over standard encoders raises efficiency questions. Whether purpose-built encoders justify their cost depends on application requirements, as we quantify in Section 4.

Paradigm Comparison. Table 3 summarizes the three training paradigms. Contrastive training dominates general-purpose applications due to text alignment and zero-shot transfer. Self-supervised encoders excel at dense prediction but require additional alignment for VLM integration. LLM-aligned encoders achieve tight vision-language coupling at the cost of higher training complexity.

Table 3: Training Paradigm Comparison for Vision Encoders

	Contrastive	Self-Supervised	LLM-Aligned
Supervision	Image-text pairs	Images only	LLM + images
Text alignment	Native	Requires bridging	Native
Dense features	Weak (image-level)	Strong	Variable
Training cost	High (batch size)	Moderate	Very high
Scale	400M–4B	300M–7B	300M–6B
Best for	General VLM	Spatial tasks, segmentation	Deep VLM integration
Examples	CLIP, SigLIP	DINOv2, MAE	InternViT

2.6 Encoder Output Integration

How vision encoder outputs are consumed affects encoder design. The connector module transforms visual features into a format compatible with language model token embeddings, and its requirements influence encoder output dimensionality, token count, and representation structure.

The simplest approach, pioneered by LLaVA [46], uses a linear projection or shallow MLP to map visual features directly into the language model’s embedding space. Let $\mathbf{Z}_v \in \mathbb{R}^{N \times D_v}$ denote the sequence of visual token embeddings from the vision encoder, where N is the number of visual tokens and D_v is the encoder’s hidden dimension. The two-layer MLP projector computes:

$$\mathbf{H}_v = \mathbf{W}_2 \cdot \text{GELU}(\mathbf{W}_1 \cdot \mathbf{Z}_v) \quad (12)$$

where $\mathbf{W}_1 \in \mathbb{R}^{D_h \times D_v}$ and $\mathbf{W}_2 \in \mathbb{R}^{D_l \times D_h}$ project through a hidden dimension D_h to the language model dimension D_l . Liu et al. demonstrated that a single linear layer, trained on 600,000 image-text pairs, could enable visual reasoning when connecting a frozen CLIP encoder to a pretrained language model, challenging assumptions that complex cross-modal fusion mechanisms were necessary.

BLIP-2 [40] introduced the Q-Former, a more sophisticated connector that uses learnable query tokens to extract relevant information from visual features through cross-attention. Given M learnable query embeddings $\mathbf{Q} \in \mathbb{R}^{M \times D}$ and visual features \mathbf{Z}_v , the Q-Former applies cross-attention:

$$\mathbf{H}_v = \text{softmax} \left(\frac{\mathbf{Q} \mathbf{W}_Q (\mathbf{Z}_v \mathbf{W}_K)^\top}{\sqrt{D_k}} \right) \mathbf{Z}_v \mathbf{W}_V \quad (13)$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$ are learned projections. The Q-Former compresses visual information into exactly M tokens regardless of image resolution, providing control over computational cost. This compression can discard fine-grained details important for tasks like document understanding or visual grounding.

The Flamingo architecture [3] employed a Perceiver Resampler that compresses visual information through learned queries, integrating visual tokens into the language model through gated cross-attention layers interleaved with the transformer blocks. This design allows visual information to influence language model processing at multiple layers rather than only at the input. The Idefics models from HuggingFace adopted this approach for open-source VLMs.

More recent work has explored direct cross-attention mechanisms where language model tokens attend to visual features without intermediate compression. NVLM [17] conducted an extensive comparison of connector paradigms (Figure 4): decoder-only architectures concatenate visual and text tokens for unified self-attention processing; cross-attention architectures integrate visual information through dedicated attention layers where text tokens attend to visual features; and hybrid architectures combine both approaches. Their findings suggest that decoder-only approaches with careful token handling can match or exceed cross-attention designs while being simpler to implement and scale.

The Ovis architecture [49] identified a tension in existing connector designs: visual features from encoders like CLIP are continuous vectors, while language models operate on discrete token embeddings drawn from a learned vocabulary. This structural mismatch may limit how effectively visual information can be integrated. Ovis addressed this through a visual embedding table that quantizes visual features into discrete codes, structurally aligning visual and textual representations. This approach achieved strong results on vision-language benchmarks with a simpler overall architecture.

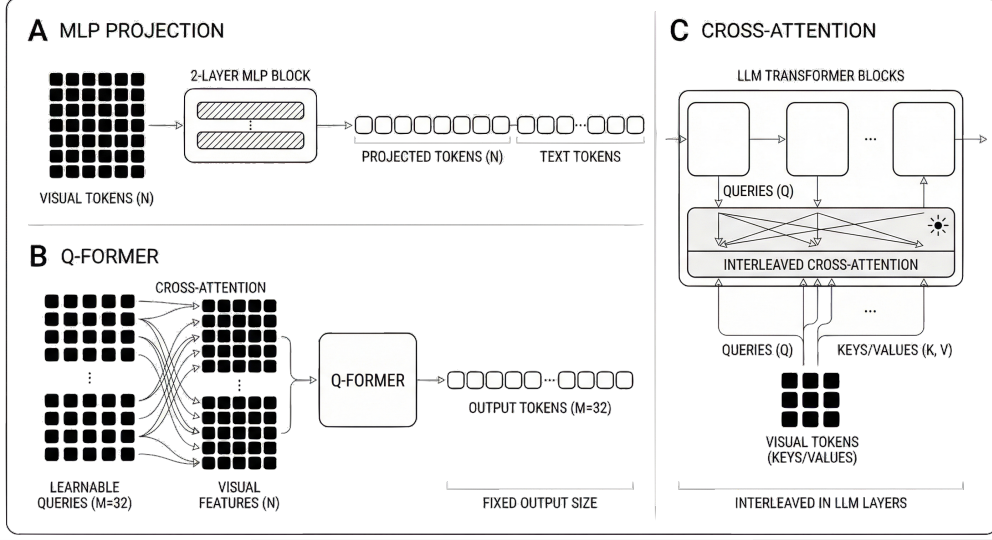


Figure 4: Vision-language connector architectures. (A) **MLP Projection:** visual tokens pass through a 2-layer MLP, preserving token count (N tokens in, N tokens out) for direct concatenation with text tokens (LLaVA). (B) **Q-Former:** learnable query tokens ($M=32$) cross-attend to visual features, compressing variable-length inputs to fixed output size (BLIP-2). (C) **Cross-Attention:** visual tokens serve as keys/values in cross-attention layers interleaved within the LLM, allowing visual information to influence processing at multiple depths (Flamingo).

Empirical studies, particularly MM1’s ablations [52], indicate that connector design has modest impact compared to encoder choice and training data quality. Simple MLP projectors match more complex designs when other factors are controlled. This finding has refocused research on encoder improvements: if connectors add little value, encoder representation quality becomes the primary determinant of VLM capability.

2.7 Resolution and Token Efficiency

Vision encoder design must balance image resolution against token count. Higher resolution preserves fine details but produces more tokens, increasing computational cost in downstream processing. This tradeoff has driven architectural innovations within encoders themselves.

Early VLMs processed all images at fixed resolutions, typically 224×224 or 336×336 pixels. Tile-based methods divide high-resolution images into grids of patches. For an image of size $H \times W$, the optimal grid configuration (n_h, n_w) minimizes aspect ratio distortion:

$$(n_h^*, n_w^*) = \arg \min_{(n_h, n_w) \in \mathcal{G}} \left| \frac{H}{W} - \frac{n_h}{n_w} \right| \quad (14)$$

where \mathcal{G} is a predefined set of valid grid configurations. LLaVA-NeXT introduced the AnyRes approach implementing this framework. The total number of visual tokens scales as $N_{\text{tokens}} = n_h \cdot n_w \cdot (P_{\text{tile}}/p)^2$ where P_{tile} is the tile size and p is the patch size.

InternVL’s dynamic tiling added Pixel Shuffle operations that reduce token counts. Given feature maps $\mathbf{Z} \in \mathbb{R}^{H' \times W' \times D}$, Pixel Shuffle with downsampling factor r rearranges spatial positions into the channel dimension:

$$\mathbf{Z}'_{i,j} = \text{concat} [\mathbf{Z}_{ri+a, rj+b}]_{a,b \in \{0, \dots, r-1\}} \in \mathbb{R}^{r^2 D} \quad (15)$$

reducing the sequence length by factor r^2 while preserving information through increased channel dimension. SmolVLM [51] uses aggressive pixel shuffle with $r = 3$, achieving $9 \times$ compression. Qwen2-VL’s “Naive Dynamic Resolution” takes an alternative approach, processing images at native resolutions into proportionally varying token counts rather than compressing post-hoc.

Beyond pixel shuffle, other compression strategies have emerged. Apple’s FastVLM [72] uses a hybrid CNN-ViT architecture (FastViTHD) that produces $3.2 \times$ fewer tokens than standard ViTs

Table 4: Resolution Strategy Comparison: Token counts and relative compute for 1MP (1024×1024) and 4MP (2048×2048) images. Assumes patch size $p = 14$ and tile size 336px. PS = Pixel Shuffle (space-to-depth); MLP = learned token merging. Compute is relative to fixed 336px baseline.

