





Turbo: Informativity-Driven Acceleration Plug-In for Vision-Language Large Models

Chen Ju^{1*}, Haicheng Wang^{1,2*}, Haozhe Cheng², Xu Chen¹, Zhonghua Zhai¹, Weilin Huang¹, Jinsong Lan¹, Shuai Xiao¹✉, and Bo Zheng¹

¹Alibaba Group, ²Shanghai Jiao Tong University
cju.void@gmail.com, anakin_skywalker@sjtu.edu.cn,
shuai.xsh@alibaba-inc.com
<https://voide1220.github.io/turbo/>

Abstract. Vision-Language Large Models (VLMs) recently become primary backbone of AI, due to the impressive performance. However, their expensive computation costs, *i.e.*, throughput and delay, impede potentials in the real-world scenarios. To achieve acceleration for VLMs, most existing methods focus on the model perspective: pruning, distillation, quantization, but completely overlook the data-perspective redundancy. To fill the overlook, this paper pioneers the severity of data redundancy, and designs one plug-and-play Turbo module guided by information degree to prune inefficient tokens from visual or textual data. In pursuit of efficiency-performance trade-offs, information degree takes two crucial factors into consideration: mutual redundancy and semantic value. Concretely, the former evaluates data duplication between sequential tokens; while the latter evaluates each token by its contribution to the overall semantics. As a result, tokens with high information degree carry less redundancy and stronger semantics. For VLMs’ calculation, Turbo works as a user-friendly plug-in that sorts data referring to information degree, utilizing only top-level ones to save costs. Its advantages are multifaceted, *e.g.*, being generally compatible to various VLMs across understanding and generation, simple use without re-training and trivial engineering efforts. On multiple VLMs benchmarks, we fully experiment to demonstrate the good acceleration of Turbo, under negligible performance drop.

1 Introduction

Vision-Language Large Models (VLMs) [20, 29–31, 45] have achieved promising progress towards AI. Inspired by superior performance and emergent abilities, VLMs are even considered as one of the future trend towards AGI. As multiple studies [4, 59] illustrated, VLMs’ capability is closely linked to: model parameters and data quality. Therefore, expanding model scale and feeding high-quality data to improve practical performance, has become the community consensus. As expected, VLMs show superior results on various downstream tasks/domains [22, 48, 67], including understanding and generation, with delicate architectures.

*: These authors contribute equally to this work.




Fig. 1: **Left:** the trouble with applying VLMs is the high-cost issue. **Right:** to accelerate VLMs, most existing ideas focus on the model perspective (pruning & quantization). While our Turbo explores de-redundancy from the data perspective.

Although it seems that building leading to AI has been completed, there is still a dark cloud hanging over VLMs, making everything a utopia. That is, the cost issue (computational throughput/latency/memory), whether it is training or deployment. Let’s imagine, for text-to-image generation, the snail’s pace of one picture every 10 seconds is one bottleneck for real-world applications. As a result, reducing costs becomes key to promoting the popularity of VLMs.

To achieve acceleration for VLMs, existing methods [15, 38, 50, 58] focus on one **Model-Center Perspective (MCP)**, *e.g.*, pruning, distillation, and quantization, as shown in Figure 1. Although effective, they suffer from tricky issues: *Specificity*: MCP is only developed for specific architectures with poor generalization, implying it’s incompatible with various VLMs, *e.g.*, MCP acceleration suitable for understanding is infeasible for generation. *Re-training*: To maintain high performance, MCP usually requires re-training or fine-tuning VLMs, which inevitably consumes additional computing overheads, being inefficient to apply. *Complexity*. The development of MCP involves considerable tricks, which raise barriers to applications. To sum up, with only efforts on the model side, MCP is trapped in a dilemma: still high costs, and poor generalization.

To jump out of the dilemma, as AI systems typically cover “data & model”, this paper raises one novel question about acceleration, from one **Data-Center Perspective (DCP)**: *Is there redundancy in the data side? And if so, how high?* To answer the question, we first define, then evaluate the informativity of token sequences in each VLMs block. The exploration motivation is that attention-based networks have emerged as a dominant architecture among VLMs, resulting in a quadratic relationship between computation and input sequence length. The results show that, token redundancy is consistently high in most blocks. Besides, token redundancy gradually increases as blocks deepen. With these observations, we conclude that acceleration by data de-redundancy is promising.

To enable data acceleration, we propose one novel Turbo plug-in, with the spirit that compressing invalid components while retaining semantic essences. Concretely, we define an information degree, covering two components: mutual redundancy and semantic value. The former evaluates the information duplication between sequential tokens, while the latter focuses on the token’s contribu-

tion to sample-wise semantics. For mutual redundancy, the insight is that tokens with dependency tend to have similar informativity, making re-use feasible. For semantic value, the insight is that tokens with core contributions maintain strong performance as principal components. Using information degree of sequential tokens, Turbo naturally sorts data to only leverage the top-level ones for VLMs’ calculation, thus saving costs from the source. To sum up, Turbo wins considerable pros. Universality. The accelerated data is compatible with various VLMs, *e.g.*, understanding/generation, multi-/uni-modality, showing the powerful generalization. Plug-and-Play. Turbo involves no training parameters, which is lightweight to avoid trivial development efforts. Practicality. Turbo is user-friendly with no need for tedious tricks. And even more valuable, it could superpose on existing model-perspective acceleration.

Under two types of VLMs, namely, image-align-text understanding and text-to-image generation, we experiment to prove the acceleration generality of Turbo, across several standard datasets. For almost all understanding tasks (retrieval, classification, caption and VQA), Turbo improves throughput by around 2X with little loss of performance. For most generation tasks (text-to-image and image-to-image), Turbo improves throughput to 1.6X without compromising quality. We also ablate the component effectiveness, both quantitatively and qualitatively.

To sum up, our contributions lie in three folds:

- We explore a data-perspective acceleration for VLMs. It’s generally adapted to understanding & generation, is orthogonal to model-perspective acceleration.
- We design one novel Turbo for data de-redundancy, by evaluating information degree of redundancy and semantics. Turbo serves as training-free plug-in for VLMs to achieve great trade-offs between efficiency-performance.
- We conduct extensive experiments and thorough ablations to reveal good acceleration of Turbo, and our superior results for understanding/generation.

2 Related Work

Vision-Align-Language Understanding gets image-text shared embeddings, by pre-training with large-scale data. Recent studies mainly divided into: single tower [8, 29], twin towers [20, 45], and bridge tower [30, 31, 68]. Structurally, attention-based Transformer has emerged as a dominant architecture, greatly increasing computing overhead while improving performance. They have brought many promising potentials in terms of understanding videos [25–27], audio [34, 36, 37] and image [7, 10, 11, 65]. They are broken down into: recognition [9, 28, 66], grounding [24, 35], segmentation [41, 42, 64], caption [39, 44], retrieval [22, 54].

