# STAMP: SEQUENCE TRANSFORMATION AND MIXED PRECISION FOR LOW-PRECISION ACTIVATION QUANTIZATION

### **Anonymous authors**

Paper under double-blind review

### **ABSTRACT**

Quantization is the key method for reducing inference latency, power and memory footprint of generative AI models. However, accuracy often degrades sharply when activations are quantized below eight bits. Recent work suggests that invertible linear transformations (e.g. rotations) can aid quantization, by reparameterizing feature channels and weights. In this paper, we propose *Sequence Transformation and Mixed Precision* (STaMP) quantization, a novel strategy that applies linear transformations along the *sequence* dimension to exploit the strong local correlation in language and visual data. By keeping a small number of tokens in each intermediate activation at higher precision, we can maintain model accuracy at lower (average) activations bit-widths. We evaluate STaMP on recent LVM and LLM architectures, demonstrating that it significantly improves low bit width activation quantization and complements established activation and weight quantization methods including recent feature transformations.

### 1 Introduction

Modern generative models, such as large language models (LLMs) and large vision models (LVMs), achieve state-of-the-art performance in text and image generation but at the cost of massive computational and memory requirements. As the demand for scalable and efficient deployment of these models grows, both in the cloud and on edge devices where computational resources are scarce, achieving inference efficiency has become a critical area of research.

Post-Training Quantization (PTQ) of weights and activations is fundamental to enhancing inference efficiency, especially for demanding operations such as large matrix multiplications in linear layers, which dominate power consumption. However, PTQ faces substantial challenges when pushing the quantization bitwidth down to 4-bits, often due to the presence of outliers in the weight and activation tensors.

To mitigate outliers, prior work applies function-preserving transformations (van Breugel et al., 2025) to weights and activations. For example, Xiao et al. (2023) scales down

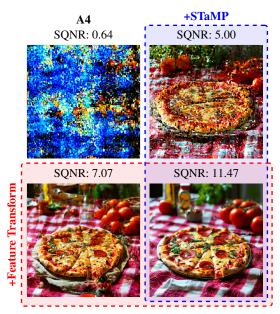


Figure 1: STaMP and feature transformations applied to  $PixArt-\Sigma$  with 4-bit activations. The benefit of STaMP is orthogonal to (Hadamard) feature transformation, drastically reducing artifacts.

outlier activations and compensates by scaling up subsequent weights, preserving the model output. Similarly, Hadamard-based feature mixing (Ashkboos et al., 2024; Liu et al., 2024; Ma et al., 2024; Zhao et al., 2025) reduces activation variance and spreads outliers across dimensions.

However, these methods operate primarily along the feature dimension and ignore correlations across the sequence dimension. Drawing inspiration from traditional media compression methods, in this paper we propose *Sequence Transformation and Mixed Precision* (STaMP) quantization, a complementary approach that leverages sequence structure to improve activation quantization.

Our main contributions are summarized below:

- We introduce a new class of activation transformations operating along the sequence dimension, complementary to existing feature transformations (Figure 1).
- We characterize the quantization error for sequence transforms and design a novel mixedprecision quantization scheme to exploit local activation correlation.
- We demonstrate that STaMP consistently improves the model accuracy when combined with feature transformations and weight quantization on both LLM and LVM.

# 2 BACKGROUND

### 2.1 ACTIVATION QUANTIZATION

Consider  $X \in \mathbb{R}^{s \times d}$  as an activation matrix of shape sequence length  $s \times feature$  size d (batch size is omitted for the sake of clarity). Integer quantization refers to the operation of converting the activations to integer values with lower bit width. The activation quantization operation  $\mathbf{Q} : \mathbb{R}^{s \times d} \to \mathbb{N}^{s \times d}$  is defined as:

$$x_{ij}^{\text{int}} \stackrel{def}{=} \mathbf{Q}(\mathbf{X})_{ij} = \text{clamp}\left(\left\lfloor \frac{x_{ij}}{s_{ij}} \right\rfloor + z_{ij}, 0, 2^{b_{ij}} - 1\right),\tag{1}$$

in which  $z_{ij}$  and  $s_{ij}$  can be interpreted as an offset and scaling parameter, while  $b_{ij}$  refers to the number of bits used to quantize the entry  $x_{ij}$ . In order to make quantization efficient, offsets, scales and bit widths are shared across all feature channels:  $s_{ij} = s_i$ ,  $z_{ij} = z_i$ , and  $b_{ij} = b_i$ .

The de-quantization function  $\mathbf{Q}^{-1}: \mathbb{N}^{s \times d} \to \mathbb{R}^{s \times d}$  maps the integer-quantized activation into the original domain:  $\mathbf{Q}^{-1}(\boldsymbol{X}^{\text{int}})_{ij} = (x_{ij}^{\text{int}} - z_{ij})s_{ij}$ . We will refer to the combination of the quantized activation into the

tization and de-quantization operation with  $\mathcal{Q}(\boldsymbol{X}) \stackrel{def}{=} \mathbf{Q}^{-1}(\mathbf{Q}(\boldsymbol{X}))$ , omitting the quantization parameters for brevity. The expected activation quantization error introduced by the quantization and de-quantization operations is commonly defined as:

$$\mathcal{L}(\boldsymbol{X}) \stackrel{def}{=} \mathbb{E}\left[\|\mathcal{Q}(\boldsymbol{X}) - \boldsymbol{X}\|_{2}^{2}\right] = \sum_{i=1}^{s} \mathbb{E}\left[\|\mathcal{Q}(\boldsymbol{x}_{i}) - \boldsymbol{x}_{i}\|_{2}^{2}\right], \tag{2}$$

in which the expectation is computed with respect to the activation distribution p(X) and  $\|\cdot\|_2^2$  refers to the squared Frobenius norm. Nagel et al. (2021) identified two main causes of quantization error: clipping error (due to the clamp operator) and rounding error (introduced by the rounding). As is common practice in the literature, we will focus on a setting in which the scale  $s_i$  and offset  $z_i$  for the i-th token are set based on the range  $(x_i) \stackrel{def}{=} \max_j x_{ij} - \min_j x_{ij}$  to prevent any clipping error:  $\overline{s_i} \stackrel{def}{=} \frac{2^{b_i}-1}{\operatorname{range}(x_i)}, \ \overline{z_i} \stackrel{def}{=} -\frac{\min_j x_{ij}}{\overline{s_i}}$ . In this setting, the quantization error for each token  $x_i$  is determined by its quantization scale  $\overline{s_i}$ :

$$\mathcal{L}(\boldsymbol{x}_i) \le \frac{d}{4} \mathbb{E}\left[\overline{s_i}^2\right] = \frac{d}{4} \frac{\mathbb{E}\left[\text{range}\left(\boldsymbol{x}_i\right)^2\right]}{(2^{b_i} - 1)^2}.$$
 (3)

As the activation quantization error increases, so does the deterioration in performance for the model outputs since the network output diverge from the original (unquantized) model.

