

This is the peer reviewed version of the following article:

Shifted Window Fourier Transform and Retention for Image Captioning / Hu, Jia Cheng; Cavicchioli, R.; Capotondi, A. - 15293:(2025), pp. 165-179. (31st International Conference on Neural Information Processing, ICONIP 2024 nzl 2024) [10.1007/978-981-96-6596-9_12].

Springer Science and Business Media Deutschland GmbH
Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

29/04/2026 11:16

(Article begins on next page)

Shifted Window Fourier Transform And Retention For Image Captioning

Jia Cheng Hu^[0009-0008-1611-966X], Roberto Cavicchioli^[0000-0003-0166-0898],
and Alessandro Capotondi^[0000-0001-8705-0761]

University of Modena and Reggio Emilia, via G.Campi 213/b, 41125, Modena, Italy
jiacheng.hu@unimore.it, roberto.cavicchioli@unimore.it,
alessandro.capotondi@unimore.it

Abstract. Image Captioning is an important Language and Vision task that finds application in a variety of contexts, ranging from healthcare to autonomous vehicles. As many real-world applications rely on devices with limited resources, much effort in the field was put into the development of lighter and faster models. However, much of the current optimizations focus on the Transformer architecture in contrast to the existence of more efficient methods. In this work, we introduce SwiFTeR, an architecture almost entirely based on Fourier Transform and Retention, to tackle the main efficiency bottlenecks of current light image captioning models, being the visual backbone’s onerosity, and the decoder’s quadratic cost. SwiFTeR is made of only 20M parameters, and requires 3.1 GFLOPs for a single forward pass. Additionally, it showcases superior scalability to the caption length and its small memory requirements enable more images to be processed in parallel, compared to the traditional transformer-based architectures. For instance, it can generate 400 captions in one second. Although, for the time being, the caption quality is lower (110.2 CIDEr-D), most of the decrease is not attributed to the architecture but rather an incomplete training practice which currently leaves much room for improvements. Overall, SwiFTeR points toward a promising direction to new efficient architectural design. The implementation code will be released in the future.

Keywords: Image-Captioning · Retention · Fourier-Transform

1 Introduction

Image Captioning describes the problem of providing a natural language description of an image without human intervention. It is an important Vision and language task that finds application in a variety of circumstances. For instance, it is adopted in surveillance [1], visual question answering [24], autonomous vehicles [5], healthcare monitoring, and assistance for the visually impaired and blind [3]. However, many applications require the deployment of Image Captioning models in memory and computationally constrained devices which can be problematic. Over the past years, Neural Image Captioning systems achieved outstanding results. Unfortunately, current State-of-the-Art models are often

based upon Large Language Models [13, 17], which are characterized by a massive number of parameters and require abundant memory and computational resources. These aspects hamper their deployment in devices with limited resources. For this reason, much effort has been put into the design of lighter or more efficient models [16, 26, 38]. However, all these architectures are based on the Transformer [30] which is currently the standard de facto method in Neural Image Captioners defining, for instance, the trend of stateless captioning models [34, 12]. Recently, the works of [16] and [26, 22] highlighted the limitations of Attention, and the respective authors proposed attractive alternatives such as the Fourier Transform and Retention, characterized by smaller computational cost and number of parameters while performing similarly to the former approach. While their effectiveness has been proven in language modeling tasks, to the best of our knowledge, their application in image captioning has not yet been tested. In Deep learning, it is hard to predict the behavior of new methods across different fields. First, when they work well in one field their effectiveness in another can not be implied (Context Shift Problem). Secondly, different methods do not always interact well with each other. As a result, when combined, their final result is of lesser quality when compared to the one produced by the single parts (Synergy Problem [10]).




Fig. 1. Computational and parameters profiling of the baseline Image Captioning model consisting of three layers of Transformer [30] and Swin-Transformer-S [20] backbone. GPU Time is calculated as the average cost to generate captions of images from the COCO [19] test set.

In this work, we introduce the SwiFTeR architecture which combines three State-of-the-Art methods, the Fourier Transform and Swin-Transformer in the design of the visual backbone and the Retention in the fusion model decoder. In particular, we highlight better scaling, higher efficiency, and lower memory occupancy compared to existing methods. Additionally, we observe that the decrease in the caption quality is reasonable given the experimental configurations and does not suggest anti-synergy. Overall, the goals and contributions of our work are the following: *(i)* we perform experiments with the recent work of Retention [26] and Fourier Transform [16] in Image Captioning and observe that the two methods enhance model performances at the cost of 4.1 lower CIDEr-D compared to the Transformer model; *(ii)* we present SwiFT, a Shifted Window Fourier Transform visual backbone; *(iii)* Using the previous results, we design

SwiFTeR, an architecture characterized by only 20M parameters and capable of generating up to 400 descriptions per-second on NVIDIA GeForce RTX 4090. (iv) SwiFTeR currently exhibits a CIDEr-D score of 110 but leaves much room for improvements as it results from a much weaker training setup.

2 Related Works

Model compression of deep learning models is a hot topic due to the discrepancy between the increasing size of the models and the devices on which they are required to operate. Because of this, many techniques were developed to decrease memory and computational costs while preserving the quality of the output. Such examples consist of Knowledge Distillation (KD) [8], Parameter-Sharing [11], Quantization [36, 25] and Alternative-Formulations [16]. While their effectiveness is often tested on Large Language Models (LLMs) they are less tested in Image Captioning. Our work adopts the KD in the backbone but is orthogonal to other approaches since we focus on the combination of the Swin-Transformer [20], the Fourier Transform layer [16] and the Retention layer [26].

