

OmniFusion: Simultaneous Multilingual and Multimodal Translations via Modular Fusion

Sai Koneru^{1,3*}, Matthias Huck², and Jan Niehues¹

¹ Karlsruhe Institute of Technology

² SAP SE, Dietmar-Hopp-Allee 16, 69190 Walldorf, Germany

³ AppTek, Aachen, Germany

{sai.koneru, jan.niehues}@kit.edu

matthias.huck@sap.com

Abstract

Open-source text-only translation LLMs have rapidly improved multilingual coverage and translation quality, but their unimodal design limits their use in multimodal translation. In speech translation, they are commonly used in cascaded ASR–MT pipelines, which introduce additional latency and prevent the translation model from directly using visual context. Conversely, pretrained multimodal foundation models provide strong audio-visual perception but generally lag behind specialized translation LLMs in multilingual generation quality. We introduce OmniFusion, a modular architecture that connects a frozen multimodal foundation model to a translation-specialized LLM through a lightweight gated fusion module. OmniFusion supports speech-to-text, speech-and-image-to-text, and text-and-image-to-text translation within a single framework. Experiments with Qwen2.5-Omni-7B and Seed-X-PPO-7B show that selected early and middle MMFM layers provide more transferable representations than late layers. On MCIF, OmniFusion improves the latency–quality trade-off over a matched cascaded SimulST pipeline when evaluated with Local Agreement decoding, while remaining competitive in offline ST and reducing severe XCOMET-XL errors. On CoMMuTE, it achieves competitive visually grounded caption translation performance.¹

1 Introduction

In recent years, NLP has shifted from unimodal foundation models (Touvron et al., 2023; Grattafiori et al., 2024) to multimodal foundation models (MMFMs) (Abdin et al., 2024; Goel et al., 2025; Xu et al., 2025). While unimodal systems work well in some settings, many tasks benefit from additional cross-modal cues.

Translation is one such task. Speech Translation (ST) converts spoken input into a target language (Ahmad et al., 2024; Agostinelli et al., 2025), while caption translation often requires image context for disambiguation. However, using visual context such as images or presentation slides in ST remains underexplored (Liu and Niehues, 2024; Gaido et al., 2024a; Sinhamahapatra and Niehues, 2025).

Although MMFMs process multimodal inputs, their multilingual translation quality is generally weaker than specialized translation LLMs (Alves et al., 2024; Cheng et al., 2025). In ST, cascaded ASR–MT pipelines (Ahmad et al., 2024; Agostinelli et al., 2025; Koneru et al., 2025) remain strong, but they cannot directly exploit multimodal context and add latency, especially in simultaneous ST (SimulST) (Agostinelli et al., 2025; Papi et al., 2025a).

These limitations motivate an end-to-end (E2E) multilingual and multimodal translation system for ST, image-aware ST, and caption translation. Existing work typically treats speech translation, speech–image translation, and text–image translation as separate tasks, rather than as a unified multimodal translation problem. Prior work (Ambilduke et al., 2025; Viveiros et al., 2025) often trains models from scratch or uses only pretrained encoders, leaving the internal cross-modal representations of MMFMs underutilized.

We address this gap with OmniFusion, which uses a frozen MMFM as a multimodal representation source for a translation-specialized LLM. OmniFusion learns a lightweight fusion interface and uses multitask training with OCR and ASR self-cascading to improve multimodal alignment and generation.

Our central hypothesis is that MMFM hidden states contain useful cross-modal signals that can enable all multimodal translation tasks with limited training.

^{*}Work done at Karlsruhe Institute of Technology.

¹Code is available at <https://github.com/saikoneru/OmniFusion>




Figure 1: Architecture of OmniFusion with supported tasks and inference modes.

Our contributions are: (1) We propose OmniFusion, a modular architecture connecting a frozen MMFM to a translation-specialized LLM, enabling speech-to-text, speech-and-image-to-text, and text-and-image-to-text translation. (2) We show that MMFM layer choice is critical where different layers transfer better for different translation tasks and gated fusion is important. (3) OmniFusion with Local Agreement decoding (Liu et al., 2020) improves the SimulST latency–quality trade-off over the same cascade (Figure 3), reduces major and critical offline ST errors (Table 3), and achieves competitive CoMMuTE performance (Table 4).