Vision-Language Generation aims to build compositional models from text & noise to pixel-wise vision. Typically, research lines are mainly divided into: Diffusion (SD) [17, 48] and DALL-E [46, 52]. Although they have shown promising potentials [6, 23, 40], expensive computing costs seriously damage user experience. For an instance, when using SD for one $1024 * 1024$ image, the inference delay is about 8 seconds. Structurally, SD’s UNet accounts for most overhead.

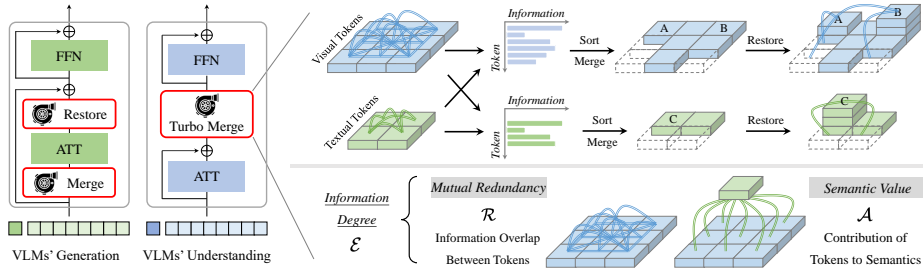


Fig. 2: Computing Architecture. As one plug-in module, Turbo compresses data to cut computing overheads for various VLMs, across understanding/generation and uni-/multi-modality. It sorts then merges tokens by information degree (mutual redundancy \mathcal{R} and semantic value \mathcal{A}) for understanding tasks; while sorts, merges and restores VLMs’ tokens for generation tasks, owning good universality and practicality.

Acceleration for VLMs. To alleviate high computing costs in the uni-modal transformer-based models, extensive acceleration studies has been proposed [3, 14, 15, 32, 33, 47, 60–63]. However, such an exploration on VLMs is still in its infancy, mainly because of the difficulty to unify different modalities under one acceleration paradigm. Nowadays, most solutions for VLMs acceleration are from the model-perspective: knowledge distillation [13, 58], floating quantization [15, 50], and model pruning [21, 51, 58]. Although achieving high sparsity with competitive performance, they all suffer from multiple issues. *Re-training.* To maintain high performance, fine-tuning VLMs is usually inevitable, which is inefficient to bring additional computing burdens. *Specificity.* They are generally developed for specific networks, which have poor generalization to compatible with various VLMs. *Complexity.* They can involve many empirical tricks, these trivial matters raises barriers to application. Overall, in contrast, Turbo is totally training-free, can be easily plugged in most VLMs with competitive performance.

Recently, ToMe [2] is designed as an acceleration plug-in for vision transformers, but Turbo differs from it in two aspects. *Universality:* ToMe is designed for uni-modal image classification; while Turbo is generally compatible to various VLMs: understanding/generation, multi-/uni-modality. *All-round:* ToMe focuses solely on mutual redundancy; while our Turbo combines mutual redundancy and semantic value, to achieve better acceleration-performance trade-offs.

3 Methods: Acceleration for VLMs

To alleviate heavy deployment costs from VLMs, we conduct thorough analysis in terms of data; then describe informativity-driven turbo plug-in for acceleration.

Revisit of Attention. For VLMs, Transformer composed of attention has emerged as a dominant architecture. For either text or image data, we generally processes it as the n -token 1D sequence $\mathbf{X} \in \mathbb{R}^{n \times D}$, then pass into the attention layer to model sequence relationships globally, where one weighted sum of values

based on the affinity over other tokens is calculated. Attention is formulated as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{D}}\right)\mathbf{V}. \quad (1)$$

Transformer achieves impressive performance, comparing to previous architectures, *e.g.*, convolution and graph, but at the cost of being more expensive. It is because the complexity of multi-head attention modules (self-attention or cross-attention) scales quadratically with the sequence length n . And thus, long tokens lead to substantial computational overheads.

With the fundamental understanding, we raise one novel question: *Is there any redundancy in the token sequence? If so, we could remove such redundancy from one data perspective, to cut off computational overheads from the source.*

3.1 Multi-Modal Informativity

Motivation & Insight. Given that most input data are likely to contain superfluous parts, *e.g.*, image background, irrelevant objects, we foresee a significant opportunity to further “compress” data. To accelerate VLMs from the data perspective, the primary process is to distinguish where the redundancy lies. And to analyse it quantitatively, we define token sequence informativity as a measure to quantify the information contained in one token sequence.

Inspired by the informativity idea in information theory, we define $(\Omega, \mathcal{F}, \mathbb{P})$ as the probability space of token sequences, where the triplet is sample space, σ -algebra, and probability measure. Following the concept of self-information [49], we define the informativity of one token sequence $\mathbf{X} = [x_1, x_2, \dots, x_n]$ as:

$$\mathcal{I}(\mathbf{X}) = -\log \mathbb{P}(\mathbf{X}) = -\log \mathbb{P}([x_1, x_2, \dots, x_n]). \quad (2)$$

By the compound probability formula, we have

$$\mathcal{I}(\mathbf{X}) = -\log \mathbb{P}(x_{k,k \in \psi} | x_{i,i \in \{\phi - \psi\}}) \cdot \mathbb{P}(x_{i,i \in \{\phi - \psi\}}), \quad (3)$$

where $\phi = \{1, \dots, n\}$ and ψ is a subset of ϕ . If we can find ψ such that

$$\mathbb{P}(x_{k,k \in \psi} | x_{i,i \in \{\phi - \psi\}}) \approx 1, \quad (4)$$

then the informativity of the token sequence can be approximated by

$$\mathcal{I}(\mathbf{X}) \approx -\log \mathbb{P}(x_{i,i \in \{\phi - \psi\}}). \quad (5)$$

As revealed by MAE series [16, 18, 56], there exists considerable redundancies in visual tokens, such that only 25% tokens almost restore the whole sequence, validating the existence of subset ψ . And hence, the subset $\{x_{k,k \in \psi}\}$, having a total dependency on its complementary set $\{x_{i,i \in \{\phi - \psi\}}\}$, enables us to aggregate $\{x_{k,k \in \psi}\}$ into $\{x_{i,i \in \{\phi - \psi\}}\}$ for efficient calculation with little information loss.

Mutual Redundancy. Based on the analysis, we aim to find the dependencies between tokens of the given sequence, so that helping discover subset ψ .