Several works tackled the optimization of Image Captioning models or part of it. Tan et al. [27] proposed a tiny model based on small Recurrent Neural Networks. Later, they leveraged sparse structures as well [28]. Wang et al. [32] combined the Knowledge Distillation with a lightweight vision backbone and a Small Language Model in the fusion model. Tan et al. [29] applied embedding token decomposition and parameter-sharing. Multiple works focused on the Transformer Vision backbone [33, 21, 15], Shashank et al. [21] proposed optimized operations and pipelines in the inference to increase hardware utilization in edge applications. Kong et al. [15] developed a latency-aware sparsity loss to perform token pruning based on the edge device constraint. Zheng et al. [37] introduced a pruning method that addresses the heterogeneous components inside the vision backbones. In contrast, our work focuses on the architectural side and less on the training strategy or the adoption of additional data to increase the caption quality.

3 SwiFTeR Design

3.1 Shifted Window Fourier Transform

In Figure 1 it can be noted that the backbone’s weights and operations occupy the most significant portion of the entire memory and forward cost respectively. To tackle this issue, we propose an adaptation of the Fourier Transform to the Swin-Transformer architecture.

The Fourier Transform consists of an operation that maps functions defined over the time dimension, into functions from the frequency domain. In [16], the authors proposed to replace the attention layer with the Discrete Fourier Transform and observed little degradation in the results in exchange for higher

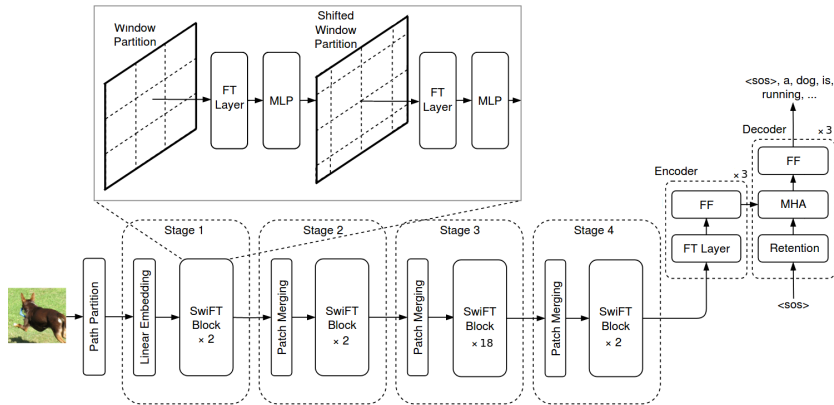


Fig. 2. The SwiFTeR architecture. Normalization layers and skip connections are omitted. FT="Fourier Transform", MHA="MultiHead-Attention", FF/MLP="FeedForward".

efficiency from both computational and memory occupation points of view. Additionally, the operation showcased better computational scalability concerning the number of input elements which is ideal in the case of Image Captioning where the number of visual features can be very large.

Let $\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{T-1}\} = X \in \mathbb{R}^{T \times H}$ be the input sequence where $\mathbf{x}_t \in \mathbb{R}^H$ for $t = \{0, \dots, T-1\}$. The Fourier Transform (FT) layer introduced in [16], applies two 1D Discrete Fourier Transform (DFT) on the input. The first, which results are denoted with $\{\mathbf{f}_0^{seq}, \mathbf{f}_1^{seq}, \dots, \mathbf{f}_{T-1}^{seq}\} = \mathcal{F}^{seq}(X) \in \mathbb{C}^{T \times H}$, is applied over the sequence length and describes the following operations:

$$\mathbf{f}_k^{seq} = \sum_{t=0}^{T-1} \mathbf{x}_t e^{-\frac{2\pi i}{T} tk}, \quad 0 \leq k \leq T-1 \quad (1)$$

The second DFT, denoted with \mathcal{F}^{hid} shares the same formulation of Equation 1 but is applied over the hidden dimension. The FT layer is ultimately formulated as:

$$FT(X) = Real(\mathcal{F}^{hid}(\mathcal{F}^{seq}(X)^\dagger))^\top \quad (2)$$

The first DFT mixes the input information along the sequence length and enables the second DFT to distribute the information of the whole sequence across the hidden dimension. The higher-order modules, e.g. the Feed-Forward, can now process the input in a bidirectional and since FT is a parameter-less operator, they hold the responsibility of learning to form meaningful compositions of the input.

Given an image or activation map of size $A = (H, W, C)$ where H is the height, W is the weight and C is the number of channels, the Windowed Fourier Transform (WFT) partitions the input into several square and rectangular windows and applies the Fourier Transform. By omitting the shifted mechanism and

border cases (we refer to [20] for details), in its essence, the layer partitions the input into N windows of size $A' = (H', W', C')$ then it treats each partition as a sequence of size $\mathbb{R}^{H'W' \times C'}$ and applies:

$$WFT(A') = Real(\mathcal{F}^{hid}(\mathcal{F}^{seq}(A')^\dagger))^\top \quad (3)$$

The advantage of the FT layer is two-fold. From the perspective of memory occupation, it does not require learned parameters. Second, the DFT can be computed using the Fast Fourier Transform Algorithm, which exhibits a complexity of $O(T \log T)$ compared to the $O(T^2)$ of the Self-Attention [30].