### 2.2 FEATURE TRANSFORMATIONS

In order to reduce the activation quantization error, recent literature has introduced linear functionpreserving transformations in the form of (right) invertible matrices R, which are applied prior to the quantization operation (van Breugel et al., 2025). Clearly, the additional flexibility introduced by **Feature Transformations** can aid in reducing the activation quantization error:

$$\mathcal{L}(X) \ge \min_{R} \underbrace{\mathbb{E}\left[\left\|\mathcal{Q}(XR)R^{-1} - X\right\|_{2}^{2}\right]}_{\mathcal{L}(X;R)}.$$
(4)

In particular, rotation matrices can effectively reduce the activation range, effectively spreading outliers across multiple channels and consequently reducing the subsequent quantization error:

$$\sum_{i=1}^{s} \mathbb{E}\left[\operatorname{range}\left(\boldsymbol{x}_{i}\boldsymbol{R}\right)^{2}\right] \leq \sum_{i=1}^{s} \mathbb{E}\left[\operatorname{range}\left(\boldsymbol{x}_{i}\right)^{2}\right] \implies \mathcal{L}\left(\boldsymbol{X};\boldsymbol{R}\right) \leq \mathcal{L}\left(\boldsymbol{X}\right) \tag{5}$$

Hadamard matrices have proven very effective in reducing the number of outliers with limited additional overhead: matrix multiplications with Hadamard matrices can be performed efficiency in  $O(sd\log d)$  thanks to the butterfly algorithm (Fino & Algazi, 1976) , and the inverse Hadamard matrix can be fused into linear layer weights (Ashkboos et al., 2024).

Existing feature transformation techniques reduce activation quantization error by redistributing activation ranges across features. However, they operate exclusively along the feature dimension and ignore correlations across the sequence dimension. Visual and textual data exhibit strong local correlations—neighboring pixels in images and adjacent tokens in text are strongly dependent. This suggests that a similar structure could exist in the model intermediate activations and could be leveraged to improve quantization efficiency, which is described in the next section.

# 3 METHOD

We define a **Sequence Transform** as a linear transformation of X across sequence dimension defined by a (left) invertible matrix L. Similarly to feature transformations R, sequence transformations can reduce the quantization error, and the two can be easily combined to further the error:

$$\mathcal{L}\left(\boldsymbol{X}\right) \ge \min_{\boldsymbol{L}} \underbrace{\mathbb{E}\left[\left\|\boldsymbol{L}^{-1}\mathcal{Q}\left(\boldsymbol{L}\boldsymbol{X}\right) - \boldsymbol{X}\right\|_{2}^{2}\right]}_{\mathcal{L}\left(\boldsymbol{X};\boldsymbol{L}\right)} \ge \min_{\boldsymbol{L},\boldsymbol{R}} \underbrace{\mathbb{E}\left[\left\|\boldsymbol{L}^{-1}\mathcal{Q}\left(\boldsymbol{L}\boldsymbol{X}\boldsymbol{R}\right)\boldsymbol{R}^{-1} - \boldsymbol{X}\right\|_{2}^{2}\right]}_{\mathcal{L}\left(\boldsymbol{X};\boldsymbol{L},\boldsymbol{R}\right)}. \tag{6}$$

Since sequence transformations are linear, they commute with other linear operations. In particular, for a linear layer:

$$\left(\boldsymbol{L}^{-1}\mathcal{Q}\left(\boldsymbol{L}\boldsymbol{X}\right)\right)\boldsymbol{W}+\boldsymbol{1}\boldsymbol{\beta}^{T}=\boldsymbol{L}^{-1}\left(\mathcal{Q}\left(\boldsymbol{L}\boldsymbol{X}\right)\boldsymbol{W}\right)+\boldsymbol{1}\boldsymbol{\beta}^{T}=\boldsymbol{L}^{-1}\left(\mathcal{Q}\left(\boldsymbol{L}\boldsymbol{X}\right)\boldsymbol{W}+\underbrace{\left(\boldsymbol{L}\boldsymbol{1}\right)}_{\boldsymbol{\ell}}\boldsymbol{\beta}^{T}\right),\tag{7}$$

in which 1 represents a vector of ones of size s, indicating that the same bias  $\beta$  is applied to all the tokens. This implies that we can invert sequence transformation (i) right before applying the bias in a linear layer, or (ii) postpone this operation and use a sequence-transformed bias  $\ell\beta^T$  in which the scale  $\ell_i$  may differ across different tokens as a function of L. The algorithm for a sequence transformed linear layer is reported in Figure 2a.

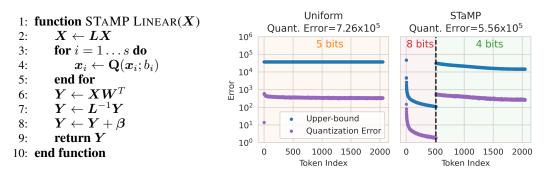
We emphasize that, contrary to feature transformations, sequence transforms do not affect weights, and hence they are orthogonal to more advanced weight quantization methods such as vector quantization, GPTQ (Guo et al., 2024), and SVDQuant (Li et al., 2025).

### 3.1 SEQUENCE TRANSFORM AND MIXED PRECISION (STAMP)

To understand how sequence transformations affect quantization error, we first formalize the relationship between quantization error and sequence-transformed tokens (proof in Appendix A.1):

**Theorem 1.** The expected quantization error for activations X transformed by an orthogonal sequence transformation L and quantized using a min-max scale for each token is upper-bounded by the weighted sum of the expected norm of the transformed tokens:

$$\mathcal{L}(\boldsymbol{X}; \boldsymbol{L}) \leq \frac{d}{2} \sum_{i=1}^{s} \frac{\mathbb{E}\left[\left\|\boldsymbol{l}_{i}^{T} \boldsymbol{X}\right\|_{2}^{2}\right]}{(2^{b_{i}} - 1)^{2}}.$$
(8)



(a) STaMP Linear Layer Pseudocode

(b) Comparison of Upper-Bound and Quantization error

Figure 2: **Summary of the STaMP Procedure**. The sequence Transform L aims to concentrate the energy in the initial tokens, which are quantized at higher precision. This reduces the value of the upper-bound in Equation 8 (blue) and, consequently, the overall activation quantization error (purple). For a fixed average bit width of 5 bits, combining energy concentration with two precision levels (2b, right) achieves lower error than a uniform quantization scheme without sequence transformations (2b, left). Activations are collected from the input to Layer 20 of the LLaMA v3 8B model.

Note that sequence transformations  $\boldsymbol{L}$  do not alter the *total energy*  $E = \sum_{i=1}^{s} e_i$ , so any improvement to the bound in Equation 8 must come from redistributing the bit width. This observation motivates a mixed precision strategy: instead of keeping energy and bit width uniform, we deliberately concentrate most of the energy into a few tokens and allocate more bits to them. Because the denominator in Equation 8 grows exponentially with  $b_i$ , allocating extra bits to tokens with large energy yields a disproportionately large reduction in their contribution to the error. In other words, redistributing a bit from a low-energy token to a high-energy token reduces the total error more than keeping bit widths and energy uniform. This property is illustrated in Figure 2b and further elaborated in Appendix A.3.