2 Approach

We propose a modular fusion framework that combines a frozen multimodal foundation model, $\mathcal{M}_{\text{MMFM}}$, for audio and image encoding with a translation-specialized language model, $\mathcal{M}_{\text{NMT}}$, for multilingual generation. A lightweight fusion module selects MMFM hidden representations and projects them into the translation LLM’s embedding space, enabling direct conditioning on audio and visual inputs.

Our main architectural contribution is a *gated fusion module* (§2.2) that integrates MMFM representations into the translation LLM. We motivate this design through layer-probing analysis (§2.1.1), which shows that early and middle MMFM layers transfer best for translation. We improve performance through the fusion architecture and multi-task training with self-cascading ASR and OCR objectives (§2.3). The system is trained as an offline E2E translation model and applied to SimulST at inference time using Local Agreement decoding

(Liu et al., 2020). Figure 1 illustrates the system architecture and supported tasks.

2.1 Architecture

Let $\mathcal{M}_{\text{NMT}}$ denote the translation model and $\mathcal{M}_{\text{MMFM}}$ denote the multimodal foundation model. Both models use a language-model backbone. As a reference, standard text-only NMT trains $\mathcal{M}_{\text{NMT}}$ on source sentence $X = (x_1, \dots, x_{L_x})$ and target sentence $Y = (y_1, \dots, y_{L_y})$ with the teacher-forcing cross-entropy objective:

$$\mathcal{L}_{\text{Trans}} = -\frac{1}{L_y} \sum_{t=1}^{L_y} \log P_{\theta}(y_t \mid y_{<t}, X), \quad (1)$$

where $X = (x_1, \dots, x_{L_x})$ and $Y = (y_1, \dots, y_{L_y})$ denote the source and target token sequences, respectively.

For multimodal translation, image and audio inputs are encoded by $\mathcal{M}_{\text{MMFM}}$ and projected into hidden states compatible with its language-model backbone. Our goal is to adapt the translation model to use these representations as additional context.

2.1.1 Probing MMFM Representations

Different MMFM layers may expose different information to the translation LLM. We use adjacent-layer cosine similarity as a diagnostic to identify candidate layers with distinct representational behavior (Figure 2). This diagnostic does not prove what semantic information each layer encodes; rather, it helps select a small candidate set instead of using all MMFM layers, which would increase sequence length and may introduce noise. We see




Figure 2: Average cosine similarity between hidden representations of adjacent layers, computed with Qwen-Omni 2.5 7B on the MCIF test set (en $\rightarrow$ de) under different modality conditions.

the shift occurs drastically at first and last layer indicating three distinct representation regions. Hence, we select the first, middle, and last MMFM layers as candidate transfer points to cover these regions. Their usefulness is validated through downstream ablations in Table 1 and Appendix Table 6.

2.2 Gated Layer Fusion

Let $Z = (z_1, \dots, z_{L_m})$ denote the multimodal token sequence produced by the MMFM input encoder, where L_m is the multimodal sequence length. We denote by $h_t^{(k)} \in \mathbb{R}^{D_{\text{MMFM}}}$ the hidden state of token z_t at layer k of $\mathcal{M}_{\text{MMFM}}$, for $k \in \{\text{first}, \text{mid}, \text{last}\}$ and $t = 1, \dots, L_m$.

Naively feeding all three layer representations as separate context tokens would triple the multimodal sequence length, treat all layers as equally informative, and ignore the dimensional mismatch between D_{MMFM} and D_{NMT} . We address these with a *token-wise gated fusion* module.

For each multimodal token z_t , we concatenate the three selected layer representations:

$$\begin{aligned} c_t &= h_t^{(\text{first})} \oplus h_t^{(\text{mid})} \oplus h_t^{(\text{last})}, \\ c_t &\in \mathbb{R}^{3D_{\text{MMFM}}}, \quad t = 1, \dots, L_m, \end{aligned} \quad (2)$$

where $\oplus$ denotes feature-wise concatenation.