3.2 Retention Layer

To address the decoding cost depicted in Figure 1 we adopt the Retention [26], which describes a novel architecture where the key idea consists of integrating a recurrent aspect in the Key-Value-Query approach proposed by [30]. While the latter approach accelerates the training phase by enabling parallelism, the Transformer is slower during the inference stage because during each time step t , the decoder is required to process the entirety of previous t input elements, which is much slower compared to the single hidden state propagated through time.

To address this issue, the Retention Layer proposes a formulation that can be computed either recurrently or in parallel. Let $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\} = X \in \mathbb{R}^{T \times H}$ be the input sequence of vectors where $\mathbf{x}_t \in \mathbb{R}^H$ for $t = \{1, \dots, T\}$. Assuming for simplicity, the Single-Scale formulation and omitting the final output gate, first, the Retention computes $Q, K, V \in \mathbb{R}^{T \times H}$ as:

$$Q = (XW_Q) \odot \Theta, \quad K = (XW_K) \odot \bar{\Theta}, \quad V = XW_V \quad (4)$$

where $W_Q, W_K, W_V \in \mathbb{R}^{H \times H}$ are linear mapping. Ignoring for the moment $\Theta, \bar{\Theta}$ and D it can be proven that the following parallelizable equation:

$$Retention_{PLL}(X) = (QK^\top \odot D)V \quad (5)$$

is equivalent to the recurrent formulation:

$$\begin{aligned} S_t &= \gamma S_{t-1} + K_t^\top V_t & 1 \leq t \leq T \\ Retention_{RNN}(\mathbf{x}_t) &= Q_t S_t & 1 \leq t \leq T \end{aligned} \quad (6)$$

where $Q_t, S_t, x_t \in \mathbb{R}^H$ and S_t is the state vector.

The equivalence between Equation 5 and Equation 6 is made possible thanks to the definition of Θ and $\bar{\Theta}$, D and γ . While we refer to [26] for the proof, in this Section we propose some insights. (i) *Causality*. A causal mask, denoted with $D \in \mathbb{R}^{T \times T}$ in Equation 5 a causal mask is multiplied to Q and K . This prohibits the parallel formulation from looking at the future elements, in the absence of such a condition, a recurrent formula would not be feasible. (ii) *Linearity*. Only linear operations are involved. Therefore, the equivalence can be expressed by the

following equation $Q_t S_t = \sum_{m=1}^t Q_t A^{t-m} K_m^\top V_m$ where $S_t = A S_{t-1} + K_t^\top V_t$ and $A \in \mathbb{R}^{H \times H}$, which leads to the last insight. (iii) *Computational feasibility*. While A is absorbed into W_K and W_Q , the elements $\gamma \in \mathbb{R}^H$, and $\Theta, \bar{\Theta} \in \mathbb{R}^{T \times H}$ are defined and chosen such that A^{t-m} can be decomposed into a more numerically stable and easier to compute formulation.

Retention is adopted in the fusion model decoder for multiple reasons. First, in [26] they showcased that Retention achieves comparable performances with Attention-based models which suggests its effectiveness in understanding image descriptions. Second, the recurrent formulation directly addresses the performance bottleneck of the decoding stage, highlighted in Figure 1. In particular, assuming the cost of a single decoder iteration is quantified by the number of vectors processed at each time step, it can be reduced to $O(1)$ thanks to the recurrent formulation, compared to $O(T)$ in the case of Transformer. If we then consider a final image description of length T , the cost of Greedy Search for the recurrent module is $O(T)$, in contrast to $O(1) + O(2) + \dots + O(T) = O(T^2)$ of the stateless decoder. Finally, thanks to the dual formulation, the layer can also be trained efficiently, which can be a desirable property in the case of online training where updates need to be made frequently.

3.3 SwiFTeR

SwiFTeR stands for Swin-Fourier Transform and Retention Network. It comprises the SwiFT_{weak} backbone, which is made by replacing the Attention with the operation described in 3.1 in the Swin-Transf.-S. And a fusion model, that consists of FNet layers within the encoder and Retention in the decoder. Formally, the input image $I \in \mathbb{R}^{C \times H \times W}$ is fed to the SwiFT_{weak} backbone which generates the initial set of processed visual features $X_0 = \{x_1^0, x_2^0, \dots, x_{|X|}^0\}$, $x_i^0 \in \mathbb{R}^{d_m}$:

$$X_0 = SwiFT_{weak}(I) \quad (7)$$

A linear mapping $W_I \in \mathbb{R}^{d_m \times H}$ is applied to conform the visual features to the hidden dimension of the fusion model. The fusion model consists of N encoders and N decoders. Let us assume for simplicity $X_0 \in \mathbb{R}^H$, the n -th encoder is defined as:

$$\begin{aligned} E_n &= X_{n-1} + FT(Norm_n^{FT}(X_{n-1})) \\ X_n &= E_n + FF_n(Norm_n^{FF}(E_n)) \end{aligned} \quad (8)$$

whereas, the n -th decoder is defined as:

$$\begin{aligned} D_n &= Y_{n-1} + Retention_n(Norm_n^D(Y_{n-1})) \\ W_n &= B_n + Attention_n(Norm_n^{CA}(D_n), X_{N_{enc}}) \\ Y_n &= W_n + FF_n(Norm_n^{FF}(W_n)) \end{aligned} \quad (9)$$

where $Y_0 = \{y_1^0, y_2^0, \dots, y_{|Y|}^0\}$, $y_i^0 \in \mathbb{R}^H$. All layers are summed through a linear projection and the final output is fed to the classification layer. The complete architecture is depicted in Fig. 2.