Therefore, to improve activation quantization performance, we propose **Sequence Transform and Mixed Precision** (STaMP) a simple yet effective strategy that concentrates activation energy into a small set of tokens and assigns them higher precision. In the next section, we describe how to design a transformation  $\boldsymbol{L}$  that achieves this efficiently and how to determine the corresponding bit width allocation.

### 3.2 EFFICIENT ENERGY CONCENTRATION

The energy of the i-th sequence-transformed token  $e_i$  can be also seen as the projection of the autocorrelation matrix  $S = \mathbb{E}[XX^T]$  along the direction  $l_i$ :

$$e_i = \mathbb{E}\left[\left\|\boldsymbol{l}_i \boldsymbol{X}\right\|_2^2\right] = \boldsymbol{l}_i^T \boldsymbol{S} \boldsymbol{l}_i.$$
 (9)

he direction that maximizes this energy is the eigenvector  $u_1$  associated with the largest eigenvalue  $\lambda_1$  of S. Similarly, the second largest energy corresponds to  $u_2$  and so on. Therefore, given the eigendecomposition  $S = U\Lambda U^T$ , the optimal orthogonal transformation L for concentrating the token norm is  $U^T$ . In this case, the energy e of the transformed tokens LX aligns with the squared eigenvalues  $\lambda^2$ . This linear transformation is also known as the **Karhunen-Loève Transform** (KLT), which requires a representative calibration set to estimate U. Despite its optimality, KLT has the same computational complexity of a full-rank matrix multiplication, which is impractical since each transform needs to be applied twice for each linear layer. Estimating S for each activation would further adds a costly calibration step.

Fortunately, the autocorrelation matrix of common LVM and LLM activations S exhibits a strong structure induced by the properties of natural images and text. Figure 3a shows that tokens corresponding to spatially or sequentially adjacent activations are highly correlated, while distant tokens are weakly correlated. As a result S is approximately (block) Toeplitz, whose eigenvectors can be

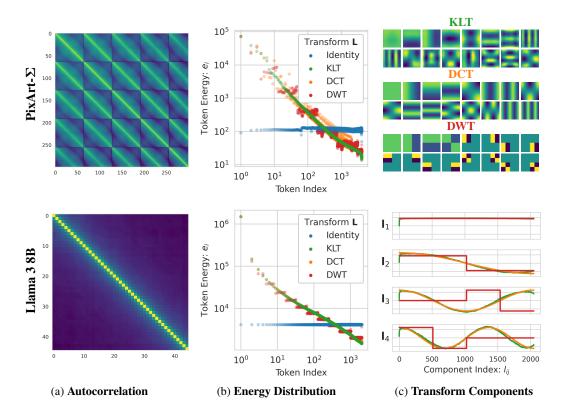


Figure 3: Visualization of (3a) a portion of the autocorrelation matrix, (3b) transformed token energy distribution, and (3c) transformation components for the input to the cross-attention layer 15 of PixArt- $\Sigma$  and the attention layer 20 of LLaMA-v3-8B, computed on COCO and Wikitext, respectively. The block structure in the LVM activations arises from flattening 2D data into a 1D sequence. Both matrices exhibit a Toeplitz-like diagonal structure, allowing their KLT eigenbases to be efficiently approximated by DCT (Figure 3c), which concentrates the token energy close to optimal distribution (Figure 3b). The DWT closely approximates the optimal energy with discrete levels.

well-approximated by a Fourier basis<sup>1</sup>. In particular, since the autocorrelation is real and symmetric, we can use a **Discrete Cosine Transform** (DCT) instead of a complex Fourier basis.

The complexity of DCT  $\mathcal{O}(ds\log s)$  is lower than a full matrix multiplication since the transformation requires only  $\log s$  steps of the Fast Fourier Transform algorithm. Further simplification is possible by retaining only the sign of Fourier coefficients, yielding the **Walsh-Hadamard Transform** (WHT), which approximates DCT while enabling more efficient hardware implementations.

Finally, the **Discrete Wavelet Transform**  $(DWT)^2$  further reduces the computational complexity to  $\mathcal{O}(ds)$  while effectively concentrating the energy of the activations at each intermediate step. Each DWT step pushes the energy in the first half (one quarter for 2D signal) of the tokens, requiring  $\log s$  steps to fully concentrate the energy.

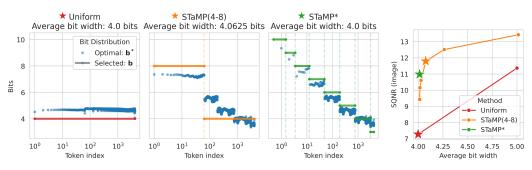
Figure 3 compares energy distributions (3b) and basis components (3c) for KLT, DCT, and DWT on intermediate activations of LVM and LLM architectures.

### 3.3 OPTIMAL BIT ALLOCATION

Given a vector of energies e, the optimal bit distribution b for a total maximum allocation of B bits follows the logarithm of the squared root of the token energy:  $b_i^* = \log_2 \sqrt{e_i} - C$ , with  $C = \left(B - \sum_{i=1}^s \log_2 \sqrt{e_i}\right)/s$ .  $b^*$  gives us an indication of the optimal bit width for each token, however, in practice we are restricted to integer bit widths  $|b^*|$ . Furthermore, due to practical limitations, it

<sup>&</sup>lt;sup>1</sup>This follows from Szegő's theorem.

<sup>&</sup>lt;sup>2</sup>We use the Haar wavelet for its simplicity and minimal padding requirements.



(a) Bit width allocation and transformation strategies

(b) Bit width SQNR tradeoff

Figure 4: Comparison of energy and bit width allocation strategies. Figure 4a compares a uniform allocation without transformation against two STaMP strategies based on the DWT transform. Restricting to only 4- and 8-bits precision (in yellow) introduces minimal overhead while significantly improving output quality. The results in Figure 4b are obtained by varying the number of high-precision tokens and adjusting the uniform quantization bit width for a per token-activation-quantized  $PixArt-\Sigma$  model with QuaRot feature transformations.

is beneficial to use only a small number of different bit precisions which are supported on hardware, such as 4 or 8 bits.

For this reason, although the DWT energy concentration is sub-optimal, its property of creating a discrete number of energy levels makes it more suitable to our use case. As illustrated in Figure 4a in yellow, we propose a simple allocation scheme that uses only two bit widths: the first 64 tokens are kept at 8 bits, while the rest uses 4 bits, resulting only in a minor bit width overhead (4.0625 on the PixArt- $\Sigma$ ), while significantly improving the model accuracy (Figure 4b). For this reason, STaMP with DWT and 2 precision level will be our main focus in the experimental section.