A bias-free linear layer $W_{\text{gate}} \in \mathbb{R}^{3 \times 3D_{\text{MMFM}}}$ maps c_t to a three-dimensional logit vector, which is normalized via softmax to give per-layer importance weights:

$$\begin{aligned} \alpha_t &= \text{softmax}(W_{\text{gate}} c_t), \\ \alpha_t &= [\alpha_t^{(\text{first})}, \alpha_t^{(\text{mid})}, \alpha_t^{(\text{last})}] \in \mathbb{R}^3. \end{aligned} \quad (3)$$

The fused MMFM representation is computed

as a weighted sum:

$$\begin{aligned} r_t &= \alpha_t^{(\text{first})} h_t^{(\text{first})} + \alpha_t^{(\text{mid})} h_t^{(\text{mid})} + \alpha_t^{(\text{last})} h_t^{(\text{last})}, \\ r_t &\in \mathbb{R}^{D_{\text{MMFM}}}, \quad t = 1, \dots, L_m. \end{aligned} \quad (4)$$

Finally, the fused representation is projected into the embedding space of the translation model using a three-layer MLP for sufficient transformation computation for all modalities with GELU activation:

$$\begin{aligned} v_t &= \text{MLP}(r_t), \\ v_t &\in \mathbb{R}^{D_{\text{NMT}}}, \quad t = 1, \dots, L_m. \end{aligned} \quad (5)$$

where D_{NMT} is the token embedding dimension of $\mathcal{M}_{\text{NMT}}$.

Let $e_i \in \mathbb{R}^{D_{\text{NMT}}}$ denote the embedding of source token x_i . The projected multimodal embeddings are prepended to the source-token embeddings (sentence or tag) along the sequence dimension:

$$\begin{aligned} \tilde{X} &= (v_1, \dots, v_{L_m}, e_1, \dots, e_{L_x}), \\ \tilde{X} &\in \mathbb{R}^{(L_m + L_x) \times D_{\text{NMT}}}. \end{aligned} \quad (6)$$

Thus, the fused multimodal representations are treated as additional input tokens for the translation model.

The multimodal translation objective is then defined as:

$$\mathcal{L}_{\text{MM-Trans}} = -\frac{1}{L_y} \sum_{t=1}^{L_y} \log P_{\theta}(y_t | y_{<t}, \tilde{X}). \quad (7)$$

2.3 Training

We align multimodal features with the translation LLM through multitask training with the objectives below.

2.3.1 Multi-task Learning

We train on three tasks. In ST, the model translates source-language audio. In Speech-Image Translation (SIT), it translates audio with an aligned image, such as a presentation slide. In Text-Image Translation (TIT), it translates source text paired with an image. These tasks expose the model to different modality combinations and encourage multimodal alignment.

2.3.2 Self-Cascading (ASR and OCR)

In addition to multi-task learning, we adopt self-cascading, which decomposes translation into intermediate steps. For example in English-to-German

ST, the model first generates the ASR transcription, then produces the final translation conditioned on it. For visual inputs, OCR text serves as an intermediate step before translation. During training, self-cascading is applied stochastically on the samples with modality-specific prompts for better MMFM representations (Algorithm 1). We freeze the MMFM and fine-tune the translation LLM with LoRA (Hu et al., 2021). We find that adding LoRA to the MMFM degrades performance due to catastrophic forgetting.

We distinguish between training-time and inference-time self-cascading. When training, self-cascading examples teach the model to generate intermediate ASR or OCR text before the final translation. At inference time, we evaluate both direct generation, where the model outputs only the translation or self-cascade generation.

3 Experimental Setup

3.1 Training Data

We train OmniFusion on speech- and image-conditioned translation data. An overview of data sources is provided in Appendix Table 7. For ST, we use subsets of Europarl-ST (Iranzo-Sánchez et al., 2020) and CoVoST (Wang et al., 2020) across multiple target languages. For image-aware ST, we use M3AV (Chen et al., 2024), which contains English scientific talks paired with slides and PaddleOCR (Cui et al., 2025) predictions. Since M3AV has limited language coverage, we augment ASR transcripts with Seed-X-PPO-7B (Cheng et al., 2025) translations to create **M3AV-Aug**, enabling speech-image translation and self-cascade training. We do not apply filtering but show that the translation quality is high even when training on this synthetic data. For TIT, we use Multi30k (Elliott et al., 2016), which pairs images with multilingual captions and encourages visual grounding beyond OCR.