Strategy	1MP Image		4MP Image		Example
	Tokens	Compute	Tokens	Compute	
Fixed Resolution					
Fixed (224px)	256	0.4×	256	0.4×	CLIP (original)
Fixed (336px)	576	1.0×	576	1.0×	LLaVA-1.5
Fixed (384px)	729	1.3×	729	1.3×	SigLIP-SO400M
Fixed (448px)	1,024	1.8×	1,024	1.8×	InternVL 1.0
Fixed (512px)	1,296	2.3×	1,296	2.3×	SigLIP2
Fixed (518px)	1,369	2.4×	1,369	2.4×	DINOv2
Tiling (AnyRes)					
AnyRes (no compression)	5,184	9.0×	20,736	36×	LLaVA-NeXT
AnyRes + PS ($r=2$)	1,296	2.3×	5,184	9×	InternVL 2.0
AnyRes + PS ($r=3$)	576	1.0×	2,304	4×	SmolVLM
Native Resolution					
Native (no compression)	5,329	9.3×	21,316	37×	Qwen2-VL
Native + MLP (4×	1,332	2.3×	5,329	9.3×	Qwen2.5-VL
Adaptive/Hybrid					
Hybrid CNN-ViT	1,620	2.8×	6,480	11×	FastVLM
Token pruning (50%)	2,592	4.5×	10,368	18×	PTP

through efficient multi-scale feature extraction. PPE [32] enables $10\times$ compression while preserving spatiotemporal structure. PS3 [60] scales vision pre-training to 4K resolution with near-constant cost through region-selective processing, achieving $4.3\times$ fewer tokens than AnyRes while improving high-resolution perception.

PTP (Pyramid Token Pruning) [43] offers training-free adaptive compression by combining bottom-up visual saliency with top-down instruction guidance. For each visual token j , the instruction-guided importance is computed from attention scores in early LLM layers:

$$c_j = \max_{q \in \mathcal{Q}} \text{Attn}_{q \rightarrow j} \quad (16)$$

where \mathcal{Q} is the set of instruction token indices and $\text{Attn}_{q \rightarrow j}$ is the attention from instruction token q to visual token j . This captures which visual regions are most relevant to the user’s query. The bottom-up saliency b_j is derived from intermediate vision encoder layers. The final importance score combines both signals:

$$s_j = \alpha c_j + (1 - \alpha) b_j \quad (17)$$

where $\alpha \in [0, 1]$ balances instruction guidance versus visual saliency. Empirically, $\alpha = 0.5$ works well for general tasks, while OCR-heavy tasks prefer lower α (relying more on visual saliency) and open-domain reasoning benefits from higher α . PTP achieves 50% token reduction with negligible accuracy loss, and can even improve performance by filtering noisy tokens.

Table 4 compares token counts and computational costs across resolution strategies for typical image sizes. Fixed resolution maintains constant cost but loses detail; tiling preserves detail but scales quadratically with resolution; compression methods offer intermediate trade-offs.

2.8 Alternative Architectures

Beyond the standard single-encoder paradigm, two alternative approaches have emerged: multi-encoder systems that combine complementary representations, and encoder-free architectures that bypass separate vision encoders entirely.

Multi-Encoder Architectures. Different vision encoders excel at different tasks, motivating architectures that combine multiple encoders. Given K encoders producing feature maps $\{\mathbf{F}_k\}_{k=1}^K$, multi-encoder fusion computes:

$$\mathbf{F}_{\text{fused}} = \text{Aggregate}(\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_K; \mathbf{Q}) \quad (18)$$

where \mathbf{Q} are learnable query tokens. Aggregation strategies range from static fusion (Cambrian-1’s Spatial Vision Aggregator [69]) to dynamic routing (SCOPE’s Mixture-of-Encoders [91]). Section 3.6 examines specific implementations and their trade-offs.

Encoder-Free Architectures. Encoder-free architectures process raw pixels directly through the language model. Given image patches $\{\mathbf{x}_p^i\}$, these are projected directly into the language model’s embedding space:

$$\mathbf{h}_i = \mathbf{W}_{\text{patch}} \cdot \text{flatten}(\mathbf{x}_p^i) + \mathbf{e}_{\text{pos}}^i \quad (19)$$

bypassing a separate vision encoder entirely. Fuyu-8B from Adept pioneered this approach, and EVE [21] advanced it through vision-centric supervision. SAIL [37] demonstrates that scaled encoder-free models can match modular MLLMs, achieving visual representation capabilities comparable to ViT-22B. Section 3.7 examines specific implementations.

3 Encoder Adoption Across VLM Families

This section examines the major vision encoder families and how they have been adopted across VLM architectures. We organize by encoder family rather than VLM family to highlight how the same foundational encoders appear across different systems, revealing patterns in encoder selection, adaptation, and evolution.

3.1 OpenAI CLIP Family

OpenAI’s CLIP ViT encoders established the foundation for vision-language modeling, with the majority of VLMs published in 2023 using CLIP variants as documented in Table 10. Two variants dominate: CLIP ViT-L/14 with 304 million parameters and CLIP ViT-H/14 with 632 million parameters.

The LLaVA family [46] demonstrated that CLIP encoders combined with architectural simplicity could achieve strong results. The original LLaVA, released in April 2023, connected a frozen CLIP ViT-L/14 encoder to Vicuna through a single linear projection layer. LLaVA keeps the vision encoder frozen throughout training: Stage 1 freezes both the encoder and LLM while training only the projection matrix, and Stage 2 freezes the encoder while training the LLM and projection. This frozen-encoder strategy preserves pretrained visual representations while reducing training compute. LLaVA-1.5 in October 2023 maintained CLIP ViT-L/14 but processed images at 336 pixels rather than 224, with a two-layer MLP replacing the linear projection. LLaVA-NeXT [45] in January 2024 introduced AnyRes dynamic resolution processing, dividing images into grid configurations selected by aspect ratio.

Apple’s MM1 [52] in March 2024 used CLIP ViT-H/14 and conducted ablation studies showing that vision encoder quality and training data matter more than connector complexity.

Other notable CLIP-based VLMs include Phi-3.5-Vision and Phi-4-Vision from Microsoft using CLIP ViT-L/14, the Ferret series from Apple for referring and grounding tasks, Yi-VL from 01.AI using CLIP ViT-H/14, and Molmo [19] from AI2 which prioritized full openness by releasing the PixMo training dataset alongside model weights and code.

The CLIP family has several limitations that motivated the development of SigLIP: softmax-based contrastive loss requiring large batch sizes, primarily English training data, and representations optimized for image-level rather than dense features. SigLIP rapidly became the preferred alternative.

3.2 Google SigLIP Family

SigLIP emerged to address CLIP’s practical limitations: the softmax contrastive loss requires large batch sizes (32K+ for optimal performance), training data was primarily English, and representations capture image-level semantics without dense spatial features. SigLIP-S0400M, with 400 million parameters, addresses these through its sigmoid loss described in Section 2.3. Operating on individual image-text pairs rather than requiring batch comparisons, this formulation improves training stability and enables better scaling.

Google’s PaliGemma [9] series in July 2024 established SigLIP-S0400M as standard, with PaliGemma 2 scaling to 28B parameters while retaining the same encoder. Gemma 3 in March 2025 added dynamic resolution processing supporting images up to 896×896 pixels.

The LLaVA family transitioned from CLIP to SigLIP-S0400M with LLaVA-OneVision [39] in August 2024, unifying image and video understanding within a single framework. DeepSeek-VL [47] in March 2024 paired SigLIP with SAM-B for fine-grained spatial details through a three-stage training pipeline, while DeepSeek-VL2 [79] in December 2024 streamlined to SigLIP-S0400M with Mixture of Experts integration.

Other adopters include Idefics2/3 from HuggingFace, SmolVLM, MiniCPM-V 2.6/4.5, Nemotron Nano V2, Baichuan-Omni, and VideoLLaMA 3. SigLIP 2 in 2025 extended the family with multilingual capabilities and improved dense features, adopted by Qwen3-VL and Jina-VLM.

3.3 EVA-CLIP Family

EVA-CLIP explored whether larger encoders with richer pretraining could improve VLM performance. While CLIP and SigLIP remain under 700M parameters, EVA-CLIP variants range from EVA-CLIP-g with 1 billion parameters to EVA2-CLIP-E at 4.4 billion, combining masked image modeling pretraining with contrastive learning.

CogVLM [76] and CogVLM2 [29] from Tsinghua/Zhipu AI use EVA2-CLIP-E with a visual expert architecture that adds dedicated vision processing pathways within the transformer, enabling deeper vision-language fusion without modifying language model weights. CogVLM2 supports resolutions up to 1344×1344 pixels. The Emu series from BAAI uses EVA-CLIP-g for the original Emu and EVA-CLIP-E for Emu2.

The trade-off is computational cost versus capacity: EVA-CLIP-E provides richer features than 400M-class encoders, but the $10 \times$ parameter increase must be justified by application requirements.

3.4 InternViT Family

InternViT-6B, introduced in InternVL 1.0 [15] in December 2023, was designed to address the disparity between LLM scale and vision encoder scale. The encoder uses progressive alignment as described in Section 2.5, training on 4.98 billion image-text pairs.

InternVL 1.5 and 2.0 in 2024 added dynamic tiling with Pixel Shuffle token reduction, particularly for document understanding. InternVL 2.5 [14] in December 2024 expanded to 78B total parameters with choices between InternViT-300M and InternViT-6B. InternVL 3.5 in August 2025 introduced Vision Reconstruction, a self-supervised objective requiring masked image reconstruction.

NVIDIA’s NVLM [17] in September 2024 adopted InternViT-6B-448px to compare architectures. Holding the encoder constant while varying connector design across decoder-only NVLM-D, cross-attention NVLM-X, and hybrid NVLM-H variants, they found decoder-only approaches match more complex designs.