### 4 RELATED WORK

Quantization is a fundamental technique for reducing the computational and memory footprint of deep neural networks, enabling efficient inference with minimal accuracy loss. The rapid growth of LLMs and LVMs has intensified interest in post training quantization (PTQ) Nagel et al. (2021); Gholami et al. (2022), as retraining these models is often impractical. Recent PTQ approaches focus on removing the outliers and reducing the dynamic range of weights and activations, to improve quantization robustness.

SmoothQuant (Xiao et al., 2023) reduces activations outliers by scaling the feature channels, shifting quantization difficulty from activations to weights. *QuaRot* and *FlatQuant* apply invertible feature transforms over weights and activations to spread outliers across channels, employing randomized Hadamard (Ashkboos et al., 2024) or learning lightweight affine transforms (Sun et al., 2025). van Breugel et al. (2025) develop transforms that commute with transformer operations, such that they can be merged into linear weights. Zhao et al. (2025) develops a Static-Dynamic Channel Balancing (SDCB) procedure based on channel scaling and mixing on Diffusion Transformer (DiT) architectures and retaining certain quantization-sensitive layers to higher bitwidth to achieve W8A8. Li et al. (2025) absorbs activation outliers into a high-precision low-rank branch via singular value decomposition, while quantizing the residuals to 4-bit with a per-token/per-group quantization granularity to achieve W4A4 mixed-precision quantization on DiT. Federici et al. (2025) reduces the dynamic range of DiT activations by subtracting the sequence average from each token and applying Hadamard feature rotations, at the cost of processing an extra token.

While these approaches have advanced PTQ for large models, they operate exclusively along the feature dimension and ignore correlations across the sequence dimension. In contrast, STaMP introduces sequence-aware transformations that exploit local token correlations by using a mixed precision allocation strategy under a fixed average budget. Our approach is strongly relates to classical media compression techniques that leverage frequency-domain transforms to concentrate energy and

Table 1: **STaMP consistently improves LVM quantization.** Image SQNR and Image Reward (IR) for W4A4 per block quantization with block size 64. For all the STaMP results we keep 64 tokens at 8-bits. STaMP consistently improves baselines.

		COCO				MJHQ			
		SQNR		IR		SQNR		IR	
	$\textbf{STaMP} \rightarrow$	X	1	X	<b>√</b>	X	1	X	<b>√</b>
	FP	$+\infty$		0.90		+∞		0.96	
PixArt- $\Sigma$	RTN ViDiT-Q SVDQuant	5.88 7.82 8.78	6.16 6.36 9.72	0.38 0.83 0.90	0.80 0.84 0.91	5.75 7.55 8.83	6.23 8.53 9.75	0.38 0.76 0.86	0.76 1.37 0.89
	FP	+	$\infty$	0.	87	+∞		+∞	
SANA	RTN ViDiT-Q SVDQuant	8.63 10.03 9.99	9.32 10.74 10.69	0.89 0.89 0.87	0.91 0.86 0.90	8.56 10.04 9.88	9.40 10.69 10.51	0.95 0.96 0.93	0.99 0.97 0.98

enable adaptive quantization. JPEG (Leger et al., 1991) and JPEG2000 (Christopoulos et al., 2002) rely on the Discrete Cosine Transform and Discrete Wavelet Transform, respectively, to decorrelate spatial data and allocate bits based on perceptual importance. Similar principles underlie modern video coding standards such AVC Wiegand et al. (2003) and HEVC (Sullivan et al., 2012), as well as audio codecs like MP3 and AAC, which use the Modified Discrete Cosine Transform (Princen et al., 1987). STaMP demonstrates that it is possible to apply similar principles to the activation space of generative models, where local correlation in the sequence dimension allows energy compaction and selective precision allocation, improving quantization efficiency without retraining.

# 5 EXPERIMENTAL RESULTS

We demonstrate the effectiveness of STaMP on both LVM and LLM transformer models, comparing against uniform activation quantization, and show its effectiveness in combination with popular feature transformation and weight quantization approaches. Throughout this section, we mostly focus on STaMP with DWT due to its computational efficiency. Additional results and visualizations are reported in Appendix C.

# 5.1 VISION LANGUAGE MODELS

**Set-up** Our LVM experiments focus on the Diffusion Transformer architecture of PixArt- $\Sigma$  (Chen et al., 2024) and SANA (Xie et al., 2025). We quantize activation before each linear layer in transformer blocks using asymmetric quantization with minmax scaling. As common practice, unless otherwise specified, we use a separate scale  $s_i$  for each token and weight output channel. Following common procedure in literature, in LVMs we do not quantize activations and weights corresponding to the cross-attention *key* and *value* since their effect accounts for less than 5% of the runtime (Li et al., 2025).

We evaluate the fidelity of the quantized LVMs by computing the Signal to Quantized Noise Ratio  $SQNR(\boldsymbol{O}_{orig}, \boldsymbol{O}_{quant}) = 10\log_{10}\left(\|\boldsymbol{O}_{orig}\|_2^2/\|\boldsymbol{O}_{orig}-\boldsymbol{O}_{quant}\|_2^2\right)$  both in the diffusion latent space and image space. We further compute CLIP Score (Hessel et al., 2021) to assess alignment with the textual prompt, CLIP IQA (Wang et al., 2023), and Image Reward (Xu et al., 2023) to evaluate the generated image quality. Following standard procedures, we compute the metrics using 1000 prompts and images from the COCO (Lin et al., 2014) and MJHQ (Li et al., 2024) datasets.

We demonstrate the effectiveness STaMP alone and by combining it with other quantization methods developed in recent literature. We combine STaMP with feature transform methods such as SmoothQuant (Xiao et al., 2023) and QuaRot (Ashkboos et al., 2024), and Static-Dynamic Channel Balancing (SDCB) described in ViDiT-Q (Zhao et al., 2025). We further demonstrate the effectiveness of STaMP with the mixed precision low-rank weight quantization method described in

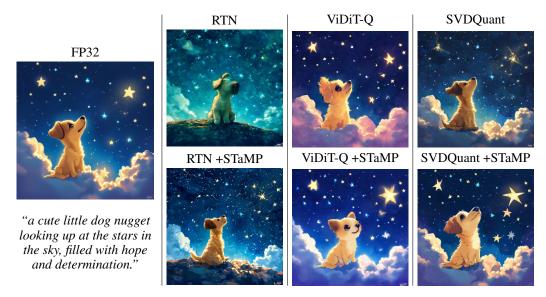


Figure 5: Visualization of PixArt- $\Sigma$  sample generation for the results reported in Table 1.

SVDQuant (Li et al., 2025). We apply STaMP before each linear layer in the Transformer blocks, inverting it right after each linear layer. We use 2-dimensional DWT with 64 8-bit (high precision) tokens unless otherwise specified.

**Results** Table 1 reports the results obtained by combining STaMP with recent LVM quantization method using the same quantization setting for all the baselines. Both activation and weights are quantized at 4 bits with blocks of size 64, following the setup described in (Li et al., 2025). We observe that STaMP consistently improves upon all the reported metrics on different models and architectures resulting in visually more accurate generations, as shown in Figure 5.

