

FastVLM: Self-Speculative Decoding for Fast Vision-Language Model Inference

Divya Jyoti Bajpai

Dept. of IEOR, IIT Bombay
Mumbai, Maharashtra, India
divyajyoti.bajpai@iitb.ac.in

Manjesh Kumar Hanawal

Dept. of IEOR, IIT Bombay
Mumbai, Maharashtra, India
mhanawal@iitb.ac.in

Abstract

Vision-language Models (VLMs) have made significant strides in visual understanding and query response generation, but often face challenges of high computational cost and inference latency due to autoregressive decoding. In this work, we introduce an imitation-learning-based Self-Speculative Decoding (SSD) framework, named FastVLM, to address these limitations. Our approach employs a lightweight draft model for token generation in an autoregressive manner, while a full model verifies these tokens non-autoregressively. Accepted tokens proceed seamlessly, while rejected tokens are corrected by the full model and used to guide the draft model’s refinement. Through an imitation network, FastVLM enhances the draft model by integrating deeper-level insights from the full model’s architecture. Also, it maintains the performance integrity of the full model while training the draft model, achieving a balance between efficiency and accuracy. Our method speeds up the inference process by $1.55 - 1.85\times$ as compared to the final layer with minimal loss in performance. The source code is available at <https://github.com/Div290/SSD>.

1 Introduction

Vision-language tasks have leveraged the benefits of large Vision-Language models (VLMs) (e.g., BLIP-2 (Li et al., 2023) LLaVA (Liu et al., 2023a), LLaVA-1.5 (Liu et al., 2024), etc.) to achieve state-of-the-art results. However, their high computation and memory requirements present challenges (Samsi et al., 2023; Pope et al., 2023) to run them on resource-constrained devices, limiting their practical utility. This highlights the need to improve the inference speed of VLMs on resource-limited devices.

One of the main efficiency bottlenecks in VLMs is autoregressive decoding. This decoding paradigm generates tokens sequentially, with each

step requiring a separate invocation of the model to predict the next token given the previously generated tokens. Memory bandwidth limitations exacerbate this inefficiency, leading to overburdening computational resources and increased latency (Shazeer, 1911). For instance, utilizing BLIP-2 with decoder FlanT5-xl (Chung et al., 2024) for generating image captions can take roughly $13\times$ longer with autoregressive decoding compared to a single forward pass for a sequence of the same length (averaged over COCO (Lin et al., 2014) validation images). This necessitates innovative decoding methods that can overcome these constraints.

Approaches such as pruning (Michel et al., 2019; Frantar and Alistarh, 2023; Wang et al., 2022), quantization (Zhang et al., 2020; Bai et al., 2020; Frantar et al., 2022; Zhang et al., 2024) and knowledge distillation (Sanh et al., 2019; Jiao et al., 2019; Touvron et al., 2021) have demonstrated effectiveness in deploying them on resource-constrained devices. However, these techniques require modifications to the model architecture, often involving parameter reductions in the backbone network, which can degrade overall performance.

Recently, Speculative Decoding (SD) (Chen et al., 2023) has gained attention for speeding up autoregressive decoding. This method employs two models: a lightweight draft model that generates token predictions quickly in an autoregressive manner and a larger verification model that assesses the quality of the generated tokens in parallel. By incorporating a verification step, SD achieves the accuracy of the larger model while reducing the frequency of its invocation, resulting in substantial efficiency gains.

Despite its advantages, SD requires storing and executing two separate models, which can be impractical for devices with limited resources. To address this, we introduce Self-Speculative Decoding (SSD) in VLMs, originally proposed by Draft and Verify (Zhang et al., 2023) to reduce

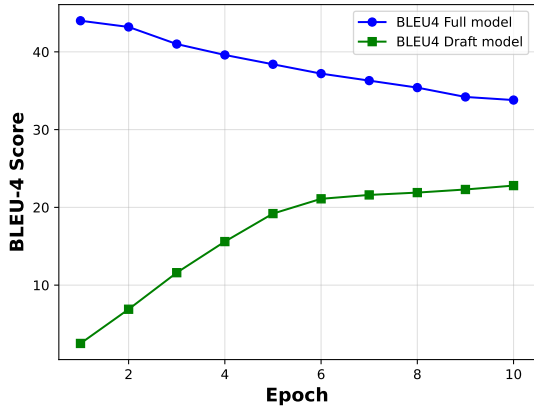


Figure 1: The effect on the final layer BLEU4 score due to combined training of the full model with the draft model.

the computational burden of autoregressive language models. It utilizes the initial layers of the primary model as a lightweight draft model, while the full model verifies the generated tokens. This parameter-sharing design addresses storage and computational challenges, making SSD appealing for resource-constrained devices. SSD’s efficiency heavily depends on the draft model’s quality: better draft predictions lead to higher token acceptance rates during verification, minimizing redundant decoding stages. However, adapting SSD to VLMs introduces significant challenges:

Loss of deeper-layer information: The draft model in SSD uses only the shallow layers of the decoder, excluding the rich multimodal features from deeper layers, which are critical for vision-language tasks. This diminishes the draft model’s accuracy and compromises overall efficiency.

Shared objectives: Sharing parameters between the draft and full model creates a trade-off between optimizing the draft model’s performance and preserving the performance of the full model. Figure 1 illustrates this trade-off using BLEU-4 scores for the draft model and the full model on the BLIP-2 architecture, with the draft model comprising the first 12 layers of FlanT5-xl. Results are based on the validation split of the COCO dataset and highlight the difficulty in jointly optimizing these objectives. These limitations necessitate a method that can improve the draft model’s performance by accessing the deeper layer representations without losing the full model performance.

To overcome these limitations, we propose FastVLM, a novel SSD framework for VLMs that trains a lightweight network to imitate the be-

haviour of the full model, combining the features of imitation learning (Hussein et al., 2017; Zare et al., 2024) and knowledge distillation. It aligns the draft model’s hidden representations with the deeper layers of the full model via cosine similarity as in (Fei et al., 2022), and further refines its predictions by distilling class probabilities. We refer to our model as an imitation-based network. The imitation network (IN) explicitly learns to recover the deeper-layer information excluded from the draft model. This is achieved by training the IN with a combination of ground-truth labels and deeper layer outputs while freezing the backbone parameters of the full model. This helps maintain its overall performance by decoupling the responsibilities of the draft model and enhancing its token generation capabilities without compromising the performance of the full model.

Combining the draft model’s shallow representations with the imitation network’s deeper-layer approximations, FastVLM significantly improves token acceptance rates during verification (see figure 6). This efficient integration ensures higher accuracy and minimizes computational overhead, offering a scalable solution. Also, since the full model and the draft model have shared parameters, the key-value (KV) cache, important for fast autoregressive decoding, can be reused for the deeper layers, where the KV cache can be reused during the verification step, significantly improving efficiency.

The main contributions of FastVLM are as follows:

- We propose FastVLM, a novel end-to-end SSD framework to improve inference latency in VLMs using a lightweight imitation-learning-based draft model.
- We develop a unique training approach for the draft model that effectively mimics deeper layers of the model, ensuring minimal information loss while maintaining the accuracy of the full model.
- By decoupling the responsibilities of the draft model, our method substantially increases token acceptance rates, leading to improved efficiency and fewer redundant computations.
- We demonstrate the effectiveness and efficiency of FastVLM through comprehensive experiments on MS-COCO (Lin et al., 2014),

NoCaps (Agrawal et al., 2019), VisDial (Das et al., 2017), MM-Vet (Yu et al., 2023), and LLaVA-Wild (Liu et al., 2024) datasets on BLIP-2 and LLaVA-1.5 model, reducing the inference time by $1.55 \times - 1.85 \times$. Some qualitative examples on BLIP-2 and LLaVA-1.5 are given in Fig. 4 and Fig. 5 respectively with discussion in Appendix A.7.