Fig. 3: Empirical Evaluation of Token Redundancy & Attention Concentration on BLIP fine-tuned for multi-modal retrieval. Results reveal the non-negligible redundancy in the token sequence from perspectives of semantics and similarity.

One intuition is, tokens with mutual dependency tend to have similar representations, *e.g.*, we can easily restore the texture of a shirt by just a few patches from it. What’s more, tokens with high similarity have analogical contributions in the attention calculation. Hence, we believe that by merging tokens with high mutual similarity, $\{x_{k,k \in \psi}\}$ with a high dependency on $\{x_{i,i \in \{\phi-\psi\}}\}$ can be constructed. Formally, for token x_i , we define its mutual redundancy \mathcal{R}_i to be:

$$\mathcal{R}_i = \text{Max}\{\mathcal{S}(x_i, x_j), j \in \{1, \dots, n\} \setminus i\}, \quad (6)$$

where $\mathcal{S}(\cdot, \cdot)$ refers to cosine similarity, and Max is the maximum operation.

Semantic Value. Mutual redundancy for data merging could already get relatively good results on simple tasks, *i.e.*, image classification. However, in this process, each token is treated equally, thus its contribution to semantic categories are neglected. This will result in an early merging for tokens with important semantics, causing a drastic performance drop on fine-grained, information-demanding tasks, *i.e.*, multi-modal understanding and cross-modal generation.

To merge tokens with preferences from trivial background to significant foreground, we use the guidance of semantic value hidden inside network. By defining \mathbf{Y} to be the *cls* token of the output sequence \mathbf{X}_{out} from \mathbf{X} through the attention block, We treat \mathbf{Y} as rich semantic guidance, calculate its informativity as:

$$\mathcal{I}(\mathbf{Y}) = -\log \mathbb{P}(\mathbf{Y}) = -\log \mathbb{P}(\mathcal{A}\mathbf{V}), \quad (7)$$

where $\mathbf{V} \in \mathbb{R}^{n \times D}$ is the affine transformation of the token sequence \mathbf{X} , while \mathcal{A} refers to the attention map for the *cls* token, calculated as follows:

$$\mathcal{A} = \text{softmax}\left(\frac{\mathbf{Q}_{cls} \mathbf{K}^T}{\sqrt{D}}\right) \in \mathbb{R}^{1 \times n}, \quad \sum_{i=1}^n \mathcal{A}_{1,i} = 1. \quad (8)$$

We let semantic value of i -th token as its attention weight $\mathcal{A}_{1,i}$. As we have no access to probability distribution of \mathbf{V} , we use a meandering way to approximate the solution. Inspired by the vector quantization [12, 57], one common sense can be summarized as follows. In the context of neighbourhood, continuous variables can be approximated by discrete quantities, illustrating that if we place the

perspective in a local neighborhood, vectors that are closer in distance possess similar semantics. Thus, we can make assumptions of local continuity for \mathcal{I} .

Proposition. Under the metric space defined by L_2 norm and R^n , for all $\eta_0 \in R_+$, there exists $\epsilon > 0$ for semantic-rich $y_1 \in R^n$, such that for all $y_2 \in R^n$ satisfying $\|y_1 - y_2\|_2 < \epsilon$, there exists $\eta \in R_+ < \eta_0$ to verify the following inequality (see detailed deductions in supplementary materials):

$$\|\mathcal{I}(y_1) - \mathcal{I}(y_2)\| \leq \eta. \quad (9)$$

Denote \mathbf{Y}' the *cls* token after pruning the tokens $\{x_{j,j \in \psi'}, \psi' \subseteq \{1, \dots, n\}\}$, satisfying that for ϵ' associated with $\eta'_0 \ll 1$,

$$\left\| \sum_{j \in \psi'} \mathcal{A}_{1,j} \mathbf{V}_j \right\|_2 < \epsilon', \quad (10)$$

Then we can deduce from Eq. (9) (deduction in supplementary materials) that

$$\mathcal{I}(\mathbf{Y}) \approx \mathcal{I}\left(\sum_{i=1}^n \mathcal{A}_{1,i} \mathbf{V}_i - \sum_{j \in \psi'} \mathcal{A}_{1,j} \mathbf{V}_j\right) = \mathcal{I}(\mathbf{Y}'). \quad (11)$$

Such analysis indicates that, pruning/merging tokens with sufficiently small semantic value, will barely affect the informativity of the *cls* token.

Semantic value captures token importance/relevance to the overall semantics. We inspect it into data de-redundancy, by giving tokens with high semantics less weights. This ensures tokens with high semantic contributions are more likely to be retained, even if their mutual redundancy is high. In Figure 4 and Section 4.2, we prove the properties of semantic value, and the existence of subset ψ' .

Remark. For uni-modal tasks, we adopt intra-modal *cls* token as guidance; while for multi-modal tasks, we use intra-modal or cross-modal *cls* token for semantic value guidance depending on architectures (self-/cross-attention). With such designs, we can leverage cross-modal information for data de-redundancy.

3.2 Empirical Evaluation of Informativity

Before, we identify mutual redundancy and semantic value as key factors for data de-redundancy. Here, we quantitatively measure informativity in the output token sequence $\mathbf{X} = [x_1, x_2, \dots, x_n]$ of one attention layer.

On the one hand, for mutual redundancy \mathcal{R} , we define token redundancy to measure information similarity between tokens: $\hat{\mathcal{R}} = \text{Avg}\{\mathcal{S}(x_i, x_j)\}$, where Avg is the averaging operation. On the other hand, to verify the long-tail distribution of semantic value \mathcal{A} , we define attention concentration $\hat{\mathcal{A}}$ as the semantic proportion in the top 1/4 tokens with highest values: $\hat{\mathcal{A}} = \sum_{i=1}^{\lfloor n/4 \rfloor} \mathcal{A}_{1,i}$. A larger $\hat{\mathcal{A}}$ implies a higher degree of concentration on few tokens, ensuring safe reduction to insignificant trailing tokens without losing much information.

To view underlying redundancies in VLMS, we do experiments on BLIP [31], using the parameters fine-tuned on COCO. Figure 3 counts token redundancy and attention concentration. We empirically draw the conclusions:

- **Basis:** token redundancy and attention concentration are consistently high in most blocks: 95% in intermediate layers, the left 3/4 tokens contributes only 5%.
- **Trend:** token redundancy gradually increases as blocks deepen, implying a tendency of clustering done by the network. Attention concentration presents a bottleneck trend, with the maximum value reached at the intermediate layer.

These observations support the fact that tokens in the attention layer possess high mutual redundancy and concentrated semantic value, thus acceleration by data de-redundancy is promising. Next, we achieve the goal by one novel Turbo.