### 5.2 Number of high precision tokens

Figure 4b demonstrates the trade-off between bit width and SQNR in image space observed by changing the number of high-precision tokens in STaMP, while fixing the high and low precision bit widths to 8 and 4 bits respectively. Only activations are quantized to focus the analysis solely on the error introduced by per token activation quantization. We observe a sharp increase in SQNR whenever even a few high precision tokens are introduced. We additionally note that even at 5 bits the STaMP scheme achieves better performance than a uniform quantization scheme despite quantizing three quarters of the tokens at 4 bits. Additional comparison with per-block activation quantization schemes can be found in Appendix C.

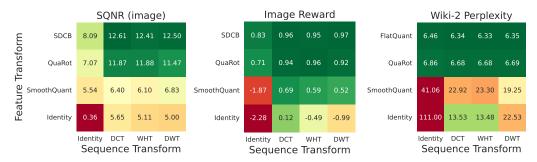
# 5.3 Large Language Models

**Set-up** We evaluate on language models of different sizes and model classes, including Llama 3 8B (Grattafiori et al., 2024), Llama 3.2 1B and 3B instruct, and Qwen 2.5 3B instruct (Qwen et al., 2025). We take popular *feature* transforms from literature, including SmoothQuant (Xiao et al., 2023), QuaRot (Ashkboos et al., 2024), and the state-of-the-art FlatQuant (Sun et al., 2025), and evaluate whether STaMP (DWT) brings additional gains to these baselines. We use the exact same quantization set-up as used in (Ashkboos et al., 2024; Sun et al., 2025), W4A4KV4 per token activation quantization, and like them, evaluate Wikitext-2 (Merity et al., 2017) perplexity at sequence length 2048. We use round-to-nearest (RTN) for weight quantization (vs. GPTQ (Frantar et al., 2022), as weight quantization is completely perpendicular to sequence transforms. We keep the first 64 tokens in 8-bits for all baselines, which means all methods use an effective activation/KV bit width of 4.125 bits. See Appendix B.2 for additional details.

**Results** We observe (Table 2) that all baselines improve consistently when STaMP is added. This is especially apparent for scenarios where baselines are far from FP performance—e.g. for the small

Table 2: **STaMP always improves LLM quantization.** Perplexity (PPL) for W4A4KV4 quantization, using the same setting as (Ashkboos et al., 2024; Sun et al., 2025). We use 64 8-bit tokens for activations and KV-cache for all methods and baselines, even if we do not apply the sequence transform (effectively W4A4.125KV4.125). The STaMP sequence transform improves all baselines.

	Llama 3 8B		Llama 3.2 1B it		Llama 3.2 3B it		Qwen 2.5 3B it	
$\textbf{STaMP} \rightarrow$	X	<b>√</b>	X	✓	X	<b>✓</b>	X	✓
FP	6.14		13.16		11.27		8.56	
RTN	668	95.3	1795	700	483	159	99723	18767
SmoothQuant	531	93.8	883	407	177	88.5	66929	29063
Quarot	9.05	8.66	25.78	23.72	18.43	17.57	94.86	71.13
FlatQuant	6.89	6.77	15.72	15.16	12.71	12.40	9.29	9.19



(a) A4 PixArt-Σ SQNR and Image Reward

(b) A4 Llama v3 8B Perplexity

Figure 6: Effect of combining Feature Transforms (rows) and STaMP (columns) with A4 on the  $PixArt-\Sigma$  and LLama v3 8B models.

Llama 3.2 1B and 3B instruct models, which are evidently hard to quantize at 4 bit. This demonstrates that STaMP is not a competitor of other quantization techniques (e.g. feature transforms), but an additional tool for achieving extremely low bit width quantization—which can be added without any manual tuning or training to existing quantization methods.

### 5.4 COMBINING FEATURE AND SEQUENCE TRANSFORMATIONS

In Figure 6 we assess the effectivenes of STaMP for three different sequence transformations in combination with popular feature transforms. Overall, we observe that improvements are largely complementary, especially for LVMs. Furthermore, results show that DCT, WHT and DWT perform similarly to each other, demonstrating that there is little price to pay for switching from the more accurate and computationally expensive DCT to the cheaper approximate DWT.

### 6 Conclusion

This work explores uncharted territories in the recent advancements of LLM and LVM quantization, applying invertible transformations over the sequence dimension to further reduce quantization error. Drawing inspiration from traditional signal processing, we introduce a novel method that exploits the autocorrelation structure that is naturally present in the intermediate activations of generative models to enable a more efficient mixed-precision activation quantization scheme.

We demonstrate that STaMP complements existing PTQ techniques such as *SmoothQuant* and *QuaRot*, and plays a critical role in advancing low-precision activation quantization, pushing the boundaries of W4A4 quantization for both LLMs and LVMs and offering training-free solution for deploying high-performance models in resource-constrained environments.

# REFERENCES

- Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L. Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefler, and James Hensman. Quarot: Outlierfree 4-bit inference in rotated llms. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 15, 2024, 2024. URL http://papers.nips.cc/paper\_files/paper/2024/hash/b5b939436789f76f08b9d0da5e8laf7c-Abstract-Conference.html.
- Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart-Σ: Weak-to-strong training of diffusion transformer for 4k text-to-image generation. In Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), Computer Vision ECCV 2024 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XXXII, volume 15090 of Lecture Notes in Computer Science, pp. 74–91. Springer, 2024. doi: 10.1007/978-3-031-73411-3\\_5. URL https://doi.org/10.1007/978-3-031-73411-3\_5.
- Charilaos Christopoulos, Athanassios Skodras, and Touradj Ebrahimi. The jpeg2000 still image coding system: an overview. *IEEE transactions on consumer electronics*, 46(4):1103–1127, 2002.
- Marco Federici, Riccardo Del Chiaro, Boris van Breugel, Paul Whatmough, and Markus Nagel. Hadanorm: Diffusion transformer quantization through mean-centered transformations. *arXiv* preprint arXiv:2506.09932, 2025.
- Fino and Algazi. Unified matrix treatment of the fast walsh-hadamard transform. *IEEE Transactions on Computers*, 100(11):1142–1146, 1976.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv* preprint arXiv:2210.17323, 2022.
- Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. In *Low-power computer vision*, pp. 291–326. Chapman and Hall/CRC, 2022.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. When Attention Sink Emerges in Language Models: An Empirical View. In *International Conference on Learning Representations*. arXiv, March 2025. doi: 10.48550/arXiv.2410.10781. URL http://arxiv.org/abs/2410.10781. arXiv:2410.10781 [cs].
- Yipin Guo, Yilin Lang, and Qinyuan Ren. GPTQT: quantize large language models twice to push the efficiency. In *IEEE International Conference on Cybernetics and Intelligent Systems, CIS 2024, and IEEE International Conference on Robotics, Automation and Mechatronics, RAM 2024, Hangzhou, China, August 8-11, 2024*, pp. 368–373. IEEE, 2024. doi: 10.1109/CIS-RAM61939.2024.10672819. URL https://doi.org/10.1109/CIS-RAM61939.2024.10672819.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pp. 7514–7528. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.EMNLP-MAIN.595. URL https://doi.org/10.18653/v1/2021.emnlp-main.595.
- Alain M Leger, Takao Omachi, and Gregory K Wallace. Jpeg still picture compression algorithm. *Optical Engineering*, 30(7):947–954, 1991.