3.5 Custom and Proprietary Encoders

Several organizations have developed custom vision encoders tailored to their specific VLM architectures. These encoders often prioritize native resolution processing, efficiency, or specialized capabilities.

NaViT-Style Native Resolution Encoders. The original Qwen-VL [4] in August 2023 employed ViT-bigG with 1.9 billion parameters, initialized from OpenCLIP’s pretrained weights. Unlike frozen-encoder approaches, Qwen-VL trains the encoder at progressively higher resolutions before freezing it during instruction tuning. Qwen2-VL [74] in October 2024 introduced a 675M NaViT encoder with native resolution processing and M-RoPE positional encoding as described in Section 2.5, shown in Figure 5. Qwen2.5-VL continued this approach with refined training.

Other NaViT-style adopters include Ovis2.5 with a 300M NaViT and learnable visual embedding table, MiMo-VL from Xiaomi using a 675M NaViT based on Qwen2.5-ViT, and QVQ-72B-Preview which inherits the Qwen2-VL encoder for reasoning tasks.

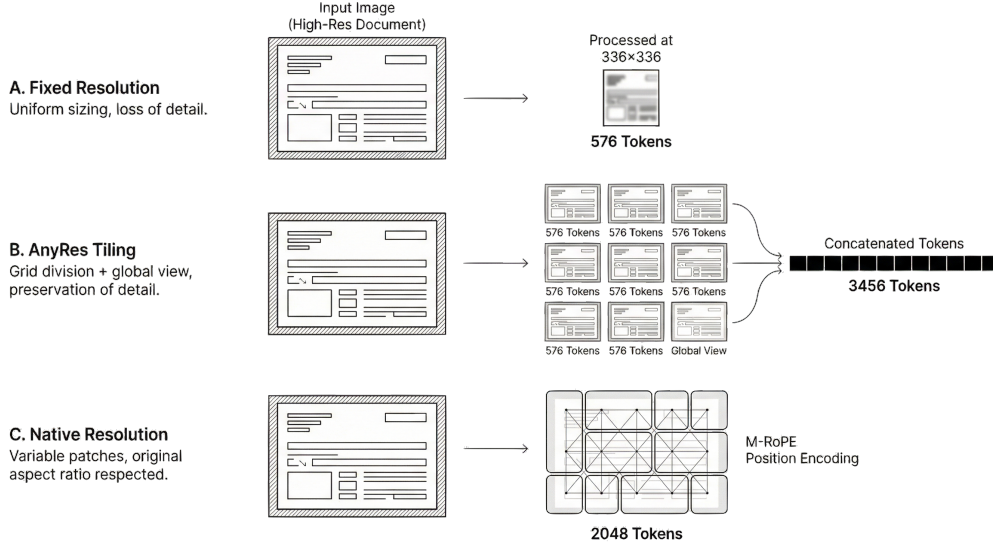


Figure 5: Resolution handling strategies for vision encoders. (A) **Fixed resolution:** images are resized to a standard size (e.g., 336×336), producing a fixed token count (576) but losing fine details. (B) **AnyRes tiling:** high-resolution images are divided into a grid of tiles plus a global view, with each tile encoded separately and concatenated (3456 tokens for a 2×3 grid). (C) **Native resolution:** images are processed at their original aspect ratio with variable patch counts, using M-RoPE position encoding to handle arbitrary dimensions (2048 tokens).

Efficiency-Focused Encoders. FastVLM [72] from Apple in December 2024 introduced FastViTHD, a 150M hybrid CNN-ViT architecture generating $3.2 \times$ fewer tokens than standard approaches while maintaining competitive performance. Kimi-VL from Moonshot AI uses MoonViT with 400 million parameters, enabling competitive reasoning performance with only 2.8B activated parameters.

Scale-Focused Encoders. Step-3 from StepFun employs EVA-CLIP 5B, one of the largest encoders in production VLMs, indicating that some applications benefit from encoder capacity beyond the typical 300M–700M range.

Cross-Attention Integration. Meta’s Llama 3.2-Vision [28] released in September 2024 uses a 632M parameter ViT-H/14 with Flamingo-style gated cross-attention layers interleaved with the language model transformer blocks. This architectural choice differs from the dominant decoder-only concatenation approach, enabling the vision encoder to remain constant while the language model scales from 11B to 90B parameters.

Custom Training. GLM-4.1V-Thinking and GLM-4.5V from Zhipu AI initialize from Apple’s AIMv2-Huge [24] and fine-tune for reasoning tasks using Reinforcement Learning with Curriculum Sampling (RLCS). MiniMax-VL-01 trains a 303M ViT encoder from scratch on 694 million image-caption pairs, demonstrating that encoder-from-scratch approaches remain viable when sufficient training data is available.

3.6 Multi-Encoder Approaches

Multi-encoder architectures combine complementary encoders through aggregation functions as formalized in Section 2.6 and illustrated in Figure 6. Cambrian-1 [69] validated the complementarity hypothesis by combining four encoders through their Spatial Vision Aggregator: CLIP ViT-L/14@336, SigLIP-S0400M/14@384, ConvNeXt-XXL@1024, and DINOv2 ViT-L/14@518. Ablation studies confirmed that no single encoder matched the combination across benchmarks.

Eagle [61] combines CLIP, ConvNeXt, Pix2Struct, and EVA-02 through Pre-Alignment training that prepares encoders for mixture before VLM training. SCOPE [91] addresses the computational




Figure 6: Multi-encoder fusion architecture (Cambrian-1). A single input image is processed by four parallel encoders: CLIP (semantic features), SigLIP (text-aligned features), DINOv2 (spatial/structural features), and ConvNeXt (multi-scale features). The Spatial Vision Aggregator (SVA) uses learnable latent tokens that cross-attend to all encoder outputs, producing unified visual tokens that capture complementary information from each encoder family before feeding into the LLM.

overhead of $4.3\times$ for Cambrian-1 through dynamic routing via Mixture of Encoder Experts, reducing cost by 24 to 49 percent while maintaining performance.

3.7 Encoder-Free Architectures

Encoder-free architectures question a fundamental assumption: is a separate pretrained vision encoder necessary, or can language models learn visual perception directly? These approaches bypass pretrained vision encoders entirely, using either direct patch projection as described in Equation 19 or discrete tokenization.

Fuyu-8B from Adept pioneered direct patch projection, processing image patches through a linear layer without vision-specific pretraining. EVE [21] and EVEv2 from BAAI advanced this approach through vision-centric supervision, demonstrating that with appropriate training, encoder-free models can match encoder-based alternatives on many benchmarks. ELVA [41] from CAS extended this work to video understanding.

Chameleon [10] and Emu3/3.5 take an alternative approach using VQ-VAE for discrete visual tokenization, treating images as token sequences analogous to text. SAIL [37] provides systematic scaling analysis, finding that encoder-free models achieve comparable performance to modular MLLMs when scaled appropriately, though with different cross-modal information flow patterns.

3.8 Document-Focused Encoders

Document understanding presents distinct challenges for vision encoding. Standard encoders trained primarily on natural images often struggle with document characteristics: dense text at varying

scales, complex layouts with tables and figures, and semantic meaning conveyed through spatial arrangements rather than visual appearance alone.

DeepSeek-OCR introduced DeepEncoder, a 380-million-parameter vision encoder specifically optimized for OCR and document layout analysis. By training on massive document datasets rather than general images, DeepEncoder develops representations attuned to the particular challenges of text recognition, including font variation, image degradation, and multilingual scripts. This specialization yields measurable improvements on document benchmarks compared to general-purpose encoders of similar size.

PaddleOCR-VL [7] from Baidu adopted NaViT-style native resolution processing specifically for multilingual document understanding. The ability to process documents at their original resolution without information-destroying downsampling proves particularly valuable for documents containing small text, fine lines in tables, or detailed figures.

DocOwl 2 [30] addressed multi-page documents through an H-Reducer architecture that compresses page-level features while preserving cross-page reasoning capability. DocVLM [53] integrates an OCR-based modality into existing VLMs, improving document understanding under tight token budgets: with 64 learned queries, DocVLM raises DocVQA accuracy from 56.0% to 86.6% when integrated with InternVL2.

Ocean-OCR [12] is a 3B-parameter MLLM employing a Native Resolution ViT for variable resolution input. It was the first MLLM to outperform professional OCR systems such as TextIn and PaddleOCR across diverse scenarios.

TextHawk2 [86] addresses OCR and visual grounding with $16\times$ token compression through a resampler architecture, achieving 78.4% accuracy on OCRBench with far fewer tokens than comparable models.

DAVE [31] (Document and web Agents Vision Encoder) introduced a purpose-built encoder for document understanding and web agent tasks through two-stage training: self-supervised pretraining with masked autoencoding followed by supervised autoregressive pretraining. On document and web benchmarks, DAVE achieves 10.5% average improvement over SigLIP 2, with particularly strong gains on web agent tasks.

4 Empirical Analysis

4.1 Encoder Adoption Trends

Vision encoder choices have shifted across the period covered by this survey. In 2023, CLIP ViT-L/14 was the predominant choice, appearing in approximately 70% of published VLMs including LLaVA, Fuyu-8B, and most models from that era. Researchers used OpenAI’s encoder for its zero-shot capabilities, availability, and community support. The encoder had become a practical default, though alternatives like EVA-CLIP and early InternVL experiments were exploring improved training procedures.