3.3 Plug-and-Play Turbo

Strategy. As mutual redundancy \mathcal{R} and semantic value \mathcal{A} are two key factors to keep information and remove redundancy, we propose information degree \mathcal{E} for balancing, by two fusion strategies: weighted difference or coupled division:

$$\mathcal{E} = \mathcal{R} - \alpha\mathcal{A}, \quad \mathcal{E} = \mathcal{R}/\mathcal{A}. \quad (12)$$

Table 8 compares in detail to show the significance of trade-offs between \mathcal{R} and \mathcal{A} . With information degree \mathcal{E} , we naturally propose Turbo to accelerate by data de-redundancy. Since the data organization of most VLMs is the attention-based sequence, Turbo could perform on any block with great flexibility.

Role. Turbo behaves in differentiated forms for generation and understanding. For high-level understanding (classification, retrieval, caption, VQA), Turbo calculates information degree for n input tokens $\mathbf{X} \in \mathbb{R}^{n \times D}$ in each layer after the attention block, then sorts them to merge the top-level ones into the rest \mathbf{X}' .

$$\mathbf{X}' = \Phi_{\text{turbo}}^{\text{high}}(\mathbf{X}, \gamma) = \Psi_{\text{meg}}(\Psi_{\text{sort}}(\mathcal{E}(\mathbf{X})), \gamma) \in \mathbb{R}^{(n-\gamma) \times D}, \quad (13)$$

where Ψ_{sort} refers to the sort operation based on the information degree, Ψ_{meg} is the merging operation, and γ is the drop ratio to all sequential tokens. For parallel computing, we set γ to one constant in each batch.

For generative tasks (text-to-image and image-to-image), fine-grained pixel requirements make such simple merging of redundant tokens infeasible. We hence divide the Turbo into more modules, making its role become merge & restore. Specifically, before each VLMs' block, Turbo calculates and records information degree between all tokens, then merges redundant tokens using Ψ_{meg} ; while after each VLMs' block, Turbo restores the merged tokens by Ψ_{rest} , that is, weighted sum related and unmerged tokens, with reference to information degree.

$$\mathbf{X}' = \Phi_{\text{turbo}}^{\text{low}}(\mathbf{X}, \gamma) = \Psi_{\text{rest}}(\Psi_{\text{meg}}(\Psi_{\text{sort}}(\mathcal{E}(\mathbf{X})), \gamma), \mathcal{E}(\mathbf{X})) \in \mathbb{R}^{n \times D}. \quad (14)$$

Such Turbo greatly reduces computing overhead for VLMs' blocks, while ensuring pixel-level generation. Besides that, Turbo also allows the bulk of computing to be done in parallel, making it friendly to the modern GPU devices. For more implementation details, please refer to supplementary materials.

Remark. Turbo wins many pros, compared to existing acceleration. *Universality.* It performs data de-redundancy, *i.e.*, reducing computing from source of the AI

Table 1: NLVR Acceleration for BLIP [31] on NLVR2 dataset. Ratio \mathcal{T} denotes the number of reduced tokens per layer. We also apply Turbo to the VLMs pruned by UPop, proving that our Turbo is perpendicular to the model-perspective accelerations. Turbo surpasses other methods by a large margin on inference speed and accuracy.

Plug-In	Performance		Ratio \mathcal{T}	Acceleration	
	Dev	Test		FLOPs	Throughput
-	82.5	83.3	-	132.5	117.5
UPop [51]	80.3	81.1	-	89.4 (-0.32 \times)	138.5 (1.18 \times)
Redundancy	79.3	79.5	35	77.2 (-0.42 \times)	182.9 (1.56 \times)
Semantics	80.2	81.0	35	77.2 (-0.42 \times)	186.7 (1.59 \times)
Turbo	81.4	82.2	35	77.2 (-0.42 \times)	182.8 (1.56 \times)
	80.5	81.5	45	62.2 (-0.53 \times)	224.2 (1.91\times)
Turbo+UPop	79.0	79.8	35	54.7 (-0.59\times)	188.2 (1.60 \times)

Table 2: VQA Acceleration for BLIP [31] on VQAv2 dataset. Through applying $\mathcal{T} = 60$, we achieve the best trade-off between performance reservation and acceleration.

Plug-In	Performance		Ratio \mathcal{T}	Acceleration	
	Dev	Std		FLOPs	Throughput
-	77.4	77.5	-	92.1	148.1
UPop [51]	76.3	76.3	-	65.2 (-0.30 \times)	167.8 (1.13 \times)
Redundancy	74.8	74.8	40	65.7 (-0.29 \times)	184.8 (1.25 \times)
Semantics	75.3	75.3	40	65.7 (-0.29 \times)	188.2 (1.27 \times)
Turbo	76.8	76.9	40	65.7 (-0.29 \times)	184.6 (1.25 \times)
	76.6	76.7	60	52.5 (-0.43 \times)	232.7 (1.57\times)
Turbo+UPop	75.4	75.4	40	44.8 (-0.51\times)	206.8 (1.40 \times)

system. And the resultant data could be generally used for various VLMs, *e.g.*, understanding/generation, multi-/uni-modality, showing powerful compatibility. *Practicality.* It fully considers the potential data-repetition guided by semantics, accelerating with little performance loss. Besides, it’s user-friendly without cumbersome tricks. *Plug-and-Play.* It works as plug-in without additional parameters for re-training, and is plain to avoid trivial development efforts.

4 Experiments

Understanding Benchmarks. We evaluate Turbo on uni-modal Image Classification, and multi-modal tasks like Cross-modal Retrieval, Natural Language for Visual Reasoning, Visual Question Answering and Image Captioning. We experiment image classification with AugReg [55] and SWAG [53] on Imagenet-1k; multi-modal tasks with BLIP [31]/BLIP2 [30] on Flickr30k, COCO and NLVR2 and MiniGPTv2 [5] on Vizwiz [1], OKVQA [43] and GQA [19].

Generation Benchmarks. We evaluate on Stable Diffusion 1.5 [48] by generating 2000 images, each resolution is set to $512 * 512$. The text classes used are from ImageNet-1k. All experiments are conducted on one 3090 GPU. For more details on different benchmarks, please refer to supplementary materials.

Table 3: Caption Acceleration on COCO dataset. The drop ratio is 16 and (30, 40) for Turbo on BLIP2 and BLIP respectively. Turbo attains competitive performance while largely enhance throughput. We also leverage Turbo upon GPTQ (4-bit model quantization) and UPop (model pruning) for further inference acceleration.