2 Related works

DNN inference: Several works have designed systems specifically engineered for DNN inference. Some methods are Orca (Yu et al., 2022), Light-Seq (Wang et al., 2020), PaLM inference (Pope et al., 2023), TurboTransformers (Fang et al., 2021), DeepSpeed Inference (Aminabadi et al., 2022), FlexGen (Sheng et al., 2023) etc. Despite these system optimizations, there are gaps in the careful co-design of algorithms and systems. This is necessary to fully harness the potential of hardware efficiency during DNN inference computation.

Faster inference methods: Various model compression techniques have been explored to accelerate inference (Zhu et al., 2024), including pruning (Frantar and Alistarh, 2023; Wang et al., 2022; Zhang et al., 2024), quantization (Zhang et al., 2020; Bai et al., 2020; Frantar et al., 2022; Wang et al., 2024; Bajpai and Hanawal, 2025d), and knowledge distillation (Sanh et al., 2019; Jiao et al., 2019; Touvron et al., 2021). These approaches statically modify the model structure.

In contrast, dynamic methods like early exiting (Bajpai et al., 2024; Bajpai and Hanawal, 2024a,b,c, 2025c,a,b; Zhou et al., 2020; Zhu, 2021; Moon et al., 2023; Huang et al., 2017; Wołczyk et al., 2021; Chataoui et al., 2023) adaptively choose sub-models during inference without altering the architecture. They leverage the fact that shallow layers often suffice for simpler inputs. While DeeCAP (Fei et al., 2022) applies imitation learning for early exit tasks, we use imitation networks to decouple the draft model’s role for speculative decoding.

Speculative Decoding methods: Speculative decoding (SD) (Chen et al., 2023; Leviathan et al., 2023; Li et al., 2024; Liu et al., 2023b; Kim et al., 2024) is a popular acceleration technique for autoregressive models. It leverages a lightweight draft model to generate multiple tokens sequentially, which are then verified in parallel by a larger, more accurate model. Recent works like Draft & Verify (Zhang et al., 2023) and LayerSkip (El-

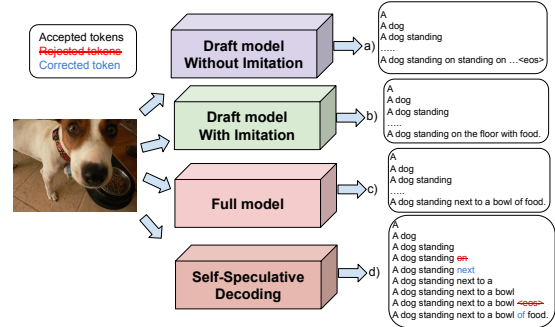


Figure 2: Inference from different model types: a) Draft without imitation; b) Draft with imitation; c) Full model; d) Self-Speculative Decoding.

houshi et al., 2024) adopt Self-Speculative Decoding (SSD), where early decoder layers act as the draft model using layer-skipping strategies.

However, applying SSD directly to multimodal models is challenging, as deeper layers are often essential for capturing cross-modal interactions. Simply skipping layers and forwarding shallow outputs to the LM Head degrades performance. To address this, we propose an imitation-based draft model that learns to replicate deeper representations while remaining decoupled from the full model. It improves draft model performance, reduces backbone degradation, and improves token acceptance rates.

3 SSD Setup for VLMs

We begin with a discussion on VLM architecture and the potential use of SSD for inference speedup. VLM consists of an encoder and a decoder where the encoder part can be an image encoder, a combination of image encoder and image-grounded query generator (Li et al., 2023) etc., and the decoder is usually an LLM. The encoder extracts the multimodal features where the input is an image and a prompt. The extracted features are then passed to the decoder for image-grounded text generation. We denote the input image as I and the input prompt as T_0 . The encoder output is denoted as $z = E(I, T_0)$. z is passed as an input to the decoder. Note that only the decoder is involved in autoregressive decoding, so we will mostly focus on the decoder of the VLM.

The decoder consists of L transformer layers with an embedding layer that maps the token indices to token embeddings x_0 . The layer l evolves embeddings output from its previous layer, $x_{l+1}^t = x_l^t + f_l(x_l^t)$ at every timestamp t and a final Language Model (LM) Head that maps the embedding

outputs of the last layer, x_L to logits $o_L^t = g(x_L^t)$ where t denotes that the output embeddings are for the t th token.

VLM decoders typically rely on an autoregressive decoding mechanism. In this framework, given a sequence of context tokens $T_0 = (y^1, y^2, \dots, y^t)$ and encoder output z , the task is to estimate the probability distribution $P(y^{t+1} | y^1, y^2, \dots, y^t, z)$ to predict the next token. The objective of the decoder is to correctly predict and estimate the output distribution’s probability. The output o_L^t is passed through a softmax layer and then the token with the highest probability is selected as the output, or alternatively, sampling can be performed based on the distribution to introduce variability and generate more diverse text outputs.

The computational burden of autoregressive decoding is comparable to performing a forward pass over the entire sequence of tokens. Each time a new token is generated, all model parameters must pass through the computational pipeline on the hardware accelerator, such as a GPU or TPU. Consequently, the model’s size, coupled with the available memory bandwidth, imposes a strict limit on decoding speed, leading to prolonged inference durations.

3.1 Self-Speculative Decoding

This approach leverages two distinct models: a primary decoder model, which represents the full-capacity model, and a faster, smaller auxiliary model referred to as the ‘draft model’. In Self-Speculative Decoding (SSD), the first n layers of the decoder model are considered the draft model. The SSD process is divided into two distinct stages:

1) Drafting Stage: The draft model predicts a batch of d tentative tokens $y^{t+1}, y^{t+2}, \dots, y^{t+d}$ based on the context sequence y^1, y^2, \dots, y^t autoregressively where $y_i = g(x_n^i)$ for $i \in \{t + 1, \dots, t + d\}$ and x_n^i denotes the hidden state of the n th layer at i th decoding step.

2) Verification Stage: The full decoder is then used to verify the correctness of the draft tokens. In a single forward pass, the full decoder computes the probability for the drafted tokens and evaluates their alignment with the context. For any token, y^j that fails verification, the primary decoder’s predicted token is used to replace y^j , and the drafting process is restarted from y^{j+1} with tokens till y^j as context.

This dual-stage process significantly reduces the number of forward passes through the full decoder model while maintaining the quality of the gener-

ated output. However, a few technical challenges hinder its direct use in VLM.

4 SSD issues in VLMs

Below, we discuss the main issue in using SSD in VLMs and how introducing an imitation network on the draft model addresses the issue.

The Role of Deeper Layers in Complex Tasks:

This section underscores the role of deeper layers in VLMs. Our analysis highlights that multimodal tasks heavily depend on deeper layers of the backbone, where interactions between visual and textual elements necessitate rich, nuanced representations. These representations are formed in the deeper layers, making them indispensable for achieving convergence toward desired performance.

However, skipping these deeper layers, as in standard SSD, can significantly impact performance in intricate tasks. For example, as shown in Figure 2, a draft model without imitation often falls into repetitive loops, predicting the same tokens without meaningful progression. To address this issue, we propose to use an imitation network. This enables the draft model to leverage deeper-layer insights, producing more accurate and coherent captions. Figure 2 is demonstrated using the BLIP-2 model with FlanT5-xl as the decoder, where our approach successfully bridges the gap between efficiency and complexity, enhancing overall system performance.