The year 2024 marked a transition as alternatives gained traction. SigLIP’s improved training objective attracted adoption from LLaVA-OneVision and the DeepSeek-VL series, while InternViT demonstrated that purpose-built encoders could outperform adapted general-purpose alternatives. Encoder choice became a deliberate design decision rather than a default.

By 2025, SigLIP 2 became the predominant choice for frontier models. Qwen3-VL, Gemma 3, and other leading VLMs adopted Google’s encoder for its multilingual capabilities, improved dense features, and strong benchmark performance. This convergence is not absolute: InternViT remains preferred for deep vision-language integration, while specialized encoders like DAVE address document-centric applications. Encoder selection increasingly depends on target use cases rather than defaults.

4.2 Connector Design Impact

As discussed in Section 2.6, connector complexity provides limited benefit compared to encoder quality. Table 8 in the Appendix confirms that most recent models adopt simple MLP projectors. Exceptions serve specific requirements: Ovis uses visual embedding tables for discrete quantization,

Table 5: Vision Encoder Cost and Performance Trade-offs. Measurements use 336×336 input resolution, batch size 1, FP16 precision, with VRAM measured on NVIDIA A100-80GB. GFLOPs are for the vision encoder only. ImageNet is zero-shot top-1; VQAv2 follows standard protocols.

Configuration	Encoder(s)	Params	GFLOPs	VRAM	Tokens	ImageNet	VQAv2
Single Encoder							
Baseline	CLIP ViT-B/16	86M	17.6	0.4GB	576	68.3	76.8
Standard	CLIP ViT-L/14	304M	81.1	1.2GB	576	75.5	79.2
Standard	SigLIP-B/16	93M	17.9	0.4GB	576	78.4	78.1
Recommended	SigLIP-S0400M	400M	95.8	1.5GB	576	83.2	81.7
Frontier	SigLIP 2 S0400M	400M	98.3	1.6GB	576	84.1	82.4
Standard	InternViT-300M	304M	82.6	1.3GB	576	79.8	80.3
Large-scale	InternViT-6B	5.9B	1,547	24GB	576	88.2	82.1
Self-supervised	DINOv2-L	304M	81.6	1.2GB	576	86.3	73.2 [†]
Self-supervised	DINOv2-g	1.1B	298	4.8GB	576	86.5	–
Multi-Encoder							
Dual	SigLIP + DINOv2-L	704M	177	2.7GB	1,152	–	–
Dual	SigLIP + SAM-B	490M	134	1.9GB	1,152	–	–
Comprehensive	Cambrian (4 enc.)	2.1B	412	5.8GB	2,048+	–	77.8
SOTA coverage	Eagle (5 enc.)	2.4B	489	6.4GB	2,500+	–	83.6

[†]DINOv2 requires text-aligned fine-tuning for VLM tasks.

DocOwl 2 uses H-Reducer for multi-page compression, and cross-attention remains useful when visual tokens must interface with language models at multiple depths.

4.3 Scaling and Token Efficiency

The relationship between encoder size and VLM performance is more nuanced than simple scaling laws might suggest. While larger encoders generally improve raw visual representation quality, these improvements do not always translate proportionally to downstream VLM performance. InternViT-6B, with twenty times the parameters of InternViT-300M, provides measurable gains on vision-centric tasks but only marginal improvements on general VLM benchmarks. Similarly, EVA-CLIP-18B shows limited practical advantage over EVA-CLIP-1B when integrated into complete VLM systems. These observations suggest that beyond a certain threshold, the language model and connector become limiting factors rather than the vision encoder’s representational capacity.

Resolution scaling follows a similar pattern of diminishing returns. The transition from 224 to 336 pixels yields improvements across nearly all tasks, while moving from 336 to 448 pixels benefits OCR and document understanding where fine text details matter. Beyond 448 pixels, the benefits become task-dependent: high-resolution processing helps for documents with small text or detailed charts but provides negligible gains for general visual question answering about natural images. Resolution requirements should be determined by target applications rather than maximized universally.

Table 5 quantifies these cost-performance trade-offs across encoder configurations, from single encoders to multi-encoder architectures. The scaling from SigLIP-B/16 to SigLIP-S0400M requires $5\times$ more computation but yields only a 4.8-point improvement on ImageNet and 3.6 points on VQAv2. InternViT-6B consumes $19\times$ the FLOPs of InternViT-300M while improving VQAv2 by less than 2 points. Multi-encoder configurations like Cambrian’s four-encoder setup incur approximately $4.3\times$ the computational cost of a single SigLIP encoder, though they achieve consistent improvements on vision-centric tasks requiring spatial understanding. For general VQA where single encoders already perform well, the multi-encoder overhead is harder to justify.

Token efficiency has emerged as equally critical to parameter efficiency in late 2024 and 2025. A 384×384 image processed at 14×14 patch size yields 729 tokens before any connector processing. High-resolution strategies compound this significantly: LLaVA-NeXT’s AnyRes can produce 2,880 tokens per image, while multi-encoder approaches generate over 2,000 tokens as shown in Table 5. For reasoning models like QVQ-72B that generate thousands of tokens during extended thinking, the visual token overhead becomes proportionally more burdensome.

Several token compression approaches address this overhead. FastV [11] prunes tokens receiving low attention scores from the language model, achieving 45% reduction with minimal accuracy loss.

VisionZip [81] uses learned selection to retain informative tokens, pushing compression ratios to 75–93%. TextHawk2 [86] achieves 16× compression through resampler architecture rather than post-hoc pruning. VTC-Bench [44] revealed that standard VLM benchmarks poorly evaluate these methods: simple downsampling often matches sophisticated compression, with compression quality mattering primarily for samples requiring fine-grained visual details, dense text, or spatial precision.

These efficiency pressures are reshaping encoder architecture. Approaches like Qwen2-VL’s NaViT naturally produce variable token counts proportional to image information content, while Pixel Shuffle operations in InternVL 2.0 reduce spatial dimensions while preserving information density. The trend toward native resolution processing partially addresses efficiency by avoiding the token overhead of tile-based approaches, though at the cost of increased architectural complexity.

4.4 Benchmarks for Vision Encoders

A distinctive challenge in vision encoder research is the absence of standardized evaluation protocols. Unlike language models, which are routinely evaluated on fixed benchmarks with reproducible metrics, vision encoders are assessed through heterogeneous methodologies that complicate cross-encoder comparison. This heterogeneity reflects tension between two evaluation philosophies: standalone evaluation that isolates encoder quality, and VLM-integrated evaluation that captures real-world utility but confounds encoder quality with connector design, LLM capability, and training procedures.

Table 6 organizes benchmarks for vision encoder evaluation by evaluation type and task category. Standalone benchmarks finetune the encoder on classic vision tasks without any language component. ImageNet classification measures category-level visual understanding and text-image alignment, with scores above 80% now expected for production encoders. COCO and ADE20K evaluate spatial localization and dense prediction respectively. For specialized applications, DocBank, DocLayNet, and RICO-SCA assess document and UI understanding. This standalone approach isolates encoder quality from downstream integration effects but may not predict VLM performance, as representations that excel at classification do not necessarily transfer optimally to language-conditioned tasks.

VLM-integrated benchmarks insert the encoder into a complete pipeline and measure downstream task performance. General visual understanding benchmarks like VQAv2, GQA, and MMBench test broad capabilities but have become relatively saturated for frontier encoders, meaning encoder improvements yield diminishing returns on these metrics. In contrast, text-heavy benchmarks including DocVQA, TextVQA, OCRBench, and ChartQA reveal encoder limitations in text recognition and document understanding, areas where encoder choice has substantial impact. Reasoning-focused benchmarks like MMMU and MathVista require combining visual perception with multi-step reasoning, testing whether encoder representations support complex inference. Spatial benchmarks including RefCOCO variants and counting tasks reveal whether encoders capture geometric structure beyond semantic content.

Research across encoder families suggests that text-image alignment quality matters most for VLM applications: encoders trained with contrastive objectives on high-quality paired data consistently outperform those with weaker supervision. Dense feature quality affects spatial tasks, with self-supervised encoders like DINOv2 excelling at segmentation despite weaker text alignment. The lack of standardization in evaluation reporting undermines reliable cross-encoder comparisons, as studies vary in LLM and connector choice, training data overlap, resolution, and whether fine-tuning was applied.

Impact of Encoder Selection. The “Encoder Sensitivity” column in Table 6 reveals that encoder choice has high impact on vision-centric tasks but diminishing returns on language-dominated benchmarks. Specifically:

- **High sensitivity:** Document understanding (DocVQA, OCRBench, ChartQA), spatial reasoning (RefCOCO, RealWorldQA, TallyQA), fine-grained recognition (BLINK), and video analysis (Video-MME). For these tasks, encoder selection and resolution handling directly determine performance ceilings.
- **Medium sensitivity:** Multi-ability benchmarks (MMBench, MME, MMMU) where both vision and language contribute. Here, encoder improvements yield measurable but not dominant gains.

Table 6: Benchmarks for Vision Encoder Evaluation. Standalone benchmarks evaluate encoder representations directly through finetuning, while VLM-integrated benchmarks measure downstream task performance within complete vision-language pipelines. Encoder sensitivity indicates how strongly benchmark performance depends on encoder choice versus other factors.