VLMs	Plug-In	Performance		Acceleration	
		B@4	CIDEr	FLOPs	Throughput
BLIP [31]	-	39.7	133.3	330.2	34.2
	UPop [51]	38.6	128.9	137.9 (-0.58×)	56.6 (1.65×)
	Redundancy	35.5	120.9	134.2 (-0.60×)	67.9 (1.99×)
	Semantics	36.4	123.8	134.2 (-0.60×)	70.3 (2.06×)
	Turbo	38.2	130.0	134.2 (-0.60×)	67.6 (1.98×)
	Turbo+UPop	37.3	126.0	60.8 (-0.82×)	74.8 (2.19×)
BLIP2 [30]	-	42.7	145.7	1379.2	29.5
	GPTQ [15]	42.2	144.8	1379.2	30.1 (1.02×)
	Redundancy	40.8	137.7	1029.5 (-0.25×)	51.1 (1.73×)
	Semantics	40.3	137.9	1029.5 (-0.25×)	52.2 (1.77×)
	Turbo	42.1	142.2	989.7 (-0.29×)	50.9 (1.72×)
	Turbo+GPTQ	41.4	142.0	989.7 (-0.29×)	52.5 (1.78×)

Table 4: Acceleration on MiniGPTv2 [5]. We utilize LLaMA2-chat 7B as the LLM backbone, and 8-bit quantization for GPTQ [15], and evaluate on the VQA benchmarks (accuracy as evaluation metric). Turbo still exhibits superior performance, and it can be orthogonal to further boost the model-perspective acceleration.

Plug-In	Dataset			Ratio γ	Acceleration	
	OKVQA	GQA	Vizwiz		FLOPs	Throughput
-	57.9	59.9	56.9	-	5736.5	6.8
GPTQ [15]	57.7	59.6	55.9	-	5736.5	5.9 (0.87×)
Redundancy	54.5	52.8	55.1	12	4411.3	9.7 (1.43×)
Semantics	54.3	53.2	55.2	12	4411.3	9.8 (1.44×)
Turbo	56.3	56.7	55.7	12	4411.3	9.7 (1.43×)
	54.7	54.2	55.2	16	3853.4	11.8 (1.74×)
Turbo+GPTQ	56.1	56.6	55.7	12	4411.3	8.2 (1.21×)

4.1 Comparison with State-Of-The-Art

We carry out full experiments on uni-/multi-modal VLMs, to prove effectiveness. Two typical methods in the model-perspective acceleration, *i.e.*, GPTQ [15] for quantization and UPop [51] for model pruning, are used for full comparisons. In addition, we introduce two baselines: Redundancy and Semantics, to represent methods with only mutual redundancy and semantic value.

Multi-Modal VLMs for Understanding. Thorough experiments on a wide range of multi-modal tasks for VLMs BLIP [31], BLIP2 [30] and MiniGPTv2 [5] are done to test the effectiveness, in Tables 1, 2, 3, 4 and 5. On all benchmarks, Turbo achieves the best trade-off for acceleration and performance preservation. We achieve better or competitive performance with less time, than the quantization (GPTQ) and the model pruning idea (UPop). Compared to model-perspective accelerations that require re-training for specific VLMs, Turbo is training-free to serve as universal plug-in for various VLMs. What’s more, due to the huge cost of memory access time, FLOPs is not inversely proportional to actual throughput. This leads to poor accelerations for model-perspective meth-

Table 5: Retrieval Acceleration on Flickr30K/COCO datasets. The drop ratio is 16 and (30,40) for Turbo on BLIP2 & BLIP. Turbo gets superior results, with a drop of 0.4% on BLIP image-text retrieval, while other methods drops 3.5%, 4.0% and 4.7%.

VLMs	Dataset	Plug-In	Image-to-Text		Text-to-Image		Acceleration	
			R@1	R@5	R@1	R@5	FLOPs	Throughput
BLIP2 [30]	Flickr30k	-	97.6	100.0	89.7	98.1	717.5	10.7
		Redundancy	96.4	100.0	86.5	97.2	376.0	19.9
		Semantics	96.0	100.0	86.2	97.2	376.0	20.2
		Turbo	96.7	100.0	87.9	97.4	376.0	19.7
	COCO	-	85.4	97.0	68.3	87.7	717.5	10.8
		Redundancy	83.6	95.8	66.4	86.7	396.5	18.7
		Semantics	83.2	95.6	66.0	86.3	396.5	19.0
		Turbo	84.2	96.2	67.1	87.1	396.5	18.6
BLIP [31]	Flickr30k	-	97.2	99.9	87.3	97.6	55.5	281.0
		Redundancy	93.7	99.5	80.1	95.5	37.0	375.6
		Semantics	92.8	99.4	79.3	95.0	37.0	380.0
		UPop [51]	92.5	99.0	78.4	94.5	39.1	322.2
		Turbo	96.8	99.8	85.1	96.8	37.0	375.0
			94.2	99.4	82.9	96.0	31.1	449.2
	COCO	-	81.9	95.4	64.3	86.1	55.5	285.0
		Redundancy	74.2	92.1	56.5	80.4	36.8	381.3
		Semantics	75.3	92.6	57.2	80.9	36.8	385.8
		UPop [51]	77.4	93.4	59.8	83.1	39.1	323.2
		Turbo	78.8	94.7	61.3	84.2	36.8	380.5
			77.5	93.8	60.5	83.7	34.1	424.3

ods in reality, since these methods can’t reduce the memory access time of data. For quantization-based method, due to the dequantization process, throughput could even be lower than the original one. On the contrary, by directly merging data, our Turbo saves both memory access time and calculation amount.

NLVR: Table 1 reports results (accuracy) under different drop ratios \mathcal{Y} . Turbo performs the best even with an acceleration of $1.91\times$ compared to $1.18\times$ for UPop. We also insert Turbo into VLMs pruned by UPop, showing that Turbo is perpendicular to model-perspective acceleration, reflecting our universality.

VQA: Table 2 & Table 4 evaluate the VQA task (accuracy) for BLIP [31] and LLM-based MiniGPTv2 [5]. In general, Turbo again exceeds the other methods either on performance or acceleration, even with several large drop ratios.

Image Captioning requires fine-grained information. Table 3 reports results on BLIP and BLIP2. Relying solely on mutual redundancy or semantic value for token merging performs poorly, implying non-negligible information loss. In contrast, Turbo retains much better results with same FLOPs and throughput, proving that it keeps key information by merging background rather than foreground. Besides, Turbo could further accompany these model-perspective accelerations (such as UPop & GPTQ) to get better efficiency.

Multi-Modal Retrieval often deals with large amount of data, *i.e.* millions or even billions of image-text pairs in real scenarios, so there is an urgent demand for inference speed-up. Table 5 demonstrates that, Turbo exceeds model-perspective methods and baselines by a large margin, *i.e.* 4.3% for image-text and 6.7% for text-image on BLIP Flickr30k compared to UPop [51]. Such acceptable accuracy drops show our method helps to make large-scale retrieval efficiently.