- Daiqing Li, Aleks Kamko, Ehsan Akhgari, Ali Sabet, Linmiao Xu, and Suhail Doshi. Playground v2.5: Three insights towards enhancing aesthetic quality in text-to-image generation. *CoRR*, abs/2402.17245, 2024. doi: 10.48550/ARXIV.2402.17245. URL https://doi.org/10.48550/arXiv.2402.17245.
  - Muyang Li, Yujun Lin, Zhekai Zhang, Tianle Cai, Xiuyu Li, Junxian Guo, Enze Xie, Chenlin Meng, Jun-Yan Zhu, and Song Han. Svdquant: Absorbing outliers by low-rank component for 4-bit diffusion models. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025.* OpenReview.net, 2025. URL https://openreview.net/forum?id=vWR3KuiOur.
  - Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), *Computer Vision ECCV 2014*, pp. 740–755, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10602-1.
  - Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. Spinquant: LLM quantization with learned rotations. *CoRR*, abs/2405.16406, 2024. doi: 10.48550/ARXIV.2405.16406. URL https://doi.org/10.48550/arXiv.2405.16406.
  - Yuexiao Ma, Huixia Li, Xiawu Zheng, Feng Ling, Xuefeng Xiao, Rui Wang, Shilei Wen, Fei Chao, and Rongrong Ji. Affinequant: Affine transformation quantization for large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024.* OpenReview.net, 2024. URL https://openreview.net/forum?id=of2rhALq81.
  - Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. In *International Conference on Learning Representations*, 2017.
  - Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *CoRR*, abs/2106.08295, 2021. URL https://arxiv.org/abs/2106.08295.
  - J. Princen, A. Johnson, and A. Bradley. Subband/transform coding using filter bank designs based on time domain aliasing cancellation. In *ICASSP '87. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 12, pp. 2161–2164, 1987. doi: 10.1109/ICASSP.1987. 1169405.
  - Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL https://arxiv.org/abs/2412.15115.
  - Gary J. Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(12):1649–1668, 2012. doi: 10.1109/TCSVT.2012.2221191.
  - Mingjie Sun, Xinlei Chen, J. Zico Kolter, and Zhuang Liu. Massive Activations in Large Language Models. In *First Conference on Language Modeling (COML)*. arXiv, August 2024. doi: 10.48550/arXiv.2402.17762. URL http://arxiv.org/abs/2402.17762. arXiv:2402.17762 [cs] version: 2.
  - Yuxuan Sun, Ruikang Liu, Haoli Bai, Han Bao, Kang Zhao, Yuening Li, JiaxinHu, Xianzhi Yu, Lu Hou, Chun Yuan, Xin Jiang, Wulong Liu, and Jun Yao. Flatquant: Flatness matters for LLM quantization. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=uTz2Utym5n.
  - Boris van Breugel, Yelysei Bondarenko, Paul Whatmough, and Markus Nagel. Fptquant: Function-preserving transforms for llm quantization. *arXiv preprint arXiv:2506.04985*, 2025.

- Jianyi Wang, Kelvin C. K. Chan, and Chen Change Loy. Exploring CLIP for assessing the look and feel of images. In Brian Williams, Yiling Chen, and Jennifer Neville (eds.), Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023, pp. 2555–2563. AAAI Press, 2023. doi: 10.1609/AAAI.V37I2.25353. URL https://doi.org/10.1609/aaai.v37i2.25353.
- T. Wiegand, G.J. Sullivan, G. Bjontegaard, and A. Luthra. Overview of the h.264/avc video coding standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7):560–576, 2003. doi: 10.1109/TCSVT.2003.815165.
- Guangxuan Xiao, Ji Lin, Mickaël Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant: Accurate and efficient post-training quantization for large language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA, volume 202 of Proceedings of Machine Learning Research, pp. 38087–38099. PMLR, 2023. URL https://proceedings.mlr.press/v202/xiao23c.html.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient Streaming Language Models with Attention Sinks. In *International Conference of Learning Representations*, April 2024. doi: 10.48550/arXiv.2309.17453. URL http://arxiv.org/abs/2309.17453. arXiv:2309.17453 [cs].
- Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li, Ligeng Zhu, Yao Lu, and Song Han. SANA: efficient high-resolution text-to-image synthesis with linear diffusion transformers. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025.* OpenReview.net, 2025. URL https://openreview.net/forum?id=N8Oj1XhtYZ.
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10-16, 2023, 2023. URL http://papers.nips.cc/paper\_files/paper/2023/hash/33646ef0ed554145eab65f6250fab0c9-Abstract-Conference.html.
- Tianchen Zhao, Tongcheng Fang, Haofeng Huang, Rui Wan, Widyadewi Soedarmadji, Enshu Liu, Shiyao Li, Zinan Lin, Guohao Dai, Shengen Yan, Huazhong Yang, Xuefei Ning, and Yu Wang. Vidit-q: Efficient and accurate quantization of diffusion transformers for image and video generation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=E1Nloxd63b.

# A Proofs

### A.1 Proof for Theorem 1

We prove the result reported in Theorem 1 in three steps. First, we demonstrate that the quantization error for an orthogonal sequence transformation is the same as the quantization error on the sequence transformed inputs:

$$\mathcal{L}\left(\boldsymbol{X};\boldsymbol{L}\right) = \mathcal{L}\left(\boldsymbol{L}\boldsymbol{X}\right) \tag{10}$$

*Proof.* This simply follows from the invariance of the Frobenious norm to orthogonal transformations.

$$\mathcal{L}(\boldsymbol{X}; \boldsymbol{L}) = \left\| \boldsymbol{L}^{-1} \mathcal{Q}(\boldsymbol{L} \boldsymbol{X}) - \boldsymbol{X} \right\|_{2}^{2}$$

$$= \left\| \boldsymbol{L}^{-1} \mathcal{Q}(\boldsymbol{L} \boldsymbol{X}) - \boldsymbol{L}^{-1} \boldsymbol{L} \boldsymbol{X} \right\|_{2}^{2}$$

$$= \left\| \boldsymbol{L}^{-1} \left( \mathcal{Q}(\boldsymbol{L} \boldsymbol{X}) - \boldsymbol{L} \boldsymbol{X} \right) \right\|_{2}^{2}$$

$$= \left\| \mathcal{Q}(\boldsymbol{L} \boldsymbol{X}) - \boldsymbol{L} \boldsymbol{X} \right\|_{2}^{2}$$

$$= \mathcal{L}(\boldsymbol{L} \boldsymbol{X})$$

Secondly, for completeness, we prove the upper-bound reported in Equation 3. For quantization scheme with minmax scales shared across all feature channels  $s_{ij} = \overline{s}_i$ :

$$\mathcal{L}\left(\boldsymbol{x}_{i}\right) \leq \frac{d}{4} \frac{\mathbb{E}\left[\operatorname{range}\left(\boldsymbol{x}_{i}\right)^{2}\right]}{(2^{b_{i}}-1)^{2}}$$
(11)