Benchmark	Task Category	Primary Measure	Encoder Sensitivity
Standalone Encoder Benchmarks			
ImageNet	Classification	Zero-shot top-1 accuracy	High
COCO	Detection/Segmentation	mAP, mask IoU	High
ADE20K	Semantic Segmentation	mIoU	High
DocBank	Document Recognition	Token-level F1	High
DocLayNet	Layout Analysis	mAP	High
RICO-SCA	UI Classification	Classification accuracy	Medium
VLM-Integrated: General Visual Understanding			
VQAv2	Visual QA	Accuracy	Low (saturated)
GQA	Compositional QA	Accuracy	Low (saturated)
MMBench	Multi-ability	Accuracy	Medium
MME	Perception + Cognition	Score (perc./cog.)	Medium
POPE	Hallucination	F1 score	Medium
MMStar	Multi-modal Reasoning	Accuracy	Medium
SEED-Bench	Generative Understanding	Accuracy	Medium
BLINK	Visual Perception	Accuracy	High
VLM-Integrated: Video Understanding			
Video-MME	Video Analysis	Accuracy	High
VLM-Integrated: Text and Document Understanding			
DocVQA	Document QA	ANLS	High
TextVQA	Scene Text QA	Accuracy	High
OCRBench	OCR Evaluation	Accuracy	High
ChartQA	Chart Understanding	Accuracy	High
InfoVQA	Infographic QA	ANLS	High
MMLongBench-Doc	Long Document QA	F1	High
VLM-Integrated: Reasoning and Knowledge			
MMMU	Multi-discipline Reasoning	Accuracy	Medium
MathVista	Mathematical Reasoning	Accuracy	Medium
MATH-V	Visual Math Problems	Accuracy	High
AI2D	Science Diagrams	Accuracy	Medium
ScienceQA	Science QA	Accuracy	Low
VLM-Integrated: Spatial and Grounding			
RefCOCO+/g	Referring Expression	Accuracy	High
RealWorldQA	Spatial Understanding	Accuracy	High
TallyQA	Object Counting	Accuracy	High

- **Low sensitivity:** General VQA (VQAv2, GQA) and knowledge-heavy tasks (ScienceQA) where benchmarks have saturated for modern encoders or language reasoning dominates. Switching from CLIP ViT-L/14 to SigLIP 2 may yield only 1–2 point improvements.

This pattern suggests practitioners should invest in encoder selection primarily when targeting document, spatial, or fine-grained visual tasks. For general-purpose VLMs, a modern baseline like SigLIP-S0400M suffices; the marginal gains from larger or specialized encoders often do not justify the computational overhead.

4.5 Bias and Safety Considerations

Vision encoders inherit biases from web-scale training data. CLIP and SigLIP models encode demographic biases, geographic imbalances, and cultural assumptions that propagate to downstream VLMs. Documented issues include lower accuracy on darker skin tones, geographic bias favoring North America and Europe, and occupational stereotyping. MetaCLIP 2 [16] addresses some issues through balanced sampling, but systematic bias auditing remains limited. Multilingual coverage is uneven, with non-Latin scripts and low-resource languages lagging behind despite claims of

multilingual support. As VLMs power agentic systems for GUI interaction and robotics, encoder failures could lead to harmful actions, and high-resolution text and face recognition raises surveillance concerns.

5 Synthesis and Future Directions

5.1 Key Findings

Three design principles emerge from this analysis, directly addressing the questions posed in Section 1:

Summary of Findings

1. **Training over Scale:** A well-trained 400M encoder outperforms a 6B encoder on most tasks.
2. **Resolution at the Encoder:** Native resolution handling yields gains that post-processing cannot recover.
3. **Complementarity over Universality:** No single encoder captures all visual features; fusion helps.

Training Methodology Dominates Scale (Questions 1–2). SigLIP 2’s sigmoid loss, multilingual data, and dense feature objectives yield gains that parameter scaling alone cannot match. A 400M-parameter encoder with superior training outperforms a 5.9B-parameter encoder on most VLM benchmarks (Table 5). This finding has practical implications: practitioners should prioritize encoder training quality over size when selecting components.

Resolution Has Become an Encoder-Level Concern (Question 3). Native resolution processing in NaViT and M-RoPE positional encoding preserve information that preprocessing discards. The shift from fixed 224px to dynamic multi-resolution handling represents a fundamental architectural change, not merely a hyperparameter adjustment. For document understanding and fine-grained recognition, resolution handling at the encoder level proves more important than downstream processing.

No Single Encoder Captures All Visual Features (Question 4). Multi-encoder fusion (Cambrian-1, Eagle) improves spatial and semantic understanding over any individual encoder, confirming that contrastive and self-supervised objectives capture complementary information. The $4\times$ computational overhead of multi-encoder approaches is justified for applications requiring both semantic understanding and spatial precision.

5.2 Future Directions

Four tensions define the field’s trajectory, each representing an unresolved design question (Question 5):

The Encoder-Free Trajectory. Encoder-free architectures (Fuyu, EVE, SAIL) demonstrate that LLMs can learn visual perception directly. As language models scale and training data grows, the value proposition of pretrained vision encoders faces fundamental questions. Will specialized encoders remain essential for efficiency, or will unified architectures subsume them? Early evidence suggests encoder-free models require substantially more training compute to match encoder-based alternatives, but this gap may narrow.

Unified Visual-Text Tokenization. Chameleon and Emu3 process images and text through a single vocabulary, eliminating the modality boundary entirely. This approach simplifies architecture but requires rethinking how visual structure is captured. The trade-offs between discrete tokenization and continuous representations remain underexplored, particularly for tasks requiring fine-grained spatial understanding.

Resolution Scaling Beyond 4K. Document understanding and video processing push resolution requirements beyond current encoder designs. Token efficiency becomes the limiting factor: a 4K image at 14px patch size yields 50,000+ tokens before any compression. Architectural innovations in selective attention, hierarchical processing, or learned compression will determine practical limits.

Encoder Specialization vs. Generalization. Purpose-built encoders like DAVE (documents) and DeepEncoder (OCR) achieve strong domain performance but sacrifice generality. The emerging question is whether the field will converge on universal encoders or fragment into task-specific variants, mirroring the broader tension between foundation models and fine-tuned specialists.

5.3 Practical Recommendations

Based on the empirical findings in Section 4, Table 7 provides decision guidance for practitioners:

Table 7: Encoder Selection Guide by Application

Application	Recommended Encoder	Rationale
General-purpose VLM	SigLIP 2 S0400M	Best training methodology, multilingual, dense features
Document understanding	NaViT-style or DeepEncoder	Native resolution preserves text details
Spatial reasoning	DINOv2+SigLIP fusion	Self-supervised features complement contrastive
Resource-constrained	SigLIP 2 Base/Large	86M–303M parameters, maintains quality
Maximum capability	Multi-encoder (Cambrian-style)	Captures complementary visual features
Research/flexibility	InternViT variants	Open weights, well-documented

For most applications, starting with SigLIP 2 variants provides a strong baseline. Add DINOv2 if spatial tasks (segmentation, depth, referring expressions) matter. Use native resolution encoders for documents. Consider encoder-free approaches only with substantial training compute budgets.

5.4 Limitations

This analysis has limitations. We cover only publicly documented models, leaving proprietary systems from Google (Gemini), OpenAI (GPT-4V), and Anthropic (Claude) opaque. Differing evaluation protocols complicate benchmark comparisons, and the rapid pace of development means some findings may not generalize to architectures published after our survey period. Our quantitative comparisons rely on reported results, which may use different evaluation setups.

5.5 Closing Perspective

Returning to our opening questions: contrastive training with modern improvements (sigmoid loss, multilingual data) yields the best general-purpose encoders (Q1); training methodology dominates parameter scale by a wide margin (Q2); native resolution handling has become essential for document and fine-grained tasks (Q3); multi-encoder fusion captures complementary features no single encoder provides (Q4); and the encoder-free trajectory remains viable but computationally expensive (Q5).

The vision encoder’s role is changing. Whether these components remain distinct or merge into natively multimodal systems, the principles identified here, namely training objective innovation, resolution flexibility, and encoder complementarity, will continue to shape how machines perceive the visual world.

References

- [1] Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024.

- [2] Praveesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, et al. Pixtral 12b. *arXiv preprint arXiv:2410.07073*, 2024.
- [3] 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. In *Advances in Neural Information Processing Systems*, volume 35, pages 23716–23736, 2022.
- [4] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.
- [5] Shuai Bai et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [6] Shuai Bai et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025.
- [7] Baidu Research. Paddleocr-vl: Boosting multilingual document parsing via a vision-language model. *arXiv preprint arXiv:2510.14528*, 2025.
- [8] Rohan Bavishi, Erich Elsen, Curtis Hawthorne, Maxwell Nye, et al. Introducing fuyu-8b: A multimodal architecture for ai agents. <https://www.adept.ai/blog/fuyu-8b>, 2023.
- [9] Lucas Beyer, Andreas Steiner, Andre Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024.
- [10] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024.
- [11] Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. In *European Conference on Computer Vision*, 2024.
- [12] Song Chen, Xinyu Guo, Yadong Li, Tao Zhang, Mingan Lin, Dongdong Kuang, Youwei Zhang, Lingfeng Ming, Fengyu Zhang, Yuran Wang, Jianhua Xu, Zenan Zhou, and Weipeng Chen. Ocean-ocr: Towards general ocr application via a vision-language model. *arXiv preprint arXiv:2501.15558*, 2025.
- [13] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025.
- [14] Zhe Chen et al. Internvl2.5: An open multimodal model outperforming gpt-4o in diverse multimodal benchmarks. *arXiv preprint arXiv:2412.05271*, 2024.
- [15] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023.
- [16] Yung-Sung Chuang, Yang Li, Dong Wang, Ching-Feng Yeh, Kehan Lyu, Ramya Raghavendra, James Glass, Lifei Huang, Jason Weston, Luke Zettlemoyer, Xinlei Chen, Zhuang Liu, Saining Xie, Wen-tau Yih, Shang-Wen Li, and Hu Xu. Meta clip 2: A worldwide scaling recipe. *arXiv preprint arXiv:2507.22062*, 2025.
- [17] Wenliang Dai, Nayeon Lee, Boxin Wang, Zhuolin Yang, Zihan Liu, Jon Barker, Tuomas Rintamaki, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Nvlm: Open frontier-class multimodal llms. *arXiv preprint arXiv:2409.11402*, 2024.
- [18] Mostafa Dehghani, Basil Mustafa, Josip Djolonga, Jonathan Heek, Matthias Minderer, Mathilde Caron, Andreas Steiner, Joan Puigcerver, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Patch n’ pack: Navit, a vision transformer for any aspect ratio and resolution. *Advances in Neural Information Processing Systems*, 36, 2023.