Table 6: Generation Acceleration for Stable Diffusion. Turbo wins higher quality (FID) and faster throughput, regardless of text-to-image or image-to-image tasks.

Plug-In	Ratio \mathcal{r}	Text-to-Image		Image-to-Image	
		FID	Throughput	FID	Throughput
-	-	32.12	0.32	30.04	0.29
Redundancy	10	32.80	0.39 (1.22 \times)	30.69	0.37 (1.28 \times)
	20	32.86	0.44 (1.38 \times)	30.75	0.40 (1.38 \times)
Turbo	20	32.63	0.43 (1.34 \times)	30.48	0.40 (1.38 \times)
	30	32.77	0.50 (1.56\times)	30.56	0.49 (1.69\times)

Table 7: Uni-Modal Visual Classification on ImageNet. Comparing to ToMe [2], Turbo shows better results on all-size models, while providing same acceleration.

Model	Backbone	Plug-In	Top-1 Acc	Acceleration
ViT [55]	ViT-S/16	ToMe	81.4/79.3	1.53 \times
		Turbo	81.4/ 79.9	
	ViT-B/16	ToMe	84.5/82.6	1.62 \times
		Turbo	84.5/ 83.1	
	ViT-L/16	ToMe	85.8/84.2	1.71 \times
		Turbo	85.8/ 84.6	
SWAG [53]	ViT-B/16	ToMe	85.3/84.5	1.85 \times
		Turbo	85.3/ 84.9	
	ViT-L/16	ToMe	88.1/87.7	1.98 \times
		Turbo	88.1/ 87.9	
	ViT-H/14	ToMe	88.6/88.2	1.91 \times
		Turbo	88.6/ 88.4	

Multi-Modal Generation. Table 6 compares across text-to-image, image-to-image tasks, for fine-grained visual generation. Comparing to acceleration only by mutual redundancy, Turbo achieves better generation quality, *i.e.*, lower FID scores. While getting similar generation results, Turbo brings great throughput gains over vanilla stable diffusion models.

Uni-Modal Foundation Models. On ImageNet-1k, we experiment with two uni-modal training plans, *i.e.*, AugReg [55] and SWAG [53]. Table 7 compares Turbo with one prevalent baseline: ToMe [2]. Note that image classification is rather simple comparing to multi-modal understanding tasks, so the performance drop is relatively small for the baseline. Nevertheless, Turbo surpasses ToMe on models of all sizes and training plans, while keeping the same acceleration rate.

4.2 Ablation Study & Discussion


We here make comprehensive ablations to dissect components. Without loss of generality, all analysis below are conducted on BLIP image captioning [31].

Effectiveness of Mutual Redundancy & Semantic Value are studied in Table 8. As reported in Section 3.1, mutual redundancy reveals the hidden dependency relationship between tokens. Adding mutual redundancy solely or on semantic value, witnesses obvious boosts in the performance. Moreover, semantic value also has an apparent enhancement for model performance. By balancing two components, we launch the information degree to achieve the best results.

Table 8: Ablation Study on Key Components. We validate the effectiveness of each component, and compare fusion strategies. This result proves the necessity of both mutual redundancy and semantic value, with a prominent performance boost.

	Redundancy \mathcal{R}	Semantics \mathcal{A}	Fusion \mathcal{E}	B@4	CIDEr	Throughput
A1			-	34.5	112.8	73.6
A2	✓		-	35.5	120.9	67.9
A3		✓	-	36.4	123.8	70.3
A4	✓	✓	+	38.2	130.0	67.6
A5	✓	✓	×	38.2	129.9	62.8

Fig. 4: Ablation Study on Drop Ratio γ . Semantic value retains superior performance when γ is small, mutual redundancy possesses better stability on the large γ . By combining these two components, Turbo obtains competitive results and stability on the whole scope.



Key Components with Respect to Drop Ratio. Figure 4 also investigates the performance fluctuation within a wide range of drop ratio γ . (a) Scheme with only mutual redundancy is relatively stable on the large scope but performs badly on small γ . This accords with our prediction in Section 3.1 that mutual redundancy can’t distinct important tokens. Merging semantic-rich tokens at early stages inevitably loses fine-grained information. (b) Scheme solely guided by semantic value possesses satisfying results on small γ but drops dramatically on large γ . This is because that the assumption set in Eq. (10) no longer holds, resulting in a dramatic drop in performance. By combing semantic value with mutual redundancy, we achieve superior performance on all the scope of γ .

Information Degree. In Eq. (12), we adopt two fusion strategies. Table 8 shows that coupled division and weighted difference achieve similar performance. Due to the complexity of multiplication over addition, coupled division is slower than weighted difference. Moreover, several complex strategies are studied: dynamic α on different layers (see supplementary materials), but they all result in minor gains, so we use weighted difference for efficient merging.

Robustness. The balancing coefficient α in Eq. (12) may affect the Turbo performance. Hence, Figure 5 experiments to choose α on image captioning for BLIP (ViT-Base and ViT-Large). From the range $\alpha \in \{1, 2, \dots, 20\}$, the performance fluctuates slightly on both models with different sizes, indicating that Turbo is unaware of α within a certain range. This proves our robustness.

4.3 Visualization for Turbo

Image Classification. Figure 6 shows the process of token merging, and we highlight foreground in red boxes. Comparing to ToMe [2], Turbo merges more

Fig. 5: Ablation Study of Balancing Coefficient α . On image captioning using BLIP (VIT-Base and VIT-Large), these results prove our robustness, as the performance varies slightly.

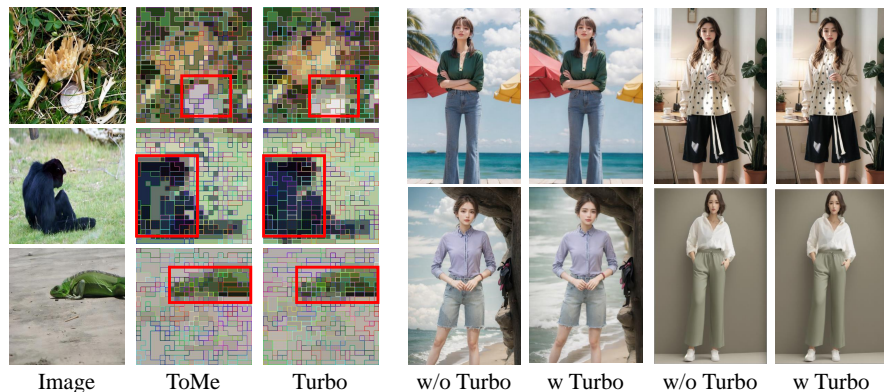



Fig. 6: Visualizations. Left: Turbo merges background patches, while retains foreground patches with semantics, preserving more key information. Right: The quality of text-to-image generation is close before and after Turbo acceleration.

background patches, while retaining most foreground patches with semantics. Thus, data compressed by Turbo keeps comprehensive/fine-grained information.