Proof.

$$\begin{split} \mathcal{L}\left(\boldsymbol{x}_{i}\right) &= \mathbb{E}\left[\left\|\mathcal{Q}\left(\boldsymbol{x}_{i}\right) - \boldsymbol{x}_{i}\right\|_{2}^{2}\right] \\ &= \mathbb{E}\left[\left\|\mathbf{Q}^{-1}\left(\left\lfloor\frac{x_{i}}{\overline{s}_{i}}\right\rfloor + z_{i}\right) - \boldsymbol{x}_{i}\right\|_{2}^{2}\right] \\ &= \mathbb{E}\left[\left\|\left(\left\lfloor\frac{\boldsymbol{x}_{i}}{\overline{s}_{i}}\right\rceil + z_{i} - z_{i}\right)\overline{s}_{i} - \boldsymbol{x}_{i}\right\|_{2}^{2}\right] \\ &= \mathbb{E}\left[\left\|\left(\left\lfloor\frac{\boldsymbol{x}_{i}}{\overline{s}_{i}}\right\rceil - \frac{\boldsymbol{x}_{i}}{\overline{s}_{i}}\right)\overline{s}_{i}\right\|_{2}^{2}\right] \\ &= \mathbb{E}\left[\left\|\left(\left\lfloor\frac{\boldsymbol{x}_{i}}{\overline{s}_{i}}\right\rceil - \frac{\boldsymbol{x}_{i}}{\overline{s}_{i}}\right)\right\|_{2}^{2}\overline{s}_{i}^{2}\right] \\ &\leq \mathbb{E}\left[\left\|\frac{1}{2}\right\|_{2}^{2}\overline{s}_{i}^{2}\right] \\ &= \frac{d}{4}\frac{\mathbb{E}\left[\text{range}\left(\boldsymbol{x}_{i}\right)^{2}\right]}{(2^{b_{i}} - 1)^{2}}. \end{split}$$

Lastly, we bound the range using the norm 2:

range 
$$(x_i)^2 \le 2 ||x_i||_2^2$$
. (12)

Equality is attained whenever  $\hat{x}_i$  consists of two non-zero entries -v and v, for which range (x) = 2v and  $||x||_2 = \sqrt{2}v$ .

Using this three steps we can write the proof for Theorem 1

Proof.

$$\mathcal{L}(\boldsymbol{X}; \boldsymbol{L}) \stackrel{10}{=} \mathcal{L}(\boldsymbol{L}\boldsymbol{X})$$

$$= \sum_{i=1}^{s} \mathcal{L}(\boldsymbol{l}_{i}\boldsymbol{X})$$

$$\stackrel{11}{\leq} \frac{d}{4} \sum_{i=1}^{s} \frac{\mathbb{E}\left[\text{range}(\boldsymbol{l}_{i}\boldsymbol{X})^{2}\right]}{(2^{b_{i}} - 1)^{2}}$$

$$\stackrel{12}{\leq} \frac{d}{2} \sum_{i=1}^{s} \frac{\mathbb{E}\left[\left\|\boldsymbol{l}_{i}\boldsymbol{X}\right\|_{2}^{2}\right]}{(2^{b_{i}} - 1)^{2}}$$

# A.2 OPTIMAL BIT WIDTH ALLOCATION

We determine the optimal bit width allocation  $b^*$  for an energy vector e by determining the bit width for which the ratio  $e_i/2^{2b_i^*} \approx e_i/(2^{b_i}-1)^2$  is constant for all tokens:

$$\begin{cases} \frac{e_i}{2^{2b_i^*}} = k\\ \sum_{i=1}^s b_i^* = B \end{cases}$$
 (13)

$$\Longrightarrow b_i^* = \frac{\log_2 e_i - \log_2 k}{2} \tag{14}$$

$$\implies \sum_{i=1}^{s} \frac{\log_2 e_i - \log_2 k}{2} = B \tag{15}$$

$$\implies \log_2 k = \frac{1}{s} \sum_{i=1}^s \log_2 e_i - \frac{2B}{s} \tag{16}$$

$$\implies b_i^* = \frac{\log_2 e_i - \frac{1}{s} \sum_{i=1}^s \log_2 e_i + \frac{2B}{s}}{2}$$
 (17)

$$\Longrightarrow b_i^* = \log_2 \sqrt{e_i} + \underbrace{\frac{B - \sum_{i=1}^s \log_2 \sqrt{e_i}}{s}}_{C} \tag{18}$$

# A.3 EFFECTIVENESS OF ENERGY CONCENTRATION

We demonstrate the effectiveness of the proposed Sequence Transform and Mixed precision scheme by considering the upper bound in Equation 8 in two different settings. In order to simplify computation, we will consider  $\frac{e_i}{2^{2b_i}}$  instead of  $\frac{e_i}{(2^{b_i}-1)^2}$  since the difference between the two quantities is negligible for practical values of  $b_i$ .

## 1. Uniform Energy:

In this scenario we have  $e_i = E/s$  and  $b_i = B/s$ . We note that the energy for each token is equal to the average of the squared eigenvalues  $\lambda_i^2$  of S:

$$e_i = E/s = \operatorname{Trace}(\mathbf{S})/s = \frac{1}{s} \sum_{i=1}^{s} \lambda_i^2 \stackrel{def}{=} \overline{\lambda^2}.$$
 (19)

Therefore:

$$\frac{d}{2} \sum_{i=1}^{s} \frac{e_i}{2^{2b_i}} = \frac{d}{2} \sum_{i=1}^{s} \frac{\overline{\lambda^2}}{2^{2B/s}} = \frac{ds}{2} 2^{\log_2 \overline{\lambda^2} - 2B/s}$$
 (20)

### 2. Maximum Energy Concentration:

We consider a scenario in which the energy corresponds with the squared eigenvalues  $e_i$ 

 $\lambda_i^2$ . The optimal bit allocation is given by Equation 18:

$$b_i^* = \log \lambda_i + \frac{B}{s} - \underbrace{\frac{\sum_{i=1}^s \log_2 \lambda_i}{s}}_{\overline{\log_2 \lambda}}$$
 (21)

Therefore:

$$\frac{d}{2} \sum_{i=1}^{s} \frac{e_i}{2^{2b_i}} = \frac{d}{2} \sum_{i=1}^{s} \frac{\lambda_i^2}{2^{\log \lambda_i^2 + \frac{2B}{s} - 2\overline{\log}_2 \lambda}} = \frac{ds}{2} 2^{\overline{\log}_2 \lambda^2 - 2B/s}$$
(22)

Comparing the two results is equivalent to comparing  $\log_2 \overline{\lambda^2}$  (uniform) and  $\overline{\log_2 \lambda^2}$  (max concentration). Using Jensen's inequality and the convexity of  $log_2$ , we have:

$$\overline{\log_2 \lambda^2} = \mathbb{E}[\log_2 \lambda_i^2] \le \log \mathbb{E}[\lambda_i^2] = \log_2 \overline{\lambda_i^2}. \tag{23}$$

Therefore, the Maximum Energy concentration strategies achieves a lower value than the uniform scheme.