3.2 Models

Several MMFMs have recently shown strong gains in quality and efficiency (Abdin et al., 2024; Team et al., 2025). As our setting requires processing both audio and image modalities with cross-modal reasoning, we use Qwen2.5-Omni 7B² (Xu et al., 2025). For translation, we adopt Seed-X-PPO-7B³

(Cheng et al., 2025), which provides broad language coverage and strong translation performance compared to closed-source alternatives. Using the training data described above, we refer to the resulting fine-tuned model as **OmniFusion**. While we report results with these models, preliminary experiments with alternative model pairings (Appendix Table 6) suggest the approach can extend to other pairings. Training hyperparameters are listed in Table 14. Validation uses CoVoST dev data, selecting the checkpoint with the lowest loss after 20k steps.

3.3 Evaluation: Data & Metrics

Evaluating multimodal SimulST requires aligned audio-visual translation data. We exclude M3AV from testing due to training overlap and its monolingual nature, and instead assess robustness on unseen domains using the MCIF test set (Papi et al., 2025b), which contains ACL talk recordings translated from English into three target languages, and FLEURS (Conneau et al., 2023) for multilingual coverage. Latency is reported using Average Lagging (Ma et al., 2020), and translation quality is measured on these datasets.

For TIT, we target the cases where visual context is crucial for disambiguation, using the CoMMuTE benchmark (Futeral et al., 2023, 2025). Translation quality is evaluated with XCOMET-XL (Guerreiro et al., 2024) for ST, enabling detailed error analysis, and with COMET⁴ (Rei et al., 2022) for TIT.

4 Results

Our evaluation is organized around four questions. First, we ask whether OmniFusion exploiting MMFM representation improves the latency-quality trade-off for SimulST. Second, we analyze whether the proposed fusion architecture and MMFM layer selection are necessary for effective transfer to the translation LLM. Third, we examine whether the multitask training strategy is necessary, especially the OCR and self-cascading objectives. Fourth, we evaluate whether the resulting model remains competitive in offline ST and TIT, where latency constraints are removed. Finally, we perform inference-time analysis of the learned gate behavior to understand layer contribution.

4.1 Simultaneous Speech Translation

A SimulST system must translate incrementally as speech arrives, deciding when to read more input

²Qwen/Qwen2.5-Omni-7B

³ByteDance-Seed/Seed-X-PPO-7B

⁴Unbabel/wmt22-comet-da




Figure 3: SimulST performance on en→de and en→it with chunk sizes of [1, 1.5, 2, 2.5, 3] seconds. We report XCOMET-XL quality scores against Average Lagging (AL; lower is better). Top figures show computation-unaware AL, while bottom figures show computation-aware AL.

and when to commit target tokens. Since OmniFusion is trained as an E2E translation model and does not natively implement streaming decoding or a learned policy, we apply Local Agreement (Liu et al., 2020) at inference time. Note that local agreement is commonly used strong strategy for SimulST in IWSLT 2025 (Agostinelli et al., 2025). Audio is processed in fixed-size chunks; at each step, the model generates a new hypothesis, commits the longest common prefix between consecutive hypotheses, and force-decodes committed tokens in subsequent steps. For the cascaded baseline, we apply the same policy in two stages: Local Agreement is first applied to ASR hypotheses, and MT is then conditioned on the agreed ASR prefix. This creates a strong but latency-sensitive cascade, since translation can only proceed after ASR prefixes stabilize following IWSLT submissions (Agostinelli et al., 2025).

We compare OmniFusion against this fine-tuned cascade built from the same pretrained components,

and include SeamlessM4T Large-v2 (Barrault et al., 2023) and StreamUni-Phi4 (Guo et al., 2025) baselines to contextualize the latency–quality trade-off. We note that en→de is the only evaluated language pair supported by both external SimulST baselines; we do not report on en-it as it is not covered by StreamUni-Phi4 and Seamless performance was below 70 XCOMET skewing the visualization.