Why Imitation Network?: Imitation learning was originally proposed for Reinforcement Learning methods to train an agent to mimic the behaviour of an expert by learning from demonstrations, without explicitly relying on a reward function. We use it to mimic the behaviour of the deeper layers. The reason for using an Imitation Network (IN) and not a specific LM Head for the draft model has multiple reasons: 1) Size: The size of the LM Head is $|x_l| \times |\mathcal{V}|$ where x_l is the hidden state and \mathcal{V} is the vocabulary of tokens which is huge. For instance, the hidden size of the FlanT5-xl is 2048 and the vocab size is 32128 which scales the LM Head size to $65M$ parameters! The IN is a much simpler network with only $39M$ parameters, reducing the overall training parameters by 40%. 2) Learning objectives: The objective of the LM Head is to map the representations given by the model to class probabilities, while the objective of the draft model is to mimic the deeper layer representations and then reuse the final layer LM Head, reducing the overall complexity. Also, the IN decouples the

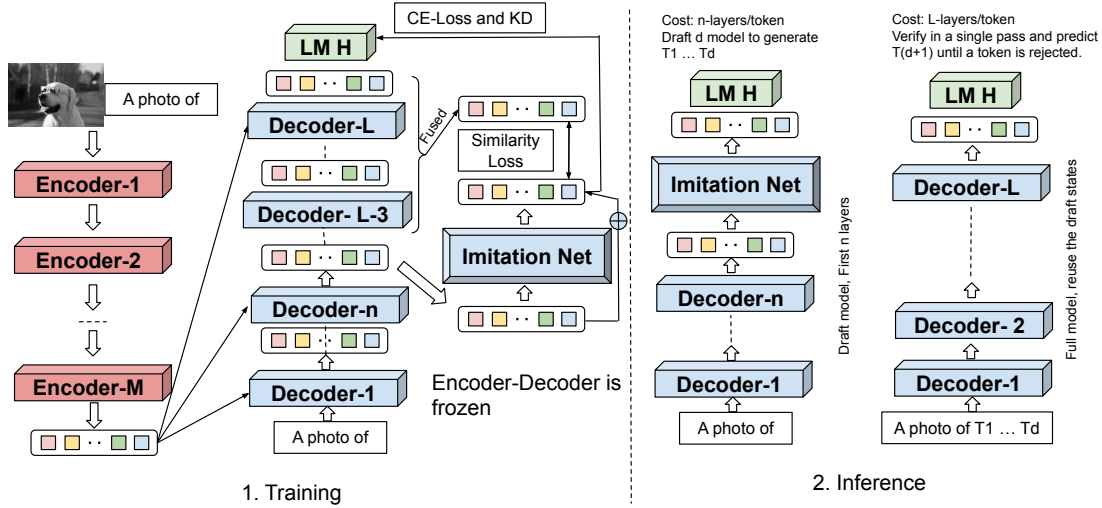


Figure 3: **Overview of our method.** **Left:** During training, the imitation network takes inputs from the n th decoder layer and learns to mimic deeper layer representations. Its output is passed to the LM Head, and training uses a combination of similarity loss, knowledge distillation, and cross-entropy loss. **Right:** At inference, the n th decoder layer and imitation network generate draft tokens, which are then verified by the full decoder.

task of the draft model, preserving the final layer performance loss as shown in Figure 1.

Our approach with IN comes as a unified solution to multiple questions. 1) It preserves the full model accuracy by decoupling the task of the draft model. 2) It is lightweight and uses minimal parameters as compared to the backbone. 3) It has the ability to mimic deep-layer representations.

5 Methodology

We start with a pre-trained VLM backbone. The procedure is divided into three parts: 1) Fine-tuning the backbone. 2) Training the imitation network. 3) Then we discuss an efficient way to re-utilize the KV-Cache without recomputing it during the verification step. First, we start with the fine-tuning of the backbone.

1) Fine-tuning the backbone: In this step, the model is fine-tuned for the final layer performance without any changes to the model architecture. The input is passed on to the full VLM and the final layer output is used to compute the loss given by

$$\mathcal{L} = \mathcal{L}_{CE}(o_L^t, y^{t*}) \quad (1)$$

where \mathcal{L}_{CE} is the cross-entropy loss, y^{t*} is the ground-truth token, and o_L^t is the output logits of the decoder’s final layer. After this step, the backbone parameters are frozen so that the optimality of the backbone is preserved.

Deep Representation Imitation: From the analysis in Figure 2 suggests that even the prediction of

simple tokens relies not only on the low-level features but also on high-level semantic information. However, directly accessing the deep representations is intractable since those are inaccessible until they feed forward the corresponding layers, which is not what we want. To bridge the gap, we approximate the uncomputed hidden states in deep layers using imitation learning. That is, we equip a lightweight network that is encouraged to predict the representation of the deeper layers based on the computed low-level representations.

Formally, the input to the imitation network \mathcal{I} is x_n^t , the output of the n th layer of the decoder at t th timestep. The imitation network is a much simpler model with fully connected layers. To find the architecture for the imitation network, we use a method similar to SelfxiT (KhademSohi et al.) (see Appendix B). Its objective is to provide a hidden representation such that it matches the deeper layer information. The output of the imitation network is denoted as $x_{\mathcal{I}}^t$. We utilize cosine-similarity loss to train the imitation network to mimic the deeper layers which is defined as:

$$\text{Cos-Sim}(x_{\mathcal{I}}^t, x_{fused}^t) = 1 - \frac{x_{\mathcal{I}}^t \cdot x_{fused}^t}{\|x_{\mathcal{I}}^t\| \cdot \|x_{fused}^t\|} \quad (2)$$

where x_{fused}^t denotes the fused hidden states from deeper layers of the backbone and $\|\cdot\|$ denotes the L2 norm. The overall loss to train the imitation network is:

$$\mathcal{L}_{imit} := \text{Cos-Sim} + KL(p_{\mathcal{I}}^t, p_L^t) \quad (3)$$

where KL is the KL divergence function defined as $KL(p_{\mathcal{I}}, p_L) = \sum_{v \in \mathcal{V}} p_{\mathcal{I}}(v) \log \frac{p_{\mathcal{I}}(v)}{p_n(v)}$. $p = \text{softmax}(g(\mathcal{I}(x_n^t)))$ and $p_L^t = \text{softmax}(o_L^t)$ and \mathcal{V} is the vocabulary, i.e., the set of all the tokens. We next discuss how we can get the fused hidden states from deeper layers.

Multi-level Representations Fusion: To perfectly guide the imitation network to learn the deeper level representation, we need to provide a good target hidden representation. For fusion, we take hidden states of the last four layers i.e. $x_L^t, x_{L-1}^t, x_{L-2}^t$ and x_{L-3}^t . The number of hidden representations that are good to fuse is explored in (Devlin et al., 2018; Horne et al., 2020) observing that the best performance is when the last three or four layers are fused together. The three methods for fusion are:

i) Averaging: All the hidden representations in different layers are averaged directly.

ii) Concatenation: All the hidden representations are concatenated in the sequence dimension and then fed into a linear transformation layer to obtain a final compressed representation.

iii) Attention-pooling: Utilizes the weighted projection of all the hidden representations as the integrated information. The attention weights are computed with the last hidden representation as the query and hold a certain robustness to noise.

These methods give the x_{fused}^t representations consisting of the fine-grained deeper layer knowledge. This guides the imitation network to generate hidden representations such that high-level information available at the deeper layers is not lost.

2) Training Draft Model: This section details the training procedure of the draft model. Draft model in our case consists of the first n layers of the backbone and the imitation model. As the backbone is frozen, the only trainable component left is the imitation network. The output of the n th layer is passed as an input to the imitation network that is used to mimic the behaviour of the remaining $L - n$ layers of the backbone. The overall loss function used for training is:

$$\mathcal{L}_{draft} := \mathcal{L}_{imit} + \mathcal{L}_{CE}(o_T^t, y^{t*}) \quad (4)$$

where \mathcal{L}_{CE} is define in Eq. 1 and \mathcal{L}_{imit} is given in Eq. 3. The left part of Figure 3 provides a visual representation of the training process.