- [19] Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv preprint arXiv:2409.17146*, 2024.
- [20] Ameya Sunil Deshmukh et al. Nvidia nemotron nano v2 vl. *arXiv preprint arXiv:2511.03929*, 2025.
- [21] Haiwen Diao, Yufeng Cui, Xiaotong Li, Yueze Wang, Huchuan Lu, and Xinlong Wang. Unveiling encoder-free vision-language models. *arXiv preprint arXiv:2406.11832*, 2024.
- [22] Haiwen Diao, Yufeng Cui, Xiaotong Wang, et al. Evev2: Improved baselines for encoder-free vision-language models. *arXiv preprint arXiv:2502.06788*, 2025.
- [23] 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. *arXiv preprint arXiv:2504.01017*, 2025.
- [24] Enrico Fini, Mustafa Shukor, Xiujun Li, Philipp Dufter, Michal Klein, David Halber, Vasu Sharma Ramber, Shashank Bhardwaj, Nikhil Jain, Hiroaki Yakura, et al. Multimodal autoregressive pre-training of large vision encoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- [25] Gemma Team. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.
- [26] GLM Team. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*, 2025.
- [27] Aaron Grattafiori, Abhimanyu Dubey, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [28] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [29] Wenyi Hong, Weihang Wang, Ming Ding, Wenmeng Yu, Qingsong Lv, Yan Wang, Yean Cheng, Shiyu Huang, Junhui Ji, Zhao Xue, et al. Cogvlm2: Visual language models for image and video understanding. *arXiv preprint arXiv:2408.16500*, 2024.
- [30] Anwen Hu, Haiyang Xu, Liang Zhang, Jiabo Ye, Ming Yan, Ji Zhang, Qin Jin, Fei Huang, and Jingren Zhou. mplug-docowl2: High-resolution compressing for ocr-free multi-page document understanding. *arXiv preprint arXiv:2409.03420*, 2024.
- [31] Siyuan Huang, Siyu Dong, Jinpeng Liu, Junyan Chen, Yuying Fan, Xiong Zeng, et al. Dave: A document-centric approach for vision encoders. *arXiv preprint arXiv:2512.17221*, 2025.
- [32] Yiming Huang, Zhen Li, Kaipeng Zhang, et al. Ppe: Positional preservation embedding for visual token compression. *arXiv preprint arXiv:2510.22936*, 2025.
- [33] Kimi Team. Kimi-vl technical report. *arXiv preprint arXiv:2504.07491*, 2025.
- [34] Alexis Koukounas et al. Jina-vlm: A small multilingual vision language model. *arXiv preprint arXiv:2512.04032*, 2025.
- [35] Hugo Laurençon, Andrés Marafioti, Victor Sanh, et al. Building and better understanding vision-language models: Insights and future directions. *arXiv preprint arXiv:2408.12637*, 2024.
- [36] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander Rush, Douwe Kiela, et al. What matters when building vision-language models? *arXiv preprint arXiv:2405.02246*, 2024.
- [37] Weixian Lei, Jiacong Wang, Haochen Wang, Xiangtai Li, Jun Hao Liew, Jiashi Feng, and Zilong Huang. The scalability of simplicity: Empirical analysis of vision-language learning with a single transformer. *arXiv preprint arXiv:2504.10462*, 2025.

- [38] Aoran Li et al. Minimax-01: Scaling foundation models with lightning attention. *arXiv preprint arXiv:2501.08313*, 2025.
- [39] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.
- [40] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [41] Shengqiong Li et al. Breaking the encoder barrier for seamless video-language understanding. *arXiv preprint arXiv:2503.18422*, 2025.
- [42] Yadong Li, Haoze Zhang, Mingan Ren, Wentao Yang, Zenan Zhang, Jianhua Xu, and Weipeng Chen. Baichuan-omni technical report. *arXiv preprint arXiv:2410.08565*, 2024.
- [43] Yuxuan Liang, Xu Li, Xiaolei Chen, Yi Zheng, Haotian Chen, Bin Li, and Xiangyang Xue. Pyramid token pruning for high-resolution large vision-language models via region, token, and instruction-guided importance. *arXiv preprint arXiv:2509.15704*, 2025.
- [44] Chenfei Liao, Wensong Wang, Zichen Wen, Xu Zheng, Yiyu Wang, Haocong He, Yuanhuiyi Lyu, Luta Jiang, Xin Zou, Yuqian Fu, Bin Ren, Linfeng Zhang, and Xuming Hu. Are we using the right benchmark: An evaluation framework for visual token compression methods. *arXiv preprint arXiv:2510.07143*, 2025.
- [45] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge. *Blog post*, 2024.
- [46] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- [47] Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yang Sun, et al. Deepseek-vl: Towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*, 2024.
- [48] Shiyin Lu et al. Ovis2.5 technical report. *arXiv preprint arXiv:2508.11737*, 2025.
- [49] Shiyin Lu, Yang Yang, Qing Li, et al. Ovis: Structural embedding alignment for multimodal large language model. *arXiv preprint arXiv:2407.04860*, 2024.
- [50] Kevis-Kokitsi Maninis, Kaifeng Chen, Soham Ghosh, Arjun Karapur, Koert Chen, Ye Xia, Bingyi Cao, Daniel Salz, Guangxing Han, Jan Dlabal, Danushen Gnanapragasam, Mojtaba Seyedhosseini, Howard Zhou, and Andre Araujo. Tips: Text-image pretraining with spatial awareness. In *International Conference on Learning Representations*, 2025.
- [51] Andrés Marafioti, Lilli Bauckholt, et al. Smolvlm: Redefining small and efficient multimodal models. *arXiv preprint arXiv:2504.05299*, 2025.
- [52] Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruti Shah, Xianzhi Du, Futang Peng, Floris Weers, et al. Mm1: Methods, analysis & insights from multimodal llm pre-training. *arXiv preprint arXiv:2403.09611*, 2024.
- [53] Mor Shpigel Nacson, Aviad Aberdam, Roy Ganz, Elad Ben-Avraham, Alona Golts, Yair Kittenplon, Shai Mazor, and Ron Litman. Docvlm: Make your vlm an efficient reader. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 29005–29015, 2025.
- [54] 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*, pages 38–55, 2024.
- [55] 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.

- [56] Limeng Qiao, Yiyang Gan, Bairui Wang, Jie Qin, Shuang Xu, Siqi Yang, and Lin Ma. Univitar: Unified vision transformer with native resolution. *arXiv preprint arXiv:2504.01792*, 2025.
- [57] Yanyuan Qiao, Zheng Yu, Longteng Guo, Sihan Chen, Zijia Zhao, Mingzhen Sun, Qi Wu, and Jing Liu. Vl-mamba: Exploring state space models for multimodal learning. *arXiv preprint arXiv:2403.13600*, 2024.
- [58] 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*, pages 8748–8763. PMLR, 2021.
- [59] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. Spreading vectors for similarity search. In *International Conference on Learning Representations*, 2019.
- [60] Baifeng Shi, Boyi Li, Han Cai, Yao Lu, Sifei Liu, Marco Pavone, Jan Kautz, Song Han, Trevor Darrell, Pavlo Molchanov, and Hongxu Yin. Scaling vision pre-training to 4k resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- [61] Min Shi, Fuxiao Xie, Ziwei Wang, Le Liu, Guangyu Chen, Xingcheng Wen, et al. Eagle: Exploring the design space for multimodal llms with mixture of encoders. *arXiv preprint arXiv:2408.15998*, 2024.
- [62] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Timothée Darcet, Daniel Haziza, Francisco Massa, Théo Moutakanni, Marc Szafraniec, Quentin Garrido, Shijie Tang, et al. Dinov3: Towards visual foundation models with self-supervised learning. *arXiv preprint arXiv:2508.10104*, 2025.
- [63] StepFun Team. Step-3 is large yet affordable: Model-system co-design for cost-effective decoding. *arXiv preprint arXiv:2507.19427*, 2025.
- [64] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiying Yu, Zhengxiong Luo, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. *arXiv preprint arXiv:2312.13286*, 2023.
- [65] Quan Sun et al. Eva-clip-18b: Scaling clip to 18 billion parameters. *arXiv preprint arXiv:2402.04252*, 2024.
- [66] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.
- [67] Quan Sun, Qiying Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Emu: Generative pretraining in multimodality. *arXiv preprint arXiv:2307.05222*, 2023.
- [68] Zineng Tang, Long Lian, Seun Eisape, Jing Shi, Trevor Darrell, and Jacob Andreas. Tulip: Towards unified language-image pretraining. *arXiv preprint arXiv:2503.15485*, 2025.
- [69] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Juhan 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.
- [70] Michael Tschannen et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.
- [71] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning, 2017.
- [72] Pavan Kumar Anasosalu Vasu, Fartash Kanber, Hadi Dehghani, Farzaneh Faghri, Peter Horton, Shlok Saxena, Jean-Baptiste Cordonnier, and Seyed Mehdi Kazemi. Fastvlm: Efficient vision encoding for vision language models. *arXiv preprint arXiv:2412.13303*, 2024.