We evaluate chunk sizes of [1, 1.5, 2, 2.5, 3] seconds on MCIF for en→de and en→it, with and without image input. We report XCOMET-XL quality and both computation-unaware and computation-aware Average Lagging (AL) in Figure 3; additional details are in Appendix A.6.

The offline systems provide an expected quality upper bound, but have high latency. Under Local Agreement decoding, OmniFusion achieves better latency and quality than the fine-tuned cascade even when applying local agreement on the ASR. In the audio-only setting (•), it reaches comparable quality with lower latency, indicating that the E2E

Model	Mode	Fusion Layers	en → de		zh → de		ja → ru		fr → ko		ur → zh	
			Chrf	Comet-22	Chrf	Comet-22	Chrf	Comet-22	Chrf	Comet-22	Chrf	Comet-22
Tower-Spire	Direct	✗	56.9	82.54	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Tower-Spire	Self-Cascade	✗	61.9	85.02	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
SeamlessM4T v2	Direct	✗	62.4	85.60	35.8	70.22	36.9	71.77	20.8	74.12	13.5	61.11
SeamlessM4T v2	Self-Cascade	✗	63.3	85.81	36.9	71.34	38.2	69.85	21.0	75.22	13.5	63.75
Our Models - OmniFusion												
Attention	Direct	All	60.9	84.02	37.4	72.96	26.6	65.62	21.2	75.36	4.3	43.66
Attention	Direct	First, Mid, Last	61.8	86.95	39.4	74.86	27.6	68.81	21.9	78.94	5.5	46.78
Gated	Direct	All	62.3	85.37	39.6	75.91	28.8	69.33	21.8	77.50	5.0	45.78
Gated	Direct	First, Mid, Last	64.0	87.12	44.5	81.82	37.2	81.95	25.5	82.74	12.5	65.60

Table 1: Fusion architectures with zero-shot source languages. Tower-Spire only supports English audio.

model avoids part of the delay from ASR stabilization. With image input (■), OmniFusion obtains further quality gains, suggesting that visual context helps stabilize translation decisions.

Compared with SeamlessM4T Large-v2 and StreamUni-Phi4, OmniFusion achieves substantially higher XCOMET-XL scores in our evaluated settings. StreamUni-Phi4 obtains lower computation-aware latency because it is designed for streaming generation and can reuse cached computation. OmniFusion is not trained with a streaming-specific objective; adding streaming-specific training and caching support is a promising direction for future work. Overall, OmniFusion is able to: preserve competitive quality and reduces latency relative to the cascade, and exploit visual context when available.

4.2 Fusion Architecture and Layer Selection

We next analyze whether OmniFusion’s performance depends on the proposed fusion design. The final model uses a gated module over three MMFM layers: the first, middle, and last layers. This design is motivated by the hypothesis that different MMFM layers encode different levels of acoustic, visual, and semantic information, and that using all layers may introduce noise rather than improve transfer.

To test this, we compare gated fusion with attention-based fusion (detailed in Appendix A.3), using either all MMFM layers or only the selected three layers. We evaluate these variants on FLEURS, including non-English source languages, to test whether the fused representations transfer beyond the English-centric training setting. Further, we also compare with Tower-Spire (Ambil-[duke et al., 2025](#)) and SeamlessM4T v2 to test the multilingual ability.

As shown in Table 1, the three-layer gated model performs best across all evaluated language pairs.

Using all MMFM layers consistently degrades performance, suggesting that later MMFM representations can introduce noise when transferred to the translation LLM. Attention-based fusion is also less robust than gated fusion, particularly in zero-shot source-language settings. This indicates that a lightweight layer-wise gate is better suited for aligning MMFM representations with the translation LLM than attending over all hidden states.

Moreover, OmniFusion is better than SeamlessM4T and Tower-Spire in all directions by exploiting the state-of-the-art pretrained models. These results show that the selected-layer fusion interface is important: simply exposing the translation LLM to more MMFM layers does not improve performance. Instead, the best results in ST come from selecting a small number of transferable layers and combining them through a lightweight gating mechanism.