# B EXPERIMENTAL DETAILS

# B.1 LVMs

**Baselines** We implement SmoothQuant, QuaRot, ViDiT-Q and SVDQuant closely following the details reported in the paper and the respective codebases. Specifically, for ViDiT-Q (Zhao et al., 2025) we use  $\alpha=0.01$ , as reported in their PixArt- $\Sigma$  setup. For SVDQuant and SmoothQuant we use a default value of  $\alpha=0.5$ .

**Quantization** In order to promote a fair comparison with the other models, for the results reported in Table 1, instead of using the 8-bits weight quatization scheme described in ViDiT-Q (8 bits for the FFN blocks, first and last transformer blocks), we use per-block weight and activation quantization, as described in SVDQuant (Li et al., 2025), which improves the ViDiT-Q results. In line with the SVDQuant methodology, we retain the depth-wise convolutions within the feed-forward layers of the SANA transformer blocks in full precision. Meanwhile, the two point-wise convolutions are quantized by treating them as linear layers. Notably, unlike the original SVDQuant implementation, our experiments do not incorporate GPTQ at any stage. The activation-quantization only results reported in Figure 4b and Figure 6 use 4 bits per token quantization.

**STaMP** For the STaMP results, excluding the specific ablation study, we use 64 high-precision tokens, which results in an effective activation bit-width of 4.0625 bits on the PixArt- $\Sigma$  model and 4.125 bits on the SANA model. Non-STaMP result do not use any mixed precision tokens.

### B.2 LLMs

**Quantization** We use dynamic quantization for the KV cache and activations. Quantization scales and offsets are determined per token, sequence, and (for KV) head. We follow (Ashkboos et al., 2024; Sun et al., 2025) and only quantize weights, KV cache, and inputs to linear layers within the transformer block. We use round-to-nearest weight quantization, as weights are unaffected by STaMP—more advanced weight quantization schemes could improve results further, but this is perpendicular to our contribution. we range set the weights by computing the weight quantization squared error for a grid of candidate ranges and selecting the candidate with lowest error. For fairness, for all experiments (incl. baselines) we keep the first 64 tokens in 8 bits, which gives an effective bit width of 4.125 for both STaMP and non-STaMP results.

**Baselines** For SmoothQuant, we calibrate the scales based on the Wikitext-2 training dataset activations. For QuaRot, we follow the original paper and reduce the minmax activation range by 10%. For FlatQuant, we use their recommended settings for all models (e.g. training for 15 epochs on 128 Wikitext-2 training sequences).

**STaMP** The first token in most LLMs acts as an attention sink Xiao et al. (2024); Gu et al. (2025), and typically contains massive outliers Sun et al. (2024). Keeping the first token in 8 bits helps to accurately represent these massive outliers. To ensure the massive outlier stays in the first token, STaMP is not applied to the first token. Note that the baselines that do not use STaMP's transform, do benefit from keeping the first token (and next 63) in high-precision in our experiments.

# C ADDITIONAL RESULTS

We report additional results for the setups described in Section 5.

# D USE OF LLMS

Large Language Models have been used only for sentence rephrasing and grammar check, which resulted in minor alterations to the paper text.

# "a cat that has a shirt on its back" Sequence: DCT Sequence: WHT Sequence: DWT Sequence: Identity Feature: SmoothQuant Feature: SDCB

Figure 7: Visualization of the Images generated with A4 activation quantization and several combination of Feature and Sequence transforms for the Pixart- $\Sigma$  architecture. The images refer to the same setup described in Figure 6.

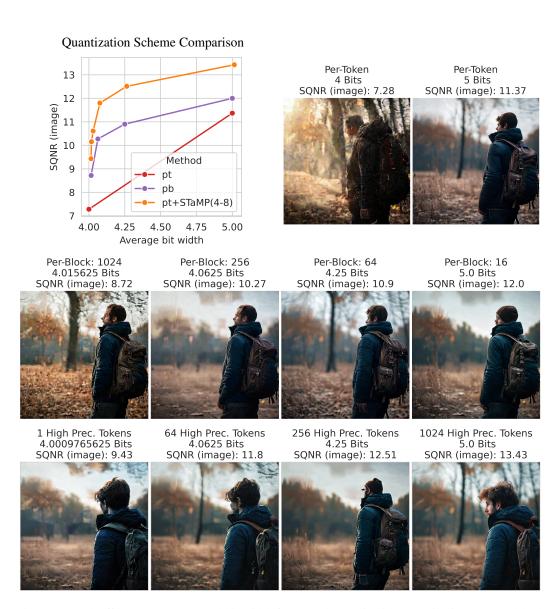


Figure 8: Tradeoff between SQNR and bit width for per-token activation quantization (pt) per-block activation quantization (pb) at different block size (from 16 to 1024), and per-token with STaMP (pb+STaMP) at varying number of high precision tokens on the PixArt- $\Sigma$  model. We consider 16 bits scales for each scale parameter. The visualization correspond to the Prompt 'A guy with a backpack looking at the ground to his left.'

Dataset	LVM	Method	STaMP	CLIP	CLIP IQA	SQNR (latent)
COCO	PixArt-Σ	FP		31.54	0.91	$+\infty$
		RTN	X	31.23	0.79	0.19
			✓	31.88	0.91	0.49
		SVDQuant	X	31.68	0.91	0.96
			✓	31.76	0.91	1.47
		ViDiT-Q	X	31.63	0.87	0.73
			✓	31.92	0.87	1.03
	SANA	FP		31.85	0.89	151.70
		RTN	Х	31.85	0.89	3.39
			✓	31.94	0.89	4.09
		SVDQuant	Х	31.83	0.90	4.61
			✓	31.88	0.90	5.30
		ViDiT-Q	X	31.79	0.89	4.63
			<b>√</b>	31.82	0.89	5.42
МЈНО	PixArt-Σ	FP	X	31.46	0.85	$+\infty$
		RTN	X	29.90	0.75	0.28
			<b>√</b>	30.96	0.87	0.71
		SVDQuant	X	31.13	0.86	1.27
			<b>√</b>	31.28	0.87	1.86
		ViDiT-Q	X	30.97	0.81	0.86
			✓	31.20	0.83	1.31
	SANA	FP		31.57	0.83	151.93
		RTN	X	31.51	0.84	4.45
			$\checkmark$	31.60	0.84	5.39
		SVDQuant	Х	31.52	0.83	5.69
			$\checkmark$	31.55	0.83	6.46
		ViDiT-Q	X	31.60	0.83	5.76
			$\checkmark$	31.60	0.83	6.65

Table 3: Additional metrics for the experiments reported in Table 1.



Figure 9: Visualization of a sample generation for the SANA for the results reported in Table 1.