Table 1. Ablation study of SwiFTeR architecture. The performance is evaluated on the COCO validation dataset. Top) reports the "fusion model only" ablation study. The Swin-Transf.-S backbone is adopted but the backbone parameters, FLOPs, and latency are excluded from the calculation. Captions for quality measuring are generated with beam search with a beam size of 3 in the case of Cross-entropy (XE) training and a beam size of 5 in the case of reinforcement (RF) training. Greedy decoding is adopted for the latency computation. Bottom) reports the values for the backbone+fusion model. The latency is measured for a single forward pass, assuming 49 visual features and 20 tokens. γ denotes the latency percentage compared to the baseline value. Latency is computed on a single Mobile NVIDIA GeForce RTX 3070 Ti GPU. "★"=Baseline.

Fusion model only	Train	B4	C	Params	FLOP	Latency (γ)
Normal						
Transf. Enc. & Transf. Dec. (★)	XE	34.3	114.5	32.7M	3.7G	37.0ms (100%)
Transf. Enc. & RetNet Dec.	XE	33.6	113.2	35.8M	4.1G	28.7ms (77.5%)
RetNet Enc. & Transf. Dec.	XE	32.2	107.5	35.8M	6.5G	36.0ms (97.0%)
RetNet Enc. & RetNet Dec.	XE	32.1	106.7	39.0M	6.9G	31.2ms (84.3%)
FNet Enc. & Transf. Dec.	XE	33.4	112.4	29.6M	2.7G	35.6ms (96.2%)
FNet Enc. & RetNet Dec.	XE	32.8	110.4	32.7M	3.1G	28.9ms (78.1%)
Small						
FNet Enc. & RetNet Dec.	XE	34.3	111.5	2.7M	0.1G	28.0ms (75.6%)

Backbone included	Train	B4	C	Params	FLOP	Latency (γ)
Normal (Swin-Transf.-L)						
FNet Enc. & RetNet Dec. (★)	XE	35.3	118.9	229.7M	107.0G	55.4ms (100%)
FNet Enc. & RetNet Dec.	RF	39.1	134.2	229.7M	107.0G	55.4ms (100%)
Small (Swin-Transf.-S)						
FNet Enc. & RetNet Dec.	XE	34.0	111.4	52.7M	8.8G	49.0ms (88.4%)
FNet Enc. & RetNet Dec.	RF	36.8	124.8	52.7M	8.8G	49.0ms (88.4%)
SwiFTeR (SwiFT _{weak}) (Ours)						
FNet Enc. & RetNet Dec.	XE	30.5	94.7	20M	3.1G	46.1ms (83.2%)
FNet Enc. & RetNet Dec.	RF	33.0	109.5	20M	3.1G	46.1ms (83.2%)

Normal={H=512, F=2048}, Small={H=96, F=384}

H=Hidden dimension, F=Feed Forward size.

4 Experimental Setup

4.1 Dataset and Training

We trained the visual backbone, SwiFT_{weak} on ImageNet-1k [4], which is made of only a million images belonging to 1000 classes. We adopted this dataset over the ImageNet-22k due to computational limitations during the development, the reason why it is a *weak* backbone. However, such a decision does not impact the discussion over efficiency and methods synergy. Images are resized into $3 \times 244 \times 244$ and each channel is normalized to mean=(0.485, 0.456, 0.406) and standard deviation of (0.229, 0.224, 0.225). The training of the fusion model and evaluation of the whole architecture is conducted on the renowned Microsoft COCO benchmark [19], which consists of 113 thousand images for training, 5000

Table 2. Offline comparison of State-of-the-Art models over the COCO Karpathy test split. "Pre-training" in the column "Type" denotes the fusion model size and the practice of pre-training the final model (on top of the already pre-trained backbones) on additional text or images before fine-tuning the COCO test split. In "Caption quality" the subscript XE and RF distinguish between the results with Cross-Entropy and Reinforcement training. In conformity with the literature, the FLOPs computation between the image backbone and fusion model differs. Fusion model FLOPs include multiplications and sums and we assume a length of 128 whereas image backbone FLOP computation does not count for additions.

Model		Backbone	Caption quality			Performance	
Type	Architecture	Image	C_{XE}	$B4_{RF}$	C_{RF}	Params	FLOP
1	Att2all [23]	R-CNN	99.4	34.2	114.0	60M	>760G
	UpDown [2]	F-RCNN ₁₀₁	113.5	36.3	120.1	69M	762G
	AoANet [14]	F-RCNN ₁₀₁	119.8	38.9	129.8	90.7M	766G
	ExpansionNet v2 [12]	Swin-Transf.-L	128.7	41.5	140.4	245M	115.3G
2	UNIMO _B [18]	F-RCNN ₁₀₁	124.4	-	-	172.8M	789.5G
	LEMON _B [13]	ResNeXt ₁₅₂	133.3	41.6	142.7	250.2M	1029.7G
	BLIP _B [17]	ViT _B	133.3	-	-	195.4M	78G
3	E2E-VLP [35]	ResNet-152	117.3	-	-	94M	12.8G
	MiniVLM [31]	Eff-DET	119.8	-	-	33M	8.3G
	DistilVLM [6]	Eff-DET	120.8	-	-	33M	8.3G
	LightCap [32]	ResNet-50	125.8	40.1	136.6	38M	5.3G
4	Ablation	Swin-Transf.-S	113.2	36.5	125.5	52.7M	8.8G
	SwiFTeR	SwiFT _{weak}	96.0	33.3	110.2	20M	3.1G

for validation, and 5000 for the evaluation test set. Each image is paired with five descriptions. Each reference caption is pre-processed by lowering casing and punctuation filtering. The vocabulary is made of tokens that occur at least five times resulting in 10000 unique words.