4.3 Impact of Training Tasks

Training Configuration	en→de	en→it	en→zh
W/o OCR	87.12	82.42	79.75
W/o (OCR + Self-Cascade)	89.19	84.35	82.14
All Tasks	89.90	86.79	82.03
All Tasks (Self-Cascade Inference)	90.05	87.20	82.47

Table 2: Training task ablations on the MCIF test set.

The previous ablation shows the benefit of lightweight gate, but the final system also relies on multitask training. We therefore ablate the training tasks from Section 2.3: one model removes OCR, and another removes both OCR and self-cascading. We evaluate all variants on MCIF.

As shown in Table 2, removing OCR gives the weakest performance across all three language pairs. Interestingly, removing both OCR and self-cascading performs better than removing OCR alone. We hypothesize that self-cascading with

Category	Config	Image	XCOMET-XL	Minor	Major	Critical	Total
Omni FT (Direct ST)	LoRA FT	✗	82.30	1079.0	852.0	90.3	2021.3
Cascade Systems (ASR → MT) Omni + SeedX	No FT	✗	85.88	1212.3	773.0	74.3	2059.7
	No FT	✓	85.94	1191.3	759.3	74.7	2025.3
	LoRA FT	✗	86.42	1239.0	751.3	75.0	2065.3
	LoRA FT	✓	86.59	1202.3	766.3	55.7	2024.3
End-to-End (Direct ST) OmniFusion	Mid Fusion	✗	83.57	1114.0	1188.3	72.3	2374.7
	Mid Fusion	✓	85.98	1256.3	749.0	61.7	2067.0
	Gated Fusion	✗	83.98	1119.3	1188.0	62.7	2370.0
	Gated Fusion	✓	86.24	1255.7	736.0	57.0	2048.7
	Gated Fusion + Self-Cascade	✓	86.57	1231.0	719.0	55.3	2005.3

Table 3: Offline ST results averaged across en→zh, en→de, and en→it on the MCIF test. ✓ indicates models using image input. Bold indicates best XCOMET-XL and lowest error counts. Language-wise results are additionally reported in Table 10, 11, and 12.

Model	Config	en → ar	en → cs	en → de	en → fr	en → ru	en → zh
Omni 2.5 7B	-	78.56	81.03	81.94	79.99	47.27	81.74
Seed-X-PPO-7B	-	80.39	83.56	82.38	80.31	82.13	83.25
ZeroMMT 3.3B	-	80.64	86.69	83.54	83.40	83.87	77.85
TowerVision 9B	-	70.85	86.62	84.97	84.29	85.81	85.24
OmniFusion 14B	Mid Fusion	80.98	87.49	85.14	84.97	85.69	85.15
OmniFusion 14B	Gated Fusion	80.81	87.58	84.88	84.28	84.62	84.34

Table 4: COMET scores (↑) on the CoMMuTE dataset. Statistically significant top-performing results per language pair are highlighted in **bold** (computed using comet-compare with default settings). Qualitative example can be found in Table 16.

images can introduce confusion when visual representations are not well aligned, and that the OCR objective helps ground these representations, enabling more effective self-cascade training.

Training on all tasks yields the strongest direct inference results for en→de and en→it, while remaining competitive for en→zh. Under self-cascade inference, the full model performs best across all language pairs, showing that OmniFusion’s gains come not only from fusion, but also from multitask and self-cascading objectives that better align MMFM representations with translation generation.

4.4 Offline Speech Translation

Having established that both the fusion architecture and training tasks are necessary, we next evaluate the final OmniFusion system in offline ST, where the full speech input is available and latency constraints are removed. This setting allows us to compare translation quality directly against cascaded systems and to assess whether the E2E model remains competitive when the cascade is not penalized for latency.

Table 3 reports XCOMET-XL scores and fine-grained error counts on MCIF, averaged over

en→zh, en→de, and en→it. Directly fine-tuning the MMFM for speech translation performs substantially worse than OmniFusion, confirming that adapting the MMFM alone is insufficient. In the cascaded pipeline, LoRA fine-tuning improves quality, while image input provides only marginal gains. In contrast, OmniFusion benefits more clearly from image context, suggesting that our model can use visual information more directly than the cascade.