Text-to-Image Generation. Figure 6 shows the results, using text prompts such as: one female model wearing purple long sleeves and blue jeans stands on the coast. Although Stable Diffusion [48] still needs improvement in generation details (such as hands), Turbo acceleration has almost no side effects.

5 Conclusion

We propose one novel Turbo plug-in for VLMs’ acceleration, from the data perspective. In pursuit of speed-performance trade-offs, Turbo defines information degree for data reduction, taking into account two critical components: mutual redundancy and semantic value. The former evaluates data overlap between sequential tokens; the latter evaluates token’s contribution to overall semantics. By eliminating data redundancy from the source, Turbo is generally compatible to various VLMs with trivial development efforts. Extensive experiments show the significance, across multi-modality/uni-modality and understanding/generation. **Acknowledgements.** This work is supported by Alibaba Group through Alibaba Research Intern Program.

References

1. Bigham, J.P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R.C., Miller, R., Tatarowicz, A., White, B., White, S., et al.: Vizwiz: nearly real-time answers to visual questions. In: Proceedings of the 23rd annual ACM symposium on User interface software and technology. pp. 333–342 (2010)
2. Bolya, D., Fu, C.Y., Dai, X., Zhang, P., Feichtenhofer, C., Hoffman, J.: Token merging: Your vit but faster. arXiv preprint arXiv:2210.09461 (2022)
3. Bolya, D., Hoffman, J.: Token merging for fast stable diffusion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4599–4603 (2023)
4. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
5. Chen, J., Zhu, D., Shen, X., Li, X., Liu, Z., Zhang, P., Krishnamoorthi, R., Chandra, V., Xiong, Y., Elhoseiny, M.: Minigpt-v2: large language model as a unified interface for vision-language multi-task learning. arXiv preprint arXiv:2310.09478 (2023)
6. Chen, M., Chen, X., Zhai, Z., Ju, C., Hong, X., Lan, J., Xiao, S.: Wear-any-way: Manipulable virtual try-on via sparse correspondence alignment. arXiv preprint arXiv:2403.12965 (2024)
7. Chen, X., Cheng, Z., Yao, J., Ju, C., Huang, W., Lan, J., Zeng, X., Xiao, S.: Enhancing cross-domain click-through rate prediction via explicit feature augmentation. In: Companion Proceedings of the ACM on Web Conference 2024 (2024)
8. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: Proceedings of the European Conference on Computer Vision. Springer (2020)
9. Cheng, H., Ju, C., Wang, H., Liu, J., Chen, M., Hu, Q., Zhang, X., Wang, Y.: Denoiser: Rethinking the robustness for open-vocabulary action recognition. arXiv preprint arXiv:2404.14890 (2024)
10. Cheng, Z., Ju, C., Chen, X., Zhai, Z., Xiao, S., Zeng, X., Huang, W.: Image to multi-modal retrieval for industrial scenarios. arXiv preprint arXiv:2305.03972 (2023)
11. Cheng, Z., Xiao, S., Zhai, Z., Zeng, X., Huang, W.: Mixer: Image to multi-modal retrieval learning for industrial application. arXiv preprint arXiv:2305.03972 (2023)
12. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12873–12883 (2021)
13. Fang, Z., Wang, J., Hu, X., Wang, L., Yang, Y., Liu, Z.: Compressing visual-linguistic model via knowledge distillation. In: Proceedings of the International Conference on Computer Vision (2021)
14. Fayyaz, M., Koohpayegani, S.A., Jafari, F.R., Sengupta, S., Joze, H.R.V., Sommerlade, E., Pirsavash, H., Gall, J.: Adaptive token sampling for efficient vision transformers. In: European Conference on Computer Vision. pp. 396–414. Springer (2022)
15. Frantar, E., Ashkboos, S., Hoefler, T., Alistarh, D.: Gptq: Accurate post-training quantization for generative pre-trained transformers. arXiv preprint arXiv:2210.17323 (2022)
16. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)

17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* (2020)
18. Huang, P.Y., Sharma, V., Xu, H., Ryali, C., Li, Y., Li, S.W., Ghosh, G., Malik, J., Feichtenhofer, C., et al.: Mavil: Masked audio-video learners. In: *Advances in Neural Information Processing Systems* (2024)
19. Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6700–6709 (2019)
20. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *Proceedings of the International Conference on Machine Learning* (2021)
21. Jiang, C., Xu, H., Li, C., Yan, M., Ye, W., Zhang, S., Bi, B., Huang, S.: Trips: Efficient vision-and-language pre-training with text-relevant image patch selection. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. pp. 4084–4096 (2022)
22. Ju, C., Han, T., Zheng, K., Zhang, Y., Xie, W.: Prompting visual-language models for efficient video understanding. In: *Proceedings of the European Conference on Computer Vision*. Springer (2022)
23. Ju, C., Li, Z., Zhao, P., Zhang, Y., Zhang, X., Tian, Q., Wang, Y., Xie, W.: Multi-modal prompting for low-shot temporal action localization. *arXiv preprint arXiv:2303.11732* (2023)
24. Ju, C., Wang, H., Liu, J., Ma, C., Zhang, Y., Zhao, P., Chang, J., Tian, Q.: Constraint and union for partially-supervised temporal sentence grounding. *arXiv preprint arXiv:2302.09850* (2023)
25. Ju, C., Zhao, P., Chen, S., Zhang, Y., Wang, Y., Tian, Q.: Divide and conquer for single-frame temporal action localization. In: *Proceedings of the International Conference on Computer Vision* (2021)
26. Ju, C., Zhao, P., Chen, S., Zhang, Y., Zhang, X., Wang, Y., Tian, Q.: Adaptive mutual supervision for weakly-supervised temporal action localization. *IEEE Transactions on Multimedia* (2022)
27. Ju, C., Zhao, P., Zhang, Y., Wang, Y., Tian, Q.: Point-level temporal action localization: Bridging fully-supervised proposals to weakly-supervised losses. *arXiv preprint arXiv:2012.08236* (2020)
28. Ju, C., Zheng, K., Liu, J., Zhao, P., Zhang, Y., Chang, J., Tian, Q., Wang, Y.: Distilling vision-language pre-training to collaborate with weakly-supervised temporal action localization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2023)
29. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. In: *Proceedings of the International Conference on Machine Learning*. PMLR (2021)
30. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597* (2023)
31. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: *International Conference on Machine Learning*. pp. 12888–12900. PMLR (2022)
32. Liang, Y., Ge, C., Tong, Z., Song, Y., Wang, J., Xie, P.: Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv preprint arXiv:2202.07800* (2022)