Inference: In Algorithm 1, we outline the inference process where the input is $z = E(I, T_0)$, the encoder output, input prompt, target sequence

Algorithm 1 Self-Speculative Imitation Decoding (greedy)

```

1: Input:  $z = E(I, T_0), y^1, y^2 \dots y^t$  could be
   input prompt to decoder or just the  $\langle \text{sos} \rangle$ 
   token, target sequence length  $T$ ; max draft
   tokens to generate  $d$ .
2:  $i \leftarrow t$ 
3: while  $i < T$  or  $\langle \text{eos} \rangle$  is generated do
4:   for  $j \leftarrow i, i + 1, \dots, i + d$  do
5:      $p_{\mathcal{I}}(y|z, y^1, \dots, y^j) \leftarrow \mathcal{S}(g(\mathcal{I}(x_n^{j+1})))$ 
6:      $y^{j+1} \leftarrow \arg \max p_{\mathcal{I}}(y|z, y^1, \dots, y^j)$ 
7:   end for
8:   for  $i \leftarrow i, \dots, j$  do
9:     if  $y^{i+1} \neq \arg \max p_L(y|z, y^1, \dots, y^i)$ 
   then
10:       $y^{i+1} \leftarrow \arg \max p_L(y|z, y^1, \dots, y^i)$ 
11:      Break
12:     end if
13:      $i \leftarrow i + 1$ 
14:   If all draft accepted, then generate next
   token  $y^{i+1} \leftarrow \arg \max p_L(y|z, y^1, \dots, y^i)$ 
   and  $i \leftarrow i + 1$ 
15:   end for
16: end while
17: return  $y^1, y^2 \dots y^T$ 

```

length and number of draft tokens to be generated before verification. The inference process could be divided into two stages: 1) Draft token generation (lines 4-7): The first step is to generate the set of d draft tokens using the imitation network and the first n layers of the backbone. 2) Verification (lines 8-14): where the model verifies each generated token and if one of these tokens is not the same as the final layer’s token the verification process is stopped with the corrected token appended to the caption and the process returns to the drafting stage. This process is repeated until the target sequence length is achieved or the end-of-sentence $\langle \text{eos} \rangle$ token is predicted. For a visual representation of inference process see right side of Figure 3.

Choice of d and n : The value of d i.e., the number of draft tokens to be generated before a verification step effects the number of acceptances. A higher value of d can lead to more wasted generation from the draft model if there are rejections while a smaller value of d can lead to more number of full model calls. We perform an ablation study in Figure 9. However, from Figure 6 and Appendix A.3, we observe that the token acceptance

Models	n	BLEU1	BLEU4	CIDEr	SPICE	METEOR	ROUGE-L	Spd.
<i>BLIP2-FlanT5</i>								
BLIP2-FlanT5	24	84.1	44.0	144.6	24.6	30.9	61.8	1.00×
Draft w/o Im	12	05.2	00.1	01.7	0.00	00.4	02.5	2.00×
Draft w Im	12	65.5	18.9	86.3	14.1	17.6	47.5	2.00×
Draft & Verify	12	74.1	32.9	127.8	21.7	28.9	55.7	1.27×
LayerSkip	12	75.6	33.4	129.1	21.9	29.5	56.5	1.25×
Our	12	82.4	42.5	139.7	23.6	29.7	59.9	1.39×
Our+CS	12	83.2	43.3	141.9	24.1	30.4	60.6	1.53×
Our+CS+KD	12	83.8	43.6	142.6	24.3	30.7	61.1	1.61×
<i>BLIP2-OPT</i>								
BLIP2-OPT	32	83.5	43.7	143.1	24.4	30.9	61.5	1.00×
Draft w/o Im	15	07.9	01.3	02.5	00.1	00.8	05.2	2.13×
Draft w Im	15	65.8	19.5	91.7	15.4	18.1	47.5	2.13×
Draft & Verify	15	72.7	31.4	125.2	20.9	28.1	54.4	1.32×
LayerSkip	15	73.5	31.9	127.0	21.2	28.5	55.9	1.39×
Our	15	81.8	42.5	139.8	23.6	29.5	59.5	1.49×
Our+CS	15	<u>82.3</u>	<u>43.0</u>	<u>142.4</u>	<u>23.9</u>	<u>30.0</u>	<u>60.3</u>	<u>1.71</u> ×
Our+CS+KD	15	82.9	43.4	142.5	24.3	30.5	61.2	1.75 ×

Table 1: Results on the test split of COCO dataset, Spd. shows the average speedup in the inference. Im in the table denotes the imitation network. CS denotes cosine similarity loss and KD denotes Knowledge Distillation.

rate increases with a larger context length. We use this information to choose a dynamic value of d in terms of context length t where the value d increases with increase in t (see Appendix A.3). For the choice of n , i.e., the layer up to which the model is treated as draft model can be chosen based on user-requirements, however, we provide a method based on a reward function (see Appendix A.4).

KV cache: Key-Value cache is an important part of efficient generation in auto-regressive decoding that allows not to recompute the KV pairs at each step of the decoding. As the draft model and the verification model come from the same backbone using the same order of layers. Also, the first n layers are shared for both methods, hence, the KV cache of the first n layers is already computed, so we are able to maintain a single KV cache for the draft and verify step, saving memory and latency.

6 Experiments

We evaluate our method on image captioning and visual dialogue using BLIP-2, which enables strong performance and flexible decoder selection. For more complex vision-language tasks requiring multi-step reasoning, we use LLaVA-1.5-7B (Liu et al., 2024), built on a vision encoder and LLaMA-2 (Touvron et al., 2023) decoder. This setup covers a range of SOTA models across VL tasks. Additionally, we include results on a task-tuned CLIP-LLaMA encoder-decoder model in Appendix A.1.

Datasets: We evaluate our method on image captioning (COCO (Lin et al., 2014), No-

Caps (Agrawal et al., 2019)), visual dialog (VisDial (Das et al., 2017)), and complex multimodal reasoning (MM-Vet (Yu et al., 2023), LLaVA-Wild (Liu et al., 2024)). Captioning metrics include BLEU-4, METEOR, CIDEr, and SPICE; VisDial is evaluated using MRR; and MM-Vet and LLaVA-Wild follow the LLM-as-a-judge protocol (Liu et al., 2024). Speedup is reported as average per-token inference time reduction (Zhang et al., 2023; Elhoushi et al., 2024). Results on VQAv2 (Goyal et al., 2017) appear in Appendix A.2.

Training: We utilize two variants of the BLIP-2 model where the encoder is kept the same (ViT-g/14 (Alexey, 2020)) and the LLMs used are OPT_{2.7B} (Zhang et al., 2022) and the FlanT5-xl (Chung et al., 2024) models. For the LLaVA model, we have used LLaMA-2 (Touvron et al., 2023) as the decoder model. We have two steps during training: 1) Finetuning the backbone: Here, the full backbone is fine-tuned. It is the usual fine-tuning step (not required for the LLaVA model). 2) Imitation Network tuning: During this step, the model weights are frozen, and we attach an imitation network to the n th layer of the decoder. The architecture of the imitation network consists of two fully connected layers and one layer norm. In this step, the model is trained for 20 epochs. More hyperparameter details can be found in Table 7, 10 and 8 in the Appendix. We show the advantages of freezing the backbone parameters in Figure 7 and a discussion in Appendix A.6.

For the architecture of the imitation network, we considered options such as fully connected layers

Models	MM-Vet							LLaVA-W	Spd.
	Rec	OCR	Know	Gen	Spat	Math	Total		
LLaVA-7B	33.1	20.0	18.9	20.1	25.8	05.1	28.0	64.8	1.00×
Draft w/o Im	00.6	00.1	00.0	00.4	00.7	00.0	00.3	04.6	2.00×
Draft w Im	21.3	10.8	11.5	12.2	12.9	01.7	15.8	45.2	2.00×
Draft & Verify	26.4	15.2	14.1	16.0	17.4	03.5	22.4	53.9	1.64×
LayerSkip	28.2	16.5	15.3	17.8	20.5	04.1	25.7	58.4	1.71×
Ours	32.8	19.5	18.7	19.9	25.7	05.0	27.8	64.1	1.85×

Table 2: Results of the LLaVA-1.5-7B (LLaMA-2) model on MM-Vet and LLaVA-Wild datasets.

(up to 2) of the same size as the backbone, and transformer layers (up to one) and used Neural Architecture Search (NAS) to obtain the best architecture (see Appendix B). We observed that the combination of two fully connected layers performed better as compared to one transformer layer. Next, the fusion strategy used was the concatenation of the hidden states as it shows better performance (see Table 5). Note that the loss in performance is small across fusion strategies.

To find the optimal layer for the draft model we use the method provided in section A.4. We provide an ablation study on different values of n in Appendix A.5. We also provide details on the number of draft tokens to be generated before the verification step in Appendix A.3.