The best cascaded system achieves the highest aggregate XCOMET-XL score, showing that cascades remain strong when full input and latency are not constraints. However, OmniFusion with gated fusion, image input, and self-cascade inference nearly matches it while producing fewer severe errors. This suggests the end-to-end fused model improves not only latency but also robustness by conditioning directly on multimodal context rather than only an ASR transcript.

4.5 Text-Image Translation

We further evaluate whether OmniFusion can exploit visual context beyond speech translation. CoMMuTE targets visually grounded ambiguity, making it a useful benchmark for assessing lexi-

cal and semantic disambiguation with image context. As shown in Table 4, the unimodal translation model SeedX performs competitively, but is generally weaker than multimodal systems. OmniFusion improves over the individual base components and matches or surpasses strong multimodal baselines such as ZeroMMT and TowerVision on most language pairs. We also conduct a small manual disambiguation evaluation in Appendix A.1. TowerVision performs better in this targeted setting, but OmniFusion remains competitive despite not being extensively trained for image-only translation and while additionally supporting audio input. This suggests that fusing a pretrained MMFM with a translation LLM can yield strong visually grounded translation without training from scratch.

Mid Fusion slightly outperforms Gated Fusion on CoMMuTE, indicating that fusion should match the modality’s role. Middle-layer fusion is sufficient for text-image translation, where images mainly disambiguate text, while gated fusion is better suited to speech translation, where audio is the primary input. Thus, MMFM layer selection matters, and no fusion strategy is universally optimal.

4.6 Inference-Time Gate Behavior




Figure 4: Gate predictions across modalities.

Finally, we analyze how the trained gate uses MMFM layers across different input settings. We compute the average gate weights assigned to the first, middle, and last MMFM layers for speech-only, image-only, and speech+image inputs. As shown in Figure 4, the gate consistently assigns low weight to the last layer, suggesting that final MMFM states are less useful as transfer points for translation. The relative use of the first and middle layers varies by task: speech inputs place more weight on early representations, while image-only and speech+image inputs rely more strongly on middle-layer representations. This indicates that

the gate adapts to the available modality rather than selecting a fixed layer for all tasks.

We see that useful transfer depends not only on connecting an MMFM to a translation LLM, but also on selecting modality-appropriate internal representations.

5 Related Work

Multimodal MT: Caption translation has been widely studied (Specia et al., 2016; Elliott et al., 2017; Barrault et al., 2018), including work that trains multimodal NMT models from scratch (Vijayan et al., 2024) or augments pretrained MT models such as NLLB (Costa-Jussà et al., 2022; Futeral et al., 2025). Recent work extends translation LLMs to images, video, or speech (Viveiros et al., 2025; Ambilduke et al., 2025), while fewer studies jointly use speech and visual context (Sinhmahapatra and Niehues, 2025; Anwar et al., 2023; Choi et al., 2024). In contrast, OmniFusion jointly leverages images and audio through a lightweight fusion interface.

Simultaneous ST: Real-time ST requires translating before the full utterance is available. Prior work often adapts offline models using fixed or learned policies such as wait- k (Ma et al., 2019; Elbayad et al., 2020), Local Agreement (Liu et al., 2020), attention-based heuristics (Papi et al., 2023), or LLM-based policies (Koshkin et al., 2024b,a; Guo et al., 2025; Fu et al., 2025). We use Local Agreement for evaluation, but OmniFusion is compatible with other SimulST policies.

Modality Fusion: Multimodal LLMs typically process modalities with separate encoders and merge them in the LLM (Abdin et al., 2024; Xu et al., 2025; Team et al., 2025), while other models enable earlier audio-visual interaction (Ye et al., 2025). Related work also studies modality projectors (Verdini et al., 2024; Gaido et al., 2024b) and discrete tokens (Zhan et al., 2024; Li et al., 2025). OmniFusion instead studies which internal MMFM layers provide transferable representations.

6 Conclusion

We introduced OmniFusion, a modular framework that connects a frozen MMFM with a translation LLM for speech, speech-image, and text-image translation. Experiments show that early and middle MMFM layers transfer better than late layers, and that selected-layer fusion improves over all-layer fusion in ST. OmniFusion improves the

SimulST latency–quality trade-off under Local Agreement decoding while remaining competitive in offline ST and CoMMuTE. Finally, the best fusion strategy is task-dependent: gated fusion is more robust for ST, while middle-layer fusion can suffice when images mainly provide disambiguation context.