33. Liu, D., Kan, M., Shan, S., Xilin, C.: A simple romance between multi-exit vision transformer and token reduction. In: *The Twelfth International Conference on Learning Representations* (2023)
34. Liu, J., Ju, C., Ma, C., Wang, Y., Wang, Y., Zhang, Y.: Audio-aware query-enhanced transformer for audio-visual segmentation. *arXiv preprint arXiv:2307.13236* (2023)
35. Liu, J., Ju, C., Xie, W., Zhang, Y.: Exploiting transformation invariance and equivariance for self-supervised sound localisation. In: *Proceedings of ACM International Conference on Multimedia* (2022)
36. Liu, J., Liu, Y., Zhang, F., Ju, C., Zhang, Y., Wang, Y.: Audio-visual segmentation via unlabeled frame exploitation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024)
37. Liu, J., Wang, Y., Ju, C., Ma, C., Zhang, Y., Xie, W.: Annotation-free audio-visual segmentation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 5604–5614 (2024)
38. Liu, Z., Wang, J., Dao, T., Zhou, T., Yuan, B., Song, Z., Shrivastava, A., Zhang, C., Tian, Y., Re, C., et al.: Deja vu: Contextual sparsity for efficient llms at inference time. In: *Proceedings of the International Conference on Machine Learning*. PMLR (2023)
39. Luo, H., Ji, L., Zhong, M., Chen, Y., Lei, W., Duan, N., Li, T.: Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neuro-computing* (2022)
40. Ma, C., Yang, Y., Ju, C., Zhang, F., Liu, J., Wang, Y., Zhang, Y., Wang, Y.: Diffusionseg: Adapting diffusion towards unsupervised object discovery. *arXiv preprint arXiv:2303.09813* (2023)
41. Ma, C., Yang, Y., Ju, C., Zhang, F., Zhang, Y., Wang, Y.: Attrseg: open-vocabulary semantic segmentation via attribute decomposition-aggregation. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)
42. Ma, C., Yang, Y., Ju, C., Zhang, F., Zhang, Y., Wang, Y.: Open-vocabulary semantic segmentation via attribute decomposition-aggregation. *arXiv preprint arXiv:2309.00096* (2023)
43. Marino, K., Rastegari, M., Farhadi, A., Mottaghi, R.: Ok-vqa: A visual question answering benchmark requiring external knowledge. In: *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*. pp. 3195–3204 (2019)
44. Mokady, R., Hertz, A., Bermano, A.H.: Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734* (2021)
45. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: *Proceedings of the International Conference on Machine Learning* (2021)
46. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* (2022)
47. Rao, Y., Zhao, W., Liu, B., Lu, J., Zhou, J., Hsieh, C.J.: Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in Neural Information Processing Systems* (2021)
48. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
49. Shannon, C.E.: A mathematical theory of communication. *The Bell system technical journal* **27**(3), 379–423 (1948)

50. Shen, S., Dong, Z., Ye, J., Ma, L., Yao, Z., Gholami, A., Mahoney, M.W., Keutzer, K.: Q-bert: Hessian based ultra low precision quantization of bert. In: Proceedings of the AAAI Conference on Artificial Intelligence (2020)
51. Shi, D., Tao, C., Jin, Y., Yang, Z., Yuan, C., Wang, J.: Upop: Unified and progressive pruning for compressing vision-language transformers. arXiv preprint arXiv:2301.13741 (2023)
52. Shi, Z., Zhou, X., Qiu, X., Zhu, X.: Improving image captioning with better use of captions. arXiv preprint arXiv:2006.11807 (2020)
53. Singh, M., Gustafson, L., Adcock, A., de Freitas Reis, V., Gedik, B., Kosaraju, R.P., Mahajan, D., Girshick, R., Dollár, P., Van Der Maaten, L.: Revisiting weakly supervised pre-training of visual perception models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 804–814 (2022)
54. Song, H., Dong, L., Zhang, W.N., Liu, T., Wei, F.: Clip models are few-shot learners: Empirical studies on vqa and visual entailment. arXiv preprint arXiv:2203.07190 (2022)
55. Steiner, A., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L.: How to train your vit? data, augmentation, and regularization in vision transformers. arXiv preprint arXiv:2106.10270 (2021)
56. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In: Advances in Neural Information Processing Systems (2022)
57. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017)
58. Wang, T., Zhou, W., Zeng, Y., Zhang, X.: Efficientvlm: Fast and accurate vision-language models via knowledge distillation and modal-adaptive pruning. arXiv preprint arXiv:2210.07795 (2022)
59. Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., et al.: Emergent abilities of large language models. arXiv preprint arXiv:2206.07682 (2022)
60. Wei, S., Ye, T., Zhang, S., Tang, Y., Liang, J.: Joint token pruning and squeezing towards more aggressive compression of vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2092–2101 (2023)
61. Wu, X., Zeng, F., Wang, X., Wang, Y., Chen, X.: Ppt: Token pruning and pooling for efficient vision transformers. arXiv preprint arXiv:2310.01812 (2023)
62. Xiao, G., Lin, J., Seznec, M., Wu, H., Demouth, J., Han, S.: Smoothquant: Accurate and efficient post-training quantization for large language models. In: International Conference on Machine Learning. pp. 38087–38099. PMLR (2023)
63. Xu, Y., Zhang, Z., Zhang, M., Sheng, K., Li, K., Dong, W., Zhang, L., Xu, C., Sun, X.: Evo-vit: Slow-fast token evolution for dynamic vision transformer. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 2964–2972 (2022)
64. Yang, Y., Ma, C., Ju, C., Zhang, Y., Wang, Y.: Multi-modal prototypes for open-set semantic segmentation. arXiv preprint arXiv:2307.02003 (2023)
65. Ye, Z., Ju, C., Ma, C., Zhang, X.: Unsupervised domain adaption via similarity-based prototypes for cross-modality segmentation. In: Domain Adaptation and Representation Transfer, and Affordable Healthcare and AI for Resource Diverse Global Health: Third MICCAI Workshop, DART 2021, and First MICCAI Workshop, FAIR 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27 and October 1, 2021, Proceedings 3. pp. 133–143. Springer (2021)

66. Zhao, P., Xie, L., Ju, C., Zhang, Y., Wang, Y., Tian, Q.: Bottom-up temporal action localization with mutual regularization. In: Proceedings of the European Conference on Computer Vision. Springer (2020)
67. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *International Journal of Computer Vision* (2022)
68. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592 (2023)