- [73] Bo Wan, Michael Tschannen, Yongqin Xian, Filip Pavetic, Ibrahim Alabdulmohsin, Xiao Wang, André Susano Pinto, Andreas Steiner, Lucas Beyer, and Xiaohua Zhai. Locca: Visual pretraining with location-aware captioners. In *Advances in Neural Information Processing Systems*, 2024.
- [74] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.
- [75] Weihan Wang et al. Internvl3.5: Advancing open-source multimodal models with scalable reinforcement learning and semantic routing. *arXiv preprint arXiv:2508.18265*, 2025.
- [76] Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023.
- [77] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024.
- [78] Hao Wei et al. Deepseek-ocr: Contexts optical compression. *arXiv preprint arXiv:2510.18234*, 2025.
- [79] Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, et al. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*, 2024.
- [80] Hu Xu, Saining Xie, Xiaoqing Ellen Tan, Po-Yao Huang, Russell Howes, Vasu Sharma, Shang-Wen Li, Gargi Ghosh, Luke Zettlemoyer, and Christoph Feichtenhofer. Demystifying clip data. *arXiv preprint arXiv:2309.16671*, 2023.
- [81] Senqiao Yang, Yukang Chen, Zhuotao Tian, Chengyao Wang, Jingyao Li, Bei Yu, and Jiaya Jia. Visionzip: Longer is better but not necessary in vision language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19792–19802, 2025.
- [82] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024.
- [83] Weijie Yin, Dingkang Yang, Hongyuan Dong, Zijian Kang, Jiacong Wang, Xiao Liang, Chao Feng, and Jiao Ran. Sailvit: Towards robust and generalizable visual backbones for mllms via gradual feature refinement. *arXiv preprint arXiv:2507.01643*, 2025.
- [84] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023.
- [85] Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. Yi: Open foundation models by 01.ai. *arXiv preprint arXiv:2403.04652*, 2024.
- [86] Ya-Qi Yu, Minghui Liao, Jiwen Zhang, and Jihao Wu. Texthawk2: A large vision-language model excels in bilingual ocr and grounding with 16x fewer tokens. *arXiv preprint arXiv:2410.05261*, 2024.
- [87] Zhiwei Yue et al. Mimo-vl technical report. *arXiv preprint arXiv:2506.03569*, 2025.
- [88] 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*, pages 11975–11986, 2023.

- [89] Boqiang Zhang, Zesen Cheng, Sicong Leng, Guanying Zhang, Junhao Luo, Xu Zhang, Haotian Luo, Jingdong Wang, Xiawu Nie, Yifei Xin, et al. Videollama 3: Frontier multimodal foundation models for image and video understanding. *arXiv preprint arXiv:2501.13106*, 2025.
- [90] Haotian Zhang, Haoxuan You, Philipp Dufter, Bowen Zhang, Chen Xu, Zhe Hong, Jianfeng Yang, Zhe Gan, Shih-Fu Chang, and Yinfei Yang. Ferret-v2: An improved baseline for referring and grounding with large language models. *arXiv preprint arXiv:2404.07973*, 2024.
- [91] Tianyu Zhang, Suyuchen Wang, Chao Wang, Juan A. Rodriguez, Ahmed Masry, Xiangru Jian, Yoshua Bengio, and Perouz Taslakian. Scope: Selective cross-modal orchestration of visual perception experts. *arXiv preprint arXiv:2510.12974*, 2025.
- [92] Han Zhao, Min Zhang, Wei Zhao, Pengxiang Ding, Siteng Huang, and Donglin Wang. Cobra: Extending mamba to multi-modal large language model for efficient inference. *arXiv preprint arXiv:2403.14520*, 2024.

A Comprehensive Model Tables

This appendix provides exhaustive tables of vision encoders in VLMs, organized by multiple dimensions to facilitate comparison and analysis.

A.1 Models by Connector Architecture

Table 8 categorizes VLMs by their vision-language connector design.

Table 8: VLMs Organized by Connector Architecture

Model	Connector Type	Description	Date
Linear / MLP Projection			
LLaVA	Linear	Single linear layer	Apr 2023
LLaVA-1.5	2-layer MLP	GELU activation	Oct 2023
LLaVA-NeXT	2-layer MLP	+ AnyRes tiling	Jan 2024
LLaVA-OneVision	2-layer MLP	Unified image/video	Aug 2024
PaliGemma/2	Linear	Direct projection	Jul 2024
Gemma 3	Soft-token MLP	Token-level projection	Mar 2025
InternVL 2.0+	Pixel Shuffle + MLP	Token reduction	Jul 2024
DeepSeek-VL2	MLP	+ Dynamic tiling	Dec 2024
NVLM-D	2-layer MLP	Decoder-only arch.	Sep 2024
Molmo	MLP	Simple projection	Sep 2024
SmolVLM	Efficient MLP	Optimized for size	Apr 2025
Q-Former / Learnable Queries			
BLIP-2	Q-Former	32 learnable queries	Jan 2023
InstructBLIP	Q-Former	Instruction-aware	Jun 2023
InternVL 1.0	QLLaMA	Query-based extraction	Dec 2023
MiniGPT-4	Q-Former	Single projection layer	Apr 2023
Cross-Attention / Perceiver			
Flamingo	Perceiver Resampler	Gated cross-attention	Apr 2022
Idefics	Perceiver Resampler	Open Flamingo-style	Aug 2023
Idefics2	Perceiver + MLP	Refined resampler	Apr 2024
Qwen-VL	Cross-attention	Single-layer xattn	Aug 2023
NVLM-X	Cross-attention	Gated xattn layers	Sep 2024
NVLM-H	Hybrid	MLP + xattn combined	Sep 2024
Llama 3.2-Vision	Cross-attn adapters	Interleaved layers	Sep 2024
Visual Embedding Table			
Ovis	Embedding table	Discrete quantization	Jul 2024
Ovis 1.6	Embedding table	+ AnyRes	Sep 2024
Ovis2/2.5	Embedding table	Improved quantization	Aug 2025
Specialized Connectors			
Cambrian-1	SVA	Spatial Vision Aggregator	Jun 2024
Eagle	Pre-Alignment	Encoder alignment stage	Aug 2024
DocOwl 2	H-Reducer	Hierarchical reduction	May 2024
MiniCPM-V	Compression	Adaptive compression	Aug 2024
MM1	C-Abstractor	Convolutional abstractor	Mar 2024
Ferret-v2	DPE	Dense Position Encoding	Apr 2024
No Connector (Native)			
Fuyu-8B	Direct input	Patches as tokens	Oct 2023
EVE/EVEv2	Direct input	Vision-centric training	Jun 2024
Chameleon	VQ tokens	Discrete visual tokens	May 2024

Continued on next page

Table 8 – continued

Model	Connector Type	Description	Date
Emu3	Native	Unified tokenization	Oct 2025

A.2 Models by Resolution Strategy

Table 9 organizes VLMs by their approach to handling image resolution.

Table 9: VLMs Organized by Resolution Strategy

Model	Strategy	Max Resolution	Notes
Fixed Resolution			
LLaVA	Fixed 224px	224×224	Single resolution
LLaVA-1.5	Fixed 336px	336×336	Higher res baseline
BLIP-2	Fixed 224px	224×224	Q-Former compression
Flamingo	Fixed 224px	224×224	Perceiver handles var.
MiniGPT-4	Fixed 224px	224×224	BLIP-2 backbone
Dynamic Tiling (AnyRes-style)			
LLaVA-NeXT	AnyRes	672–1344px	Grid selection
LLaVA-OneVision	AnyRes	672–1344px	Unified img/video
InternVL 1.5+	Dynamic tiling	448–4096px	Pixel Shuffle reduction
Qwen2-VL	Naive Dynamic	Variable	Proportional tokens
DeepSeek-VL2	Dynamic tiling	1024px+	MoE integration
Phi-3.5-Vision	Dynamic	1344px	Aspect-aware
MiniCPM-V 2.6	Adaptive	1344px	Compression
NVLM-D/X/H	Dynamic	448–1792px	Multi-tile
MM1.5	Dynamic	1344px	Improved AnyRes
Native Resolution (NaViT-style)			
Ovis2.5	Native	Variable	No resize
PaddleOCR-VL	Native	Variable	Document focus
Fuyu-8B	Native	Variable	Direct patches
High Fixed Resolution			
DeepSeek-VL	Fixed 1024px	1024×1024	Hybrid encoder
Qwen-VL	Fixed 448px	448×448	Large ViT
InternVL 1.0	Fixed 448px	448×448	InternViT-6B
CogVLM2	Fixed 490px	490×490	EVA2–CLIP–E
Multi-Scale / Pyramid			
Ferret-v2	Multi-scale	Multiple	DPE connector
Cambrian-1	Multi-encoder	384–512px	4 encoders
Eagle	Multi-scale	Variable	Pre-alignment

A.3 Complete Chronological Model Database

Table 10 provides a complete chronological listing of all VLMs with their full specifications.