Baselines: 1) Full Model: Standard BLIP-2 inference using OPT and Flan-T5-xl decoders. **2) Draft (w/o Imitation):** Output from the n th layer is directly passed to the LM head, skipping remaining layers. **3) Draft (w/ Imitation):** The n th layer output is refined through an imitation network before passing to the LM head. **4) LayerSkip:** We implement LayerSkip as proposed, adapting it to vision-language models. **5) Draft and Verify:** Extended to VLMs; unlike LayerSkip, it halts draft token generation when confidence drops below a threshold. **6) Ours:** Our method with imitation trained using (i) cross-entropy loss, (ii) +cosine similarity (Our+CS), and (iii) +cosine similarity & knowledge distillation (Our+CS+KD).

6.1 Results

In Tables 1 and 3, we present results on the Karpathy COCO test split, VisDial, and NoCaps datasets, where our method consistently outperforms baselines. We observe a similar behaviour in Table 2, where we provide results on the MM-Vet and LLaVA-Wild datasets. The draft model without the imitation network (IN) is analogous to skipping the last $L - n$ layers and directly passing the n th layer’s representations to the LM Head, resulting in poor performance due to the mismatch with the

LM Head’s training. In contrast, the draft model with the IN demonstrates the IN’s impact by significantly improving draft performance. The performance drop in Draft & Verify and LayerSkip stems from their joint training of the draft and final layers, which compromises final-layer performance.

Our method achieves superior results due to: (1) Decoupling the draft model, which avoids competing objectives by using a small imitation network and freezing backbone parameters post-fine-tuning, preserving final-layer performance; and (2) Deeper layer mimicking, where the imitation network learns patterns from deeper layers, enhancing token acceptance and overall speedup. The better speedup in our method comes from factors such as lightweight imitation network, better performance of the draft model and mimicking deep-layer representation. We can also observe the importance of different components in the loss function used to train the backbone. Using cross-entropy loss directly, the speedup drops significantly as the draft model is not well trained. This shows that adding the Cosine Similarity and Knowledge Distillation loss makes the model more efficient. Example instances showcasing the speedup achieved by our method are shown in Figure 2.

7 Conclusion

We proposed an SSD method that can speed up the inference process in Vision-Language Models (VLMs). Our method introduces a lightweight imitation network trained to mimic the behaviour of deeper layers. We also provide a unique training approach such that the imitation network can learn to transform the input hidden representation into a representation that is a better approximation of the deeper layers. The full model performance is also preserved by decoupling the task of improving the draft model performance and the final layer performance. Experiments on multiple vision-language tasks and decoders further support our claims.



Figure 4: Some instances of the inference performed by our method where struck-out text is the words rejected at the verification step while the text in blue colour is the replacement by the full model. The struck-out text denotes the rejected tokens and the tokens in blue colour denote the corrected token by the full model while the remaining tokens are the accepted ones.

8 Limitations

Our method extends self-speculative decoding (SSD), where a draft model—comprising a subset of the full model and an imitation network—generates preliminary outputs that are later verified by the full model. This approach accelerates long-form generation by reducing full model invocations. However, in tasks like Visual Question Answering (VQA), where outputs are typically one or two tokens, the benefit diminishes since verification occurs almost as often as draft generation. Nevertheless, our method remains more efficient than using the full model alone due to shared parameters and architecture between the draft and full models—an advantage absent in traditional speculative decoding, where the two models are distinct.

9 Impact Statement

This paper presents a method to improve inference latency in Vision Language Models. The work can potentially improve the inference of large machine language models.

As usual, these can improve societal conditions, but of course, as with any technology, specific deployments need care. However, this is outside of

the scope of the present work, which is aimed at improving the basic machine-learning models.

Acknowledgements

Divya Jyoti Bajpai is supported by the Prime Minister’s Research Fellowship (PMRF), Govt. of India. Manjesh K. Hanawal thanks funding support from SERB, Govt. of India, through the Core Research Grant (CRG/2022/008807) and MATRICS grant (MTR/2021/000645), and DST-Inria Targeted Programme. We also thank funding support from Amazon IIT-Bombay AI-ML Initiative (AIAIMLI).

References

- Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. 2019. No-caps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957.
- Dosovitskiy Alexey. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv: 2010.11929*.
- Reza Yazdani Aminabadi, Samyam Rajbhandari, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Olatunji Ruwase, Shaden Smith, Minjia Zhang, Jeff

- Rasley, and 1 others. 2022. Deepspeed-inference: enabling efficient inference of transformer models at unprecedented scale. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15. IEEE.
- Haoli Bai, Wei Zhang, Lu Hou, Lifeng Shang, Jing Jin, Xin Jiang, Qun Liu, Michael Lyu, and Irwin King. 2020. Binarybert: Pushing the limit of bert quantization. *arXiv preprint arXiv:2012.15701*.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2024a. **CapEEN: Image captioning with early exits and knowledge distillation**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6458–6472, Miami, Florida, USA. Association for Computational Linguistics.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2024b. **CeeBERT: Cross-domain inference in early exit BERT**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1736–1748, Bangkok, Thailand. Association for Computational Linguistics.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2024c. **DAdEE: Unsupervised domain adaptation in early exit PLMs**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6389–6400, Miami, Florida, USA. Association for Computational Linguistics.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2025a. **Beem: Boosting performance of early exit dnns using multi-exit classifiers as experts**. *arXiv preprint arXiv:2502.00745*.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2025b. **Beyond greedy exits: Improved early exit decisions for risk control and reliability**. *arXiv preprint arXiv:2509.23666*.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2025c. **FREE: Fast and robust vision language models with early exits**. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 23599–23615, Vienna, Austria. Association for Computational Linguistics.
- Divya Jyoti Bajpai and Manjesh Kumar Hanawal. 2025d. **A survey of early exit deep neural networks in nlp**. *arXiv preprint arXiv:2501.07670*.
- Divya Jyoti Bajpai, Vivek Kumar Trivedi, Sohan L Yadav, and Manjesh Kumar Hanawal. 2024. **Splittee: Early exit in deep neural networks with split computing**. In *Proceedings of the Third International Conference on AI-ML Systems, AIMLSystems '23*, New York, NY, USA. Association for Computing Machinery.
- Joud Chataoui, Mark Coates, and 1 others. 2023. **Jointly-learned exit and inference for a dynamic neural network**. In *The Twelfth International Conference on Learning Representations*.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. 2023. **Accelerating large language model decoding with speculative sampling**. *arXiv preprint arXiv:2302.01318*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, and 1 others. 2024. **Scaling instruction-finetuned language models**. *Journal of Machine Learning Research*, 25(70):1–53.
- Jiequan Cui, Pengguang Chen, Ruiyu Li, Shu Liu, Xiao-yong Shen, and Jiaya Jia. 2019. **Fast and practical neural architecture search**. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6509–6518.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. **Visual dialog**. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 326–335.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. **Bert: Pre-training of deep bidirectional transformers for language understanding**. *arXiv preprint arXiv:1810.04805*.
- Mostafa Elhoushi, Akshat Shrivastava, Diana Liskovich, Basil Hosmer, Bram Wasti, Liangzhen Lai, Anas Mahmoud, Bilge Acun, Saurabh Agarwal, Ahmed Roman, and 1 others. 2024. **Layer skip: Enabling early exit inference and self-speculative decoding**. *arXiv preprint arXiv:2404.16710*.
- Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2019. **Neural architecture search: A survey**. *Journal of Machine Learning Research*, 20(55):1–21.
- Jiarui Fang, Yang Yu, Chengduo Zhao, and Jie Zhou. 2021. **Turbotransformers: an efficient gpu serving system for transformer models**. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 389–402.
- Zhengcong Fei, Xu Yan, Shuhui Wang, and Qi Tian. 2022. **Deecap: Dynamic early exiting for efficient image captioning**. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12216–12226.
- Elias Frantar and Dan Alistarh. 2023. **Sparsegpt: Massive language models can be accurately pruned in one-shot**. In *International Conference on Machine Learning*, pages 10323–10337. PMLR.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2022. **Gptq: Accurate post-training quantization for generative pre-trained transformers**. *arXiv preprint arXiv:2210.17323*.
- Yonatan Geifman and Ran El-Yaniv. 2019. **Deep active learning with a neural architecture search**. *Advances in Neural Information Processing Systems*, 32.

- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Leo Horne, Matthias Matti, Pouya Pourjafar, and Zuowen Wang. 2020. Grubert: A gru-based method to fuse bert hidden layers for twitter sentiment analysis. In *Proceedings of the 1st conference of the Asia-Pacific chapter of the association for computational linguistics and the 10th international joint conference on natural language processing: Student research workshop*, pages 130–138.
- Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844*.
- Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. 2017. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*.
- Haifeng Jin, Qingquan Song, and Xia Hu. 2019. Auto-keras: An efficient neural architecture search system. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1946–1956.
- Hossein KhademSohi, Mohammadamin Abedi, Yani Ioannou, Steve Drew, Pooyan Jamshidi, and Hadi Hemmati. Selfxit: An unsupervised early exit mechanism for deep neural networks. *Transactions on Machine Learning Research*.
- Sehoon Kim, Karttikeya Mangalam, Suhong Moon, Jitendra Malik, Michael W Mahoney, Amir Gholami, and Kurt Keutzer. 2024. Speculative decoding with big little decoder. *Advances in Neural Information Processing Systems*, 36.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pages 19274–19286. PMLR.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. 2024. Eagle: Speculative sampling requires rethinking feature uncertainty. *arXiv preprint arXiv:2401.15077*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916.
- Xiaoxuan Liu, Lanxiang Hu, Peter Bailis, Alvin Cheung, Zhijie Deng, Ion Stoica, and Hao Zhang. 2023b. Online speculative decoding. *arXiv preprint arXiv:2310.07177*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32.
- Microsoft. 2021. NNI: Neural network intelligence. <https://github.com/microsoft/nni>. Accessed: 2024-08-17.
- Taehong Moon, Moonseok Choi, EungGu Yun, Jongmin Yoon, Gayoung Lee, and Juho Lee. 2023. Early exiting for accelerated inference in diffusion models. In *ICML 2023 Workshop on Structured Probabilistic Inference* {\&} *Generative Modeling*.
- Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Jonathan Heek, Kefan Xiao, Shivani Agrawal, and Jeff Dean. 2023. Efficiently scaling transformer inference. *Proceedings of Machine Learning and Systems*, 5:606–624.
- Siddharth Samsi, Dan Zhao, Joseph McDonald, Baolin Li, Adam Michaleas, Michael Jones, William Bergeron, Jeremy Kepner, Devesh Tiwari, and Vijay Gadeppally. 2023. From words to watts: Benchmarking the energy costs of large language model inference. In *2023 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–9. IEEE.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Noam Shazeer. 2019. Fast transformer decoding: One write-head is all you need, 2019. URL <https://arxiv.org/abs>.
- Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Beidi Chen, Percy Liang, Christopher Ré, Ion Stoica, and Ce Zhang. 2023. Flexgen: High-throughput generative inference of

- large language models with a single gpu. In *International Conference on Machine Learning*, pages 31094–31116. PMLR.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. 2021. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Changyuan Wang, Ziwei Wang, Xiuwei Xu, Yansong Tang, Jie Zhou, and Jiwen Lu. 2024. Q-vlm: Post-training quantization for large vision-language models. *arXiv preprint arXiv:2410.08119*.
- Tiannan Wang, Wangchunshu Zhou, Yan Zeng, and Xinsong Zhang. 2022. Efficientvlm: Fast and accurate vision-language models via knowledge distillation and modal-adaptive pruning. *arXiv preprint arXiv:2210.07795*.
- Xiaohui Wang, Ying Xiong, Yang Wei, Mingxuan Wang, and Lei Li. 2020. Lightseq: A high performance inference library for transformers. *arXiv preprint arXiv:2010.13887*.
- Maciej Wołczyk, Bartosz Wójcik, Klaudia Bałazy, Igor T Podolak, Jacek Tabor, Marek Śmieja, and Tomasz Trzcinski. 2021. Zero time waste: Recycling predictions in early exit neural networks. *Advances in Neural Information Processing Systems*, 34:2516–2528.
- Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soo-jeong Kim, and Byung-Gon Chun. 2022. Orca: A distributed serving system for {Transformer-Based} generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 521–538.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*.
- Maryam Zare, Parham M Kebria, Abbas Khosravi, and Saeid Nahavandi. 2024. A survey of imitation learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics*.
- Jun Zhang, Jue Wang, Huan Li, Lidan Shou, Ke Chen, Gang Chen, and Sharad Mehrotra. 2023. Draft & verify: Lossless large language model acceleration via self-speculative decoding. *arXiv preprint arXiv:2309.08168*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, and 1 others. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Wei Zhang, Lu Hou, Yichun Yin, Lifeng Shang, Xiao Chen, Xin Jiang, and Qun Liu. 2020. Ternarybert: Distillation-aware ultra-low bit bert. *arXiv preprint arXiv:2009.12812*.
- Yuan Zhang, Chun-Kai Fan, Junpeng Ma, Wenzhao Zheng, Tao Huang, Kuan Cheng, Denis Gudovskiy, Tomoyuki Okuno, Yohei Nakata, Kurt Keutzer, and 1 others. 2024. Sparsevlm: Visual token sparsification for efficient vision-language model inference. *arXiv preprint arXiv:2410.04417*.
- Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. 2020. Bert loses patience: Fast and robust inference with early exit. *Advances in Neural Information Processing Systems*, 33:18330–18341.
- Wei Zhu. 2021. Leebert: Learned early exit for bert with cross-level optimization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2968–2980.
- Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. 2024. A survey on model compression for large language models. *Transactions of the Association for Computational Linguistics*, 12:1556–1577.
- Lucas Zimmer, Marius Lindauer, and Frank Hutter. 2021. Auto-pytorch: Multi-fidelity metalearning for efficient and robust autodl. *IEEE transactions on pattern analysis and machine intelligence*, 43(9):3079–3090.

A Appendix

A.1 Results using the encoder-decoder backbone

In Table 4, we provide the results on the encoder-decoder backbone where the encoder used is the CLIP-base model and the decoder is the LLaMA-7B model. The model was trained under the same hyperparameter setting. The results observed align to the findings on the BLIP-2 model. The input prompt in this case is ‘Describe the image’ which is passed to the CLIP model that generates the prompt-based embeddings. These embeddings are then transformed by a linear layer to be given as input to the decoder. LLAMA model then utilizes the SSD approach to draft the tokens using the draft model with imitation and then the full LLAMA backbone is used for verification.

A.2 Results on the VQA dataset

In Table 9, we provide the results on the VQA dataset. Observe that over the same setup, the speedup in this case is quite low. This happens due to the small answer length in the VQA tasks. Instead of keeping the draft token generated at once to a small value i.e., 5, we have observed that as there were many yes/no type answers and one-word answers, the speedup was not as significant as in tasks where the generated texts has a large number of words. However, the accuracy behaviour is similar to the other tasks. Note that all the baselines also have a lower speedup for similar reasons.

A.3 Choice of d

The choice of d , representing the number of draft tokens generated before a verification step, is a critical hyperparameter. Selecting an appropriate value for d can significantly reduce redundant generations from the draft model and minimize the number of calls to the full model. To evaluate the effect of a static choice of d across time steps, we conduct an ablation study, which is illustrated in Figure 9.

However, as shown in Figure 6, we observe a non-linear behavior in the token acceptance rate as the context length increases. Initially, the acceptance rate decreases, but it later increases and saturates at a higher value. This indicates that tokens generated with a larger context length are more likely to be accepted during the verification step. Based on this observation, we opt for a dynamic value of d that gradually increases as the context

length grows. Specifically, the value of d at the t -th time step is defined as:

$$d(t) = d(0) + \left\lfloor 10 \cdot \frac{1}{1 + e^{-0.01(t - \frac{N}{2})}} \right\rfloor \quad (5)$$

where we set $d(0) = 5$, ensuring that $d(t)$ remains within the range $\{5, 6, \dots, 15\}$, N denote the maximum sequence length of the model and $\lfloor \cdot \rfloor$ denotes the floor function. By allowing d to increase with the context length, we reduce the number of calls to the full model. However, setting a higher static value for d without considering context length could lead to an increase in rejections of generated tokens by the draft model, resulting in efficiency losses. On the other hand, increasing d as t progresses allows for greater draft model reliability with longer context lengths, thus maintaining efficiency without sacrificing performance.