We train the backbone with Knowledge Distillation (KD) [8] using the Swin-Transformer-S (Small) [20] as the teacher. The student model SwiFT_{weak}, is trained on the ImageNet-1k [4] classification task and computes the loss as the KL-divergence on the output prediction layers summed to the MSE of the last feature maps. No image augmentation is adopted. The KD is adopted to accelerate the training rather than increase the performances, due to the heterogeneity between the teacher and student architectures.

Regarding the Fusion Model, we follow the standard practice of training models on Cross-Entropy Loss for Image Captioning [12], followed by CIDEr-D optimization [23], described by the following minimization:

$$\min_{\theta} L_{SCST}(\theta) = -\mathbb{E}_{p_{\theta}}[r(y_{1:T}^s)] \quad (10)$$

where $Y_{1:T}^s = (y_1, y_2, \dots, y_T)$ is the sampled description and r is the CIDEr-D reward function. We adopt the Standard Self-Critical Sequence Training [9]¹ and follow the training strategy of [12] up until the third step.

¹ Signature: STANDARD_wInit+Cider-D[n4,s6.0]+average[nspi5]+1.0.0

5 Results

5.1 Ablation Study

In the ablation study, we compare the performance and costs of our proposed architectures with several model variants. Since SwiFTeR changes both the backbone and the fusion block of the baseline end-to-end Transformer architecture, we split the discussion into two paragraphs. The first discusses the changes in the fusion network only. The latter analyzes the impact of changing the backbone.

In Table 1-Top we report the first ablation study which focuses on the cross-modal fusion network and uses the Base Transformer as a baseline. In detail, we keep the backbone of the Swin-Transformer-S and observe the impact of changing the encoder and decoder architectures. It can be noted that Retention in the decoder leads to a negligible performance degradation of 1.3 CIDEr-D, whereas, a significant drop of 7 CIDEr-D can be observed when adopted in the encoder. While the first degradation is acceptable as a trade-off with a much more scalable inference time, the latter is caused by the formulation of the retention which introduces intrinsic sequentiality bias and is not helpful in the visual processing. Instead, adopting the Fourier Transform in the encoder leads to an increase of 2.9 CIDEr-D compared to the Retention case and a decrease of only 1.9 CIDEr compared to the self-attentive encoder. Overall, the proposed fusion network of FNet encoder & RetNet decoder leads to a degradation of 4.1 CIDEr-D, however, we adopt this configuration in SwiFTeR as it compensated by a decrease in terms of FLOPs and latency as reported in Table 1-Top. Notably, the metrics of FLOPs and latency hide the true advantages of adopting the Retention layer. For instance, Retention exhibits higher FLOPs but latency is smaller. Additionally, the superior scalability of retentive networks to long sequence lengths is understated in Table 1. This aspect will be clarified more in Section 5.3.

Overall, we chose the configuration of "FNet Enc. and RetNet Dec." as the fusion model for SwiFTeR for efficiency reasons and because the combination of the two strategies does not showcase a significant quality degradation, therefore anti-synergy. However, it does not provide a noticeable reduction in memory occupancy. This aspect can be addressed by decreasing the model size. In the typical configuration, models have three layers and a hidden dimension of at least $H=512$ [14, 12]. Instead, we set the hidden dimension to $H=96$ and label this version with "Small" in Table 1-Top. This configuration is not only lighter but, surprisingly, it achieves slightly better CIDEr compared to the larger model, which suggests a high degree of overfitting in the previous setup.

One of the most significant costs of a Captioning system, as depicted in Figure 1 is given by the backbone. In Table 1-Bottom we report the impact of using different backbones on the performances. As expected, the Swin-Transformer-L (Large) backbone leads to the highest caption quality. When the fusion model is tested on Swin-Transformer-S, it achieves about 10 lower CIDEr-D score but is $4.35\times$ smaller and $12.15\times$ less computationally expensive. Finally, replacing the Swin-Transformer-S backbone with SwiFT_{weak}, we further reduce the parameters, achieving an $11.48\times$ compression, and $34.51\times$ reduction in FLOPs.

However, a significant decrease in caption quality is observed due to a weak training configuration, which makes SwiFT a much weaker visual backbone. This resulted from adopting ImageNet-1k instead of ImageNet-22k in contrast to the Swin-Transformer-S. Additionally, no data augmentation was performed, which is key to achieving a robust backbone [20]. These aspects will be addressed in future works but the weak training configuration does not affect the memory requirements and efficiency, the focus of our work.