Table 10: Complete Chronological VLM Database (2023–2025). *Encoder trained from scratch.

Model	Date	Org.	Encoder	Connector	Params
2023					

Continued on next page

Table 10 – continued

Model	Date	Org.	Encoder	Connector	Params
BLIP-2	Jan	Salesforce	EVA-CLIP-g	Q-Former	3B–12B
MiniGPT-4	Apr	KAUST	EVA-CLIP-g	Q-Former	13B
LLaVA	Apr	UW-Madison	CLIP ViT-L/14	Linear	7B–13B
InstructBLIP	Jun	Salesforce	EVA-CLIP-g	Q-Former	7B–13B
Emu	Jul	BAAI	EVA-CLIP-g	Causal	14B
Qwen-VL	Aug	Alibaba	ViT-bigG	Cross-attn	9.6B
Idefics	Aug	HuggingFace	OpenCLIP	Perceiver	9B–80B
Fuyu-8B	Oct	Adept	None	Direct	8B
LLaVA-1.5	Oct	UW-Madison	CLIP ViT-L/14	2-layer MLP	7B–13B
Ferret	Oct	Apple	CLIP ViT-L/14	Spatial	7B–13B
CogVLM	Nov	Tsinghua	EVA2-CLIP-E	Visual expert	17B
Emu2	Dec	BAAI	EVA-CLIP-E	Causal	37B
InternVL 1.0	Dec	Shanghai AI	InternViT-6B	QLLaMA	26B
2024 Q1					
LLaVA-NeXT	Jan	ByteDance	CLIP ViT-L/14	MLP+AnyRes	7B–110B
TinyLLaVA	Feb	–	CLIP/SigLIP	MLP	1.4B–3.1B
Yi-VL	Mar	01.AI	CLIP ViT-H/14	MLP	6B–34B
DeepSeek-VL	Mar	DeepSeek	SigLIP+SAM-B	Hybrid	1.3B–7B
MM1	Mar	Apple	CLIP ViT-H/14	C-Abstractor	3B–64B
VL-Mamba	Mar	–	CLIP ViT-L/14	Mamba	7B
Cobra	Mar	–	CLIP+DINOv2	Mamba	7B
2024 Q2					
Phi-3.5-Vision	Apr	Microsoft	CLIP ViT-L/14	MLP	4.2B
InternVL 1.5	Apr	Shanghai AI	InternViT-6B	PixelShuffle	2B–26B
Idefics2	Apr	HuggingFace	SigLIP-S0400M	Perceiver+MLP	8B
Ferret-v2	Apr	Apple	CLIP (multi)	DPE	7B–13B
Chameleon	May	Meta	VQ-VAE	Native	7B–34B
CogVLM2	May	Tsinghua	EVA2-CLIP-E	Visual expert	19B
DocOwl 2	May	Alibaba	ViT	H-Reducer	8B
Cambrian-1	Jun	NYU	4 encoders	SVA	8B–34B
EVE	Jun	BAAI	None	Direct	7B
2024 Q3					
InternVL 2.0	Jul	Shanghai AI	InternViT	PixelShuffle	1B–76B
Ovis	Jul	Alibaba	SigLIP-S0400M	Embed. table	9B
PaliGemma	Jul	Google	SigLIP-S0400M	Linear	3B
Idefics3	Aug	HuggingFace	SigLIP	PixelShuffle	8B
LLaVA-OneVision	Aug	ByteDance	SigLIP-S0400M	MLP	0.5B–72B
Eagle	Aug	NVIDIA	Multi-encoder	Pre-Align	7B–13B
MiniCPM-V 2.6	Aug	Tsinghua	SigLIP	Compression	8B
MM1.5	Sep	Apple	CLIP variants	C-Abstractor	1B–30B
NVLM-D/X/H	Sep	NVIDIA	InternViT-6B	MLP/xattn	72B
Llama 3.2-Vision	Sep	Meta	ViT-H/14	xattn adapter	11B–90B
Ovis 1.6	Sep	Alibaba	SigLIP-S0400M	Embed. table	9B
Molmo	Sep	AI2	CLIP ViT-L/14	MLP	7B–72B
Emu3	Sep	BAAI	VQ-VAE	Native	8B
2024 Q4					
Qwen2-VL	Oct	Alibaba	NaViT	M-RoPE	2B–72B

Continued on next page

Table 10 – continued

Model	Date	Org.	Encoder	Connector	Params
Baichuan-Omni	Oct	Baichuan	SigLIP	MLP	7B
Pixtral 12B	Oct	Mistral	ViT*	Native	12B
FastVLM	Dec	Apple	FastViTHD	Efficient	0.5B–7B
InternVL 2.5	Dec	Shanghai AI	InternViT	PixelShuffle	1B–78B
PaliGemma 2	Dec	Google	SigLIP-S0400M	Linear	3B–28B
DeepSeek-VL2	Dec	DeepSeek	SigLIP-S0400M	MLP+MoE	3B–27B
Ovis2	Dec	Alibaba	SigLIP	Embed. table	8B
Phi-4-Vision	Dec	Microsoft	CLIP+	MLP	14B
2025					
Janus-Pro	Jan	DeepSeek	SigLIP-L	Decoupled	1.5B–7B
EVEv2	Feb	BAAI	None	Direct	7B
Gemma 3	Mar	Google	SigLIP-S0400M	Soft-token	4B–27B
ELVA	Mar	CAS	None	Direct	7B
Llama 4	Apr	Meta	Proprietary	Early fusion	109B+
SmolVLM	Apr	HuggingFace	SigLIP-S0400M	Efficient	256M–2B
Kimi-VL	Apr	Moonshot AI	MoonViT	MLP	16B
MiMo-VL	Jun	Xiaomi	NaViT	MLP	7B
GLM-4.1V	Jul	Zhipu AI	AIMv2-Huge	MLP	9B
InternVL 3.5	Aug	Shanghai AI	InternViT+ViR	PixelShuffle	1B–78B
Ovis2.5	Aug	Alibaba	NaViT	Embed. table	2B–9B
DeepSeek-OCR	Oct	DeepSeek	DeepEncoder	MLP	7B
PaddleOCR-VL	Oct	Baidu	NaViT-style	MLP	0.9B
Emu3.5	Oct	BAAI	Native	Unified	–
Qwen3-VL	Nov	Alibaba	SigLIP 2 S0400M	M-RoPE	2B–72B
Nemotron Nano	Nov	NVIDIA	SigLIP	MLP	4B
Jina-VLM	Dec	Jina AI	SigLIP 2 S0400M	MLP	1.5B

A.4 Video-Capable VLMs

Table 11 lists VLMs with explicit video understanding capabilities and their temporal processing strategies.

Table 11: Video-Capable VLMs and Temporal Processing

Model	Date	Vision Encoder	Temporal Strategy	Max Frames	Params
Video-ChatGPT	Jun 2023	CLIP ViT-L/14	Spatial-temporal pooling	100	7B
VideoLLaVA	Nov 2023	LanguageBind	Joint image-video	8	7B
LLaVA-NeXT-Video	Apr 2024	CLIP/SigLIP	AnyRes + temporal	32	7B–34B
LLaVA-OneVision	Aug 2024	SigLIP-S0400M	Unified frames	32	0.5B–72B
Qwen2-VL	Oct 2024	NaViT + M-RoPE	3D position encoding	768	2B–72B
InternVL 2.0	Jul 2024	InternViT	Frame sampling	64	8B–76B
LongVA	Jun 2024	SigLIP	Long context	2000	7B
Oryx	Sep 2024	OryxViT	Dynamic resolution	64	7B–34B
ELVA	Mar 2025	Encoder-free	Direct video	64	7B
VideoLLaMA 2	Jun 2024	CLIP ViT-L	Spatial-temporal	16	7B–72B
VideoLLaMA 3	Jan 2025	SigLIP/DFN	Temporal pooling	64	2B–72B

A.5 Benchmark Performance Summary

Table 12 provides representative benchmark scores for major VLMs across key evaluation dimensions.

Table 12: Representative Benchmark Performance (2024–2025 Models)

Model	MMBench	MMMU	MathVista	DocVQA	OCRBench	RealWorldQA
Proprietary Models (Reference)						
GPT-4o (Nov 2024)	83.4	69.1	63.8	92.8	736	75.4
Claude 3.5 Sonnet	78.5	68.3	61.6	95.2	788	60.1
Gemini 1.5 Pro	73.9	62.2	57.7	93.1	754	67.5
Open Models ($\leq 10\text{B}$)						
Qwen2-VL-7B	83.0	54.1	58.2	94.5	845	70.1
InternVL2-8B	81.7	51.8	58.3	91.6	794	64.4
LLaVA-OneVision-7B	80.8	48.8	63.2	87.5	622	60.2
MiniCPM-V 2.6	78.0	49.8	60.6	90.8	852	65.2
Idefics3-8B	77.4	47.6	54.3	87.1	710	60.5
Open Models ($> 10\text{B}$)						
Qwen2-VL-72B	86.5	64.5	70.5	96.5	855	77.8
InternVL2.5-78B	85.8	68.0	72.3	95.7	857	74.1
Llama 3.2-90B	80.5	60.3	57.3	90.1	763	68.2
NVLM-D-72B	82.6	59.7	65.2	92.6	785	69.5