A.4 Choice of n

In SSD, the first n layers of the backbone are used as the draft model. The choice of n models the effectiveness versus efficiency trade-off. The effectiveness comes from a higher token acceptance rate during verification where a larger n can improve it. While efficiency decreases as more computational cost is required to get the draft tokens. Thus, the choice of n is important to achieve the overall performance of the SSD method. Also, for the candidate set of layers, we take $n > L/2$ so that the draft model is not computationally heavy.

We choose the value of n by maximizing the reward function based on the confidence at an intermediate layer l i.e., $C_l := \max_{v \in \mathcal{V}} \text{softmax}(g(x_l))$ and the computational cost to predict the token from the l th layer set as $\lambda \cdot l$. We choose $\lambda = 1/L$, which denotes the per layer computational cost of the decoder. If \mathcal{D} represents the distribution of the incoming images, the total expected reward for layer l is defined as $R(I, l) = E_{I \sim \mathcal{D}}[r(I, l)]$ where $r(I, l) = C_l - \lambda \cdot l$. The layer that is chosen is defined as: $n = \arg \max_{l \in [L/2]} E_{I \sim \mathcal{D}}[r(I, l)]$ where $[L/2] = \{1, 2, \dots, L/2\}$. The empirical reward could be defined as $\hat{R}(I, l) = \frac{\sum_{i=1}^{|\mathcal{D}|} r(I, l)}{|\mathcal{D}|}$ is used to find the value of n over the validation split of the dataset. Note that the reward is such that it will maximize confidence while minimizing the computational aspects. Also, it will restrict the draft model to a layer after which there is very little improvement in confidence.

Model	VisDial	NoCaps Zero-shot									
		in-domain			near-domain		out-domain		full-dataset		Spd.
		C	S	C	S	C	S	C	S		
BLIP-2-FlanT5-xl	45.8	123.7	16.3	120.7	16.0	124.5	15.1	121.6	15.8	1.00×	
Draft w/o imi	00.3	02.4	0.00	02.1	0.00	02.0	0.00	02.3	0.00	2.00×	
Draft w imi	23.1	60.5	09.7	58.6	09.1	61.0	09.8	60.3	09.6	2.00×	
Draft & Verify	31.5	108.1	13.8	107.5	13.4	109.5	13.0	107.9	12.8	1.38×	
LayerSkip	33.2	109.9	14.1	108.2	13.7	110.0	13.2	108.3	13.0	1.45×	
Ours	43.9	122.0	15.9	119.4	15.5	123.1	14.8	119.8	15.3	1.55×	

Table 3: Results on the VisDial and the NoCaps dataset. For the NoCaps dataset, a zero-shot evaluation is performed using a model trained on the COCO dataset.

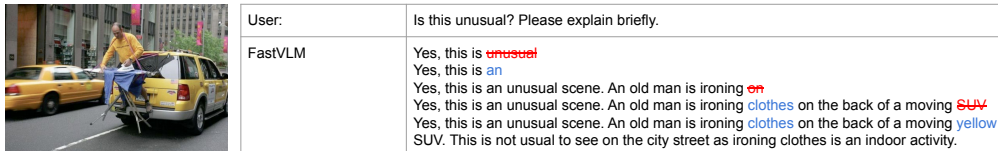


Figure 5: An instance of our method applied to LLaVA-1.5 on a long-context example.

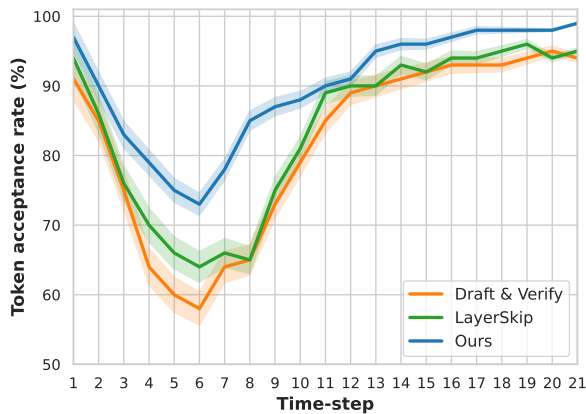


Figure 6: The token acceptance rates across different time-steps, i.e., at the t th token generation when $t - 1$ tokens are used as a context.

A.5 Ablation over λ

In table 6, we show the impact of the λ values over the full model performance and the draft model size. A higher value of λ restricts the draft model to a lower value of n as the reward function has a higher penalty for going deeper to the backbone and vice versa. Also, to show the impact of the deeper layers, if chosen as a draft model, we remove the restriction of $n \leq L/2$ only for the results in Table 6. This is done to show how the model performs in terms of efficiency as deeper layers are chosen for the draft model. One interesting observation is that as we increase λ , the value of n decreases, the model still does not take a huge hit in terms of performance that comes due to the imitation network, as it has to take input from an intermediate layer. It learns to map the representa-

tions at any intermediate layer to representations similar to the deeper layers. This makes our method robust to different values of n .

A.6 Impact of freezing the backbone

In Figure 7, we show that we save a lot of loss in performance by freezing the backbone. The loss incurred was due to shared parameters of the draft model and full model. Freezing the full model and learning only the imitation network to mimic the final layers maintains the backbone optimality and also has minimal loss in draft model performance while being substantially lower in the number of training parameters i.e., if the full backbone is fine-tuned the training parameters are $146M$ parameters and when the backbone is frozen the number of trainable parameters are down to $39M$. However, due to the deep layer imitation, the loss in draft model performance is quite small.

A.7 Some instances of SSD inference

In Figure 4, we show the multiple inference scenarios where SSD speeds up the inference process. Observe that the draft model is not able to correctly generate the very fine-grained features of the image. For example, see the first image on the left side of the figure, it shows that the draft model correctly predicts that it is a tree but the full model corrects it by predicting a ‘bare tree branch’ instead of the tree. Once the full model rejects the ‘tree’ prediction of the draft model and corrects it with bare, the draft model then correctly predicts ‘tree branch’ after ‘bare’. We observe that the draft model has some difficulty in generating rarely oc-

Baselines	n	BLEU1	BLEU4	CIDEr	SPICE	METEOR	ROUGE-L	Spd.
CLIP-LLAMA	32	81.6	41.5	138.5	22.8	29.9	60.2	1.00×
Draft w/o Im	13	07.9	01.4	03.1	00.0	00.7	07.1	2.46×
Draft w Im	13	52.3	13.7	59.2	15.9	20.6	47.0	2.46×
Draft & Verify	13	75.3	27.2	109.5	18.1	25.8	53.7	1.49×
LayerSkip	13	76.1	27.4	111.8	18.5	26.0	55.2	1.45×
Our	13	80.3	40.7	136.5	22.1	29.3	59.0	1.77×

Table 4: Results on the COCO test split where the encoder-decoder are CLIP and Llama model respectively.

Fusion-Strategy	n	B4	CIDEr	MR	Spd.
Average	15	43.1	141.0	30.3	1.68×
Concatenation	15	43.6	142.6	30.7	1.67×
Attention-pool	15	43.4	141.8	30.7	1.59×

Table 5: Results showing the impact of different fusion strategies for better knowledge transfer from deeper layers to draft model on the COCO dataset.

curing words such as ‘bare’, ‘grassy’ etc. Still, once the full model makes it correct, it starts predicting the correct tokens to complete the captions. Sometimes, due to smaller architectures, the draft model predicts common words such as ‘bunch’ in the third figure on the left side then the full model corrects it by replacing it with herd after which the draft model takes over. For longer context lengths, we use the instance of the COCO dataset image and ask it to explain briefly about the image. Observe that as the context length increases, the rejection of a token was quite small, motivating us to choose a higher value of draft token generation d before a verification step.