5.2 Comparison With State-of-the-Art

In Table 2 we existing categorized State-of-the-Art models into three classes: Type-1: "Normal without Pre-train", Type-2: "Large with Pre-train", and Type-3: "Light with Pre-train". Type-1 models [12, 34, 2], refer to models that consist of a strong image backbone and a standard-sized fusion model trained from scratch on the MS-COCO dataset only, without additional training data. Whereas, Type-2 models [13, 17, 18], are characterized by a strong image backbone and a large fusion model which is further pre-trained on a large amount of additional data to build the foundation model for multiple fine-tuning tasks, image captioning being one of them. Both categories focus on achieving high caption quality regardless of the computational and memory cost. As a result, while they achieve up to 142.7 CIDEr-D, they are associated with hundreds of GFLOPs and hundreds of millions of parameters. In contrast, the Type-3 category [35, 31, 6, 32] focuses on efficiency but also preserving a high description quality. To achieve that, they adopt a lightweight image backbone and a small fusion model architecture that is first initialized with language modeling weights and then pre-trained on a large amount of additional image-text data. The most remarkable example is LightCap [32] which exhibits 136.6 CIDEr-D on the MS-COCO test set, using only 38M parameters and up to 5.3 GFLOPs for a single image, which is the State-of-the-Art architecture among efficient models. Our proposal falls into a fourth model category, Type-4: "Light without Pre-train". SwiFTeR aims for high inference speed and model size reduction. It comprises 20M parameters and requires up to 3.1 GFLOPs, up to $12.5\times$ smaller and $331\times$ fewer operations than large models. It is 18M parameters smaller and requires $1.7\times$ fewer operations than LightCap. In terms of caption quality, SwiFTeR exhibits 110.2 CIDEr-D, which is similar to Att2al [23], E2E-VLP [35], MiniMLM[31] and DistilVLM [6] while being much faster and smaller. On the downside, SwiFTeR is currently less descriptive than LightCap by 26.4 CIDEr-D. However, this is attributed to a much weaker training configuration in the backbone SwiFT_{weak} compared to LightCap. In the ideal training configuration (backbone trained on ImageNet-22k), Retention and Fourier Transform achieve the expected results of 125.5 (see "SwiFTeR ablated" in Table 2). To further support this claim, we note that the decrease is not to be attributed to the lack of parameters since EfficientVLM [7] achieves 120.8 using a backbone of a similar number of parameters. Therefore, the current results leave much room for improvement.

5.3 Efficiency

Despite the Retention layer contributing to a small fraction of the architecture, it plays a key role in decreasing the overall latency of the captioning model. In Table 2 we reported the FLOPs assuming an output sequence length of 128, to be processed in a single passage. However, in reality, captions are generated step by step, and the real computation is higher than what is typically reported for evaluation in literature. Additionally, FLOPs do not reflect the time cost which depends instead on synchronizations and parallel operations. For these reasons, the recurrent formulation of SwiFTeR, enabled by the Retention layer, cannot be fully appreciated in the previously discussed comparisons.

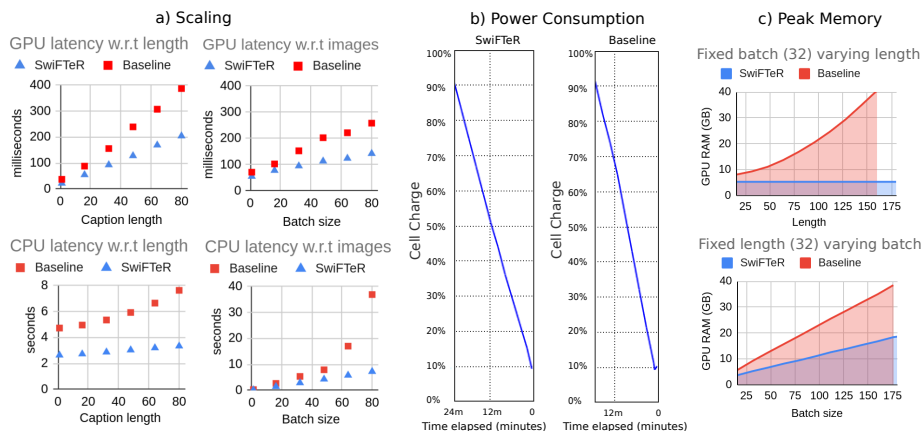


Fig. 3. Latency of SwiFTeR and baseline models with different length and batch size configurations. Unless differently stated, the length is set to 32, and the batch size is set to 32. The GPU consists of an NVIDIA GeForce RTX 4090 24GB. Each configuration is accommodated into one batch. Intel(R) Core i9-13900K is adopted as CPU.

To address this issue, we create a baseline architecture consisting of the Swin-Transf.-S + Transformer fusion model (the baseline) with the same hidden dimension as SwiFTeR’s fusion model ($H=96$). Such a baseline well represents the behavior of SotA models in Table 2 since they are based upon the same architecture.

Figure 3 highlights the time costs and several observations can be made from it. By looking at the GPU times (Figure 3-Top Row), it can be seen that SwiFTeR scales better with increased caption length and batch size. However, both latency plots appear to vary linearly to the input change. This behavior is due to the selected GPU, the NVIDIA GeForce RTX 4090 24GB, a high-performing Desktop GPU. Because of the low computational requirements of both SwiFTeR and the baseline, most of the time is spent on memory transfer

and synchronization, hence the linearity. The GPU was selected for its capacity to store all the experimental configurations in RAM, but it is not a good representative of limited resources devices, such as IoT and embedded systems. Instead, CPU times (Figure 3-Bottom Row) reflect the exact number of operations required by the two architectures in each configuration. Hence, it provides a better idea of their efficiency on devices with limited computational capabilities. The time cost on the CPU showcases the benefit of both the Retention-based fusion model and the SwiFT backbone. When the caption length increases, the final cost is dominated by the cost of the fusion decoder. As a result, the first design choice translates into a linear time increase since the cost of handling a recurrent state is constant. In contrast, the baseline architecture exhibits a polynomial increase with respect to the caption length (Figure 3-Bottom-Left). When the batch size or the number of images increases, the cost is dominated by the cost of the backbone. In this case, SwiFTeR latency showcases a linear behavior as expected. However, the baseline’s latency is not only significantly higher, but the high peak memory requirements led by the baseline, forced the RAM into paging, which explains the quadratic behavior.

Overall, running SwiFTeR on unoptimized PyTorch code² requires up to 180 ms on Intel(R) Core i9-13900K CPU, and about 54 ms on NVIDIA GeForce RTX 4090 GPU, to process a single image of size 224×244 and produce a caption of length 32, which is comparable to State-of-the-Art proposals such as DistilVLM [6], MiniVLM [31], and LightCap [32]. However, it showcases scalability to caption length and the smaller memory enables more images to be processed simultaneously, making it suitable for applications where many executions are needed frequently.

5.4 Power Consumption and Memory

While FLOPs in Table 2 do not reflect well the latency due to parallelism and synchronization aspects, it can be indicative of how much power is consumed to run the model. To test this aspect, we execute the Baseline architecture defined in Section 5.3 and SwiFTeR on Asus VivoBook equipped with NVIDIA GTX GeForce 1050 and Intel Core i7-7700HQ. In Figure 3-b) it appears that generating on GPU, descriptions of length 32 from batches of 8 images at 10Hz, the Baseline discharges the cell from 90% to 10% in 14 minutes and 37 seconds, which means it is $1.5 \times$ more power-consuming than the case of SwiFTeR, which does so in 22 minutes and 7 seconds.

Another advantage of our proposal lies in the peak memory usage, which can be seen in Figure 3-c). SwiFTeR keeps the memory peak constant regardless of the caption length. Whereas, the memory slope is lower in case of increased batch size, which enables more images to be processed simultaneously.

² PyTorch is not optimized for efficiency. Latency can be reduced significantly using TensorRT and CUDA implementations.

6 Conclusion and Future Works

In this work, we proposed a high-efficiency and lightweight Image Captioning model named SwiFTeR, characterized by Shifted Window Fourier Transform as the visual backbone and Retention in the fusion model. We assessed that Retention and Fourier Transform work well in the context of Image Captioning beyond the Language Modelling tasks they were originally designed for. Experimental results on the fusion model only showcased that the two methods led to negligible degradation in caption quality compared to the Transformer. Whereas, experimental results on the entire SwiFTeR architecture highlighted greater efficiency, compared to established architectures, thanks to superior scalability to the caption length and the image size. These aspects in addition to the lower power consumption, lower peak RAM usage, and small storage memory requirements make our proposal particularly suitable for resource-limited devices. Since the present paper focused mostly on designing a lightweight and fast architecture, future works will address the caption quality by designing a stronger visual backbone. This will be done in four steps: (1) training on ImageNet-22k, (2) exploring different training hyper-parameters configurations, (3) fine-tuning over data augmentation methods, and (4) pre-training on additional data beyond ImageNet, following the practice of the more established works [32, 7, 13].

7 Acknowledgments

This work has received funding from the European Union’s Horizon 2020 programme dAIEDGE (G.A. No 101120726).

References

1. Albertsson, V.G.: Analyzing descriptive captions for crime recognition in surveillance footage. B.S. thesis, University of Twente (2023)
2. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6077–6086 (2018)
3. Arystanbekov, B., Kuzdeuov, A., Nurgaliyev, S., Varol, H.A.: Image captioning for the visually impaired and blind: A recipe for low-resource languages. In: 2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). pp. 1–4. IEEE (2023)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
5. Dong, J., Chen, S., Miralinaghi, M., Chen, T., Li, P., Labi, S.: Why did the ai make that decision? towards an explainable artificial intelligence (xai) for autonomous driving systems. Transportation Research Part C: Emerging Technologies **156**, 104358 (2023)