A.8 Runtime

For conducting the experiments, we used a set of 5 NVIDIA A6000 (48 GB) and 3 NVIDIA RTX 1080 (12 GB) GPUs. The highest training time observed was on the BLIP-2 model for fine-tuning the backbone on the COCO dataset, requiring 7 hours of runtime for 5 epochs. During the inference phase, the runtime was less than an hour for all the datasets.

B Imitation network architecture

The imitation network can adopt architectures of varying depth and width, provided its computational overhead remains lower than the efficiency it introduces. Essentially, its complexity must be significantly less than that of the subsequent layers in the backbone following the chosen value of n . The search space for such networks includes configurations with diverse layer compositions, such as combinations of dense layers and self-attention

mechanisms. However, all architectures within this space are required to produce an output that aligns with the backbone’s output distribution, ensuring the same number of classes.

In our experiments, the search space was restricted to architectures comprising up to two linear layers and a single self-attention layer with a fixed hidden size (same as the hidden size of the model). These components mirrored the structural elements of the decoder model. While the dropout rate was left as a hyperparameter to be optimized, the search process allows customization based on specific requirements or constraints, such as computational budgets.

Several tools exist for performing Neural Architecture Search (NAS), including Auto-Keras (Jin et al., 2019), Auto-PyTorch (Zimmer et al., 2021), and Neural Network Intelligence (NNI) (Microsoft, 2021). In line with (KhademSohi et al.), we utilized NNI to conduct the NAS process, customizing the evaluation to prioritize both model accuracy and computational complexity. Computational complexity was assessed using floating-point operation (FLOPs) counts. However, the NAS process proved to be computationally intensive. Note that, as the input representations vary across different n values, the optimal architecture identified by NAS may also vary. Thus, selecting an appropriate n is a prerequisite.

We found the best architecture for the COCO dataset and then fixed it to other datasets without computing it for every dataset. This is a common practice where an optimal structure is found on a dataset and then fixed across various datasets (Geifman and El-Yaniv, 2019; Cui et al., 2019; Elsken et al., 2019). This process required 15 hours of GPU runtime with the same set of GPUs given in section A.8 (faster than manual checking). The NAS process identified the optimal design as a stack of two linear layers. This is intuitive as the attention mechanisms are not meant for mimicking scenarios while linear layers can perform better when the task comes to replicate some behaviour.

λ	n	BLEU4	CIDEr	METEOR	SPICE	Spd.
1/10	5	42.4	138.7	29.8	23.9	1.88×
1/15	8	43.2	140.4	30.2	24.1	1.75×
1/L	12	43.6	142.6	30.7	24.3	1.61×
1/30	16	43.7	143.1	30.8	24.4	1.42×
1/50	21	43.9	144.0	30.9	24.6	1.17×

Table 6: Impact of the value of λ on the draft model size and the overall model performance.

Also, they have reduced computational complexity.

C Mixed Training Setting

We have explored two training setups: 1) where the full model and the draft model with imitation are trained simultaneously and 2) where the full model and the imitation network are trained separately. In figure 8, we show the final layer and draft model with imitation performance when the model is initially trained simultaneously for 5 epochs and then the backbone is frozen while only training the imitation network for later epochs on the BLIP-2-FlanT5-xl model. This method can provide a better speedup but with overall performance loss, where the loss will be due to a hit in full model performance caused by simultaneous training. Still, this method is important because it has a higher speedup, as the performance of the draft model with imitation has been improved. The reason for the better model performance of the draft model is that now the imitation network gets help from the Q-former of the BLIP-2 model, which now generates representations that are better for both the draft model as well as final layer. However, as the loss is performance is not that high, the user can treat the number of epochs for unfrozen training as a hyperparameter and achieve the desired speedup with a slight loss in performance.

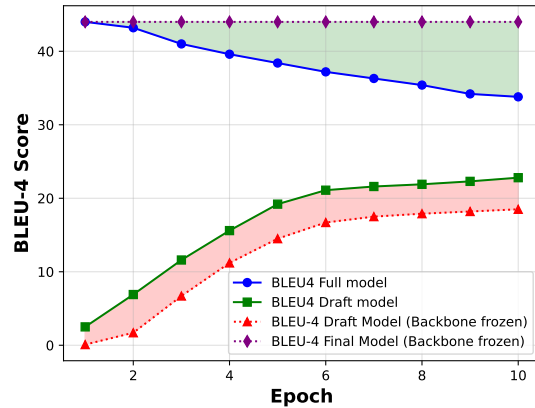


Figure 7: The figure shows how we preserve the model optimality by freezing the backbone while not losing much in terms of draft model performance.

LLM	FlanT5-xl	OPT-2.7B
Finetuning epochs	5	
Finetuning dataset	Train split	
Imitation epochs	20	
Warmup steps	1000	
Learning rate	1e-5	
Batch size	16	
AdamW beta	(0.9, 0.999)	
Weight decay	0.05	
Drop path	0	
Image resolution	364	
Prompt	"a photo of"	
Inference beam size	5	
d	8	
Fusion-strategy	Concatenation	

Table 7: Hyperparameters for training the BLIP2 backbone on image captioning datasets.

LLM	LLaMA-2
Imitation training epochs	20
Dataset-used	Instructllava150k
Warmup steps	1000
Learning rate	2e-5
Batch size	16
AdamW beta	(0.9, 0.999)
Weight decay	0.00
Image resolution	490
Inference beam size	5
d	8
n	16
Fusion-strategy	Concatenation

Table 8: Hyperparameters for training the LLaVA backbone on Instructllava-150k dataset.

Model	VQAv2		Speed
	dev	test	
BLIP2-FlanT5-xl	84.2	84.0	1.00×
Draft w/o imi	09.3	08.5	2.00×
Draft w imi	41.9	40.7	2.00×
Draft & Verify	73.5	71.8	1.29×
LayerSkip	75.8	75.0	1.33×
Our	83.9	83.4	1.39×

Table 9: Results on the VQA dataset.

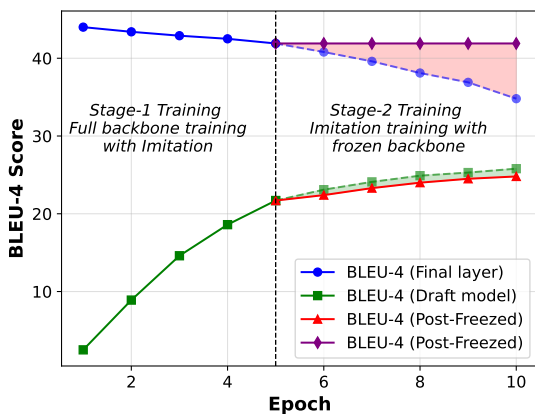


Figure 8: Mixed training where the model is simultaneously trained first and then backbone parameters are frozen and only imitation network is trained.

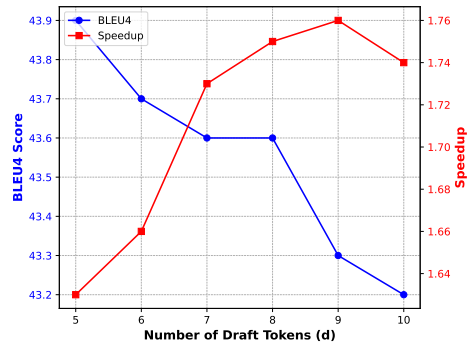


Figure 9: The impact of speedup and BLEU4 score over the change in the values of the number of draft tokens to be generated before a verification step.

LLM	FlanT5-xl	OPT-2.7B
Finetuning epochs	5	
Finetuning dataset	Train split	
Imitation training epochs	20	
Warmup steps	1000	
Learning rate	1e-5	
Batch size	16	
AdamW beta	(0.9, 0.999)	
Weight decay	0.05	
Drop path	0	
Image resolution	490	
Prompt	"Question: { } Answer:"	
Inference beam size	5	
d	5	
Fusion-strategy	Concatenation	

Table 10: Hyperparameters for training the BLIP2 backbone on VQA datasets.