6. Fang, Z., Wang, J., Hu, X., Wang, L., Yang, Y., Liu, Z.: Compressing visual-linguistic model via knowledge distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1428–1438 (2021)
7. Fei, Z.: Efficient modeling of future context for image captioning. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 5026–5035 (2022)
8. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
9. Hu, J., Cavicchioli, R., Capotondi, A.: A request for clarity over the end of sequence token in the self-critical sequence training". Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) **14233 LNCS**, 39 – 50 (2023). https://doi.org/10.1007/978-3-031-43148-7_4
10. Hu, J.C., Cavicchioli, R., Berardinelli, G., Capotondi, A.: Heterogeneous encoders scaling in the transformer for neural machine translation. In: Proceedings of the Seventh Workshop on Natural Language for Artificial Intelligence (NL4AI 2023) co-located with 22th International Conference of the Italian Association for Artificial Intelligence (AI* IA 2023) (2023)
11. Hu, J.C., Cavicchioli, R., Berardinelli, G., Capotondi, A.: Sharebert: Embeddings are capable of learning hidden layers. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 18225–18233 (2024)
12. Hu, J.C., Cavicchioli, R., Capotondi, A.: Exploiting multiple sequence lengths in fast end to end training for image captioning. In: 2023 IEEE International Conference on Big Data (BigData). pp. 2173–2182. IEEE Computer Society (2023)
13. Hu, X., Gan, Z., Wang, J., Yang, Z., Liu, Z., Lu, Y., Wang, L.: Scaling up vision-language pre-training for image captioning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17980–17989 (2022)
14. Huang, L., Wang, W., Chen, J., Wei, X.Y.: Attention on attention for image captioning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4634–4643 (2019)
15. Kong, Z., Dong, P., Ma, X., Meng, X., Niu, W., Sun, M., Shen, X., Yuan, G., Ren, B., Tang, H., et al.: Spvit: Enabling faster vision transformers via latency-aware soft token pruning. In: European Conference on Computer Vision. pp. 620–640. Springer (2022)
16. Lee-Thorp, J., Ainslie, J., Eckstein, I., Ontanon, S.: Fnet: Mixing tokens with fourier transforms. arXiv preprint arXiv:2105.03824 (2021)
17. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. arXiv preprint arXiv:2201.12086 (2022)
18. Li, W., Gao, C., Niu, G., Xiao, X., Liu, H., Liu, J., Wu, H., Wang, H.: Unimo: Towards unified-modal understanding and generation via cross-modal contrastive learning. arXiv preprint arXiv:2012.15409 (2020)
19. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
20. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
21. Nag, S., Datta, G., Kundu, S., Chandrachoodan, N., Beerel, P.A.: Vita: A vision transformer inference accelerator for edge applications. arXiv preprint arXiv:2302.09108 (2023)

22. Peng, B., Alcaide, E., Anthony, Q., Albalak, A., Arcadinho, S., Biderman, S., Cao, H., Cheng, X., Chung, M., Grella, M., GV, K.K., He, X., Hou, H., Lin, J., Kazienko, P., Kocon, J., Kong, J., Koptyra, B., Lau, H., Mantri, K.S.I., Mom, F., Saito, A., Song, G., Tang, X., Wang, B., Wind, J.S., Wozniak, S., Zhang, R., Zhang, Z., Zhao, Q., Zhou, P., Zhou, Q., Zhu, J., Zhu, R.J.: Rwkv: Reinventing rnns for the transformer era (2023)
23. Rennie, S.J., Marcheret, E., Mroueh, Y., Ross, J., Goel, V.: Self-critical sequence training for image captioning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7008–7024 (2017)
24. Shao, Z., Yu, Z., Wang, M., Yu, J.: Prompting large language models with answer heuristics for knowledge-based visual question answering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14974–14983 (2023)
25. 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. vol. 34, pp. 8815–8821 (2020)
26. Sun, Y., Dong, L., Huang, S., Ma, S., Xia, Y., Xue, J., Wang, J., Wei, F.: Retentive network: A successor to transformer for large language models. arXiv preprint arXiv:2307.08621 (2023)
27. Tan, J.H., Chan, C.S., Chuah, J.H.: Comic: Toward a compact image captioning model with attention. *IEEE Transactions on Multimedia* **21**(10), 2686–2696 (2019)
28. Tan, J.H., Chan, C.S., Chuah, J.H.: Image captioning with sparse recurrent neural network. arXiv preprint arXiv:1908.10797 (2019)
29. Tan, J.H., Tan, Y.H., Chan, C.S., Chuah, J.H.: Acort: A compact object relation transformer for parameter efficient image captioning. *Neurocomputing* **482**, 60–72 (2022)
30. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
31. Wang, J., Hu, X., Zhang, P., Li, X., Wang, L., Zhang, L., Gao, J., Liu, Z.: Minivlm: A smaller and faster vision-language model. arXiv preprint arXiv:2012.06946 (2020)
32. Wang, N., Xie, J., Luo, H., Cheng, Q., Wu, J., Jia, M., Li, L.: Efficient image captioning for edge devices. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 2608–2616 (2023)
33. Wang, X., Zhang, L.L., Wang, Y., Yang, M.: Towards efficient vision transformer inference: A first study of transformers on mobile devices. In: Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications. pp. 1–7 (2022)
34. Wang, Y., Xu, J., Sun, Y.: End-to-end transformer based model for image captioning. arXiv preprint arXiv:2203.15350 (2022)
35. Xu, H., Yan, M., Li, C., Bi, B., Huang, S., Xiao, W., Huang, F.: E2e-vlp: End-to-end vision-language pre-training enhanced by visual learning (2021)
36. Zafrir, O., Boudoukh, G., Izsak, P., Wasserblat, M.: Q8bert: Quantized 8bit bert. In: 2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS). pp. 36–39. IEEE (2019)
37. Zheng, C., Zhang, K., Yang, Z., Tan, W., Xiao, J., Ren, Y., Pu, S., et al.: Savit: Structure-aware vision transformer pruning via collaborative optimization. *Advances in Neural Information Processing Systems* **35**, 9010–9023 (2022)
38. Zhou, Y., Hu, Z., Liu, D., Ben, H., Wang, M.: Compact bidirectional transformer for image captioning. arXiv preprint arXiv:2201.01984 (2022)