Limitations

Our training data is primarily English-centric, which may limit performance for other source languages. However, our zero-shot experiments on non-English source languages suggest encouraging transfer, but broader source-language training and evaluation are needed. Another limitation of our work is that OmniFusion is trained with fixed task-specific prompts across the three supported translation settings, and therefore does not exhibit general instruction-following capabilities. It remains unclear whether our fusion strategy can be extended to support more flexible, instruction-driven multimodal behavior. Our manual CoMMuTE analysis is small-scale and shows that OmniFusion achieves reasonable performance on targeted visual disambiguation. This suggests that broader human evaluation is needed to fully assess visually grounded translation behavior.

Acknowledgments

The research leading to these results was partially supported by European Union’s Horizon Europe programme grant agreement No. 101213369 (DVPS).

References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, James R Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C T Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. Phi-4 technical report. *arXiv [cs.CL]*.
- Victor Agostinelli, Tanel Alumäe, Antonios Anastasopoulos, Luisa Bentivogli, Ondřej Bojar, Claudia Borg, Fethi Bougares, Roldano Cattoni, Mauro Cettolo, Lizhong Chen, and 1 others. 2025. Findings of the iwslt 2025 evaluation campaign. In *Proceedings of the 22nd International Conference on Spoken Language Translation (IWSLT 2025)*, pages 412–481.
- Ibrahim Said Ahmad, Antonios Anastasopoulos, Ondřej Bojar, Claudia Borg, Marine Carpuat, Roldano Cattoni, Mauro Cettolo, William Chen, Qianqian Dong, Marcello Federico, Barry Haddow, Dávid Javorský, Mateusz Krubiński, Tsz Kin Lam, Xutai Ma, Prashant Mathur, Evgeny Matusov, Chandresh Maurya, John McCrae, and 25 others. 2024. [FINDINGS OF THE IWSLT 2024 EVALUATION CAMPAIGN](#). In *Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024)*, pages 1–11, Bangkok, Thailand (in-person and online). Association for Computational Linguistics.
- Duarte M Alves, José Pombal, Nuno M Guerreiro, Pedro H Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G C de Souza, and André F T Martins. 2024. Tower: An open multilingual large language model for translation-related tasks. *arXiv [cs.CL]*.
- Kshitij Ambilduke, Ben Peters, Sonal Sannigrahi, Anil Keshwani, Tsz Kin Lam, Bruno Martins, André F T Martins, and Marcely Zanon Boito. 2025. From TOWER to SPIRE: Adding the speech modality to a translation-specialist LLM. *arXiv [cs.CL]*.
- Mohamed Anwar, Bowen Shi, Vedanuj Goswami, Wei-Ning Hsu, Juan Pino, and Changhan Wang. 2023. Muavac: A multilingual audio-visual corpus for robust speech recognition and robust speech-to-text translation. In *Proc. Interspeech 2023*, pages 4064–4068.
- Loïc Barrault, Fethi Bougares, Lucia Specia, Chiraag Lala, Desmond Elliott, and Stella Frank. 2018. [Findings of the third shared task on multimodal machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 304–323, Belgium, Brussels. Association for Computational Linguistics.
- Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, and 1 others. 2023. Seamless4t: massively multilingual & multimodal machine translation. *arXiv preprint arXiv:2308.11596*.
- Zhe Chen, Heyang Liu, Wenyi Yu, Guangzhi Sun, Hongcheng Liu, Ji Wu, Chao Zhang, Yu Wang, and Yanfeng Wang. 2024. [M³AV: A multimodal, multi-genre, and multipurpose audio-visual academic lecture dataset](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9041–9060, Bangkok, Thailand. Association for Computational Linguistics.
- Shanbo Cheng, Yu Bao, Qian Cao, Luyang Huang, Liyan Kang, Zhicheng Liu, Yu Lu, Wenhao Zhu, Jingwen Chen, Zhichao Huang, Tao Li, Yifu Li, Huiying Lin, Sitong Liu, Ningxin Peng, Shuaijie She, Lu Xu, Nuo Xu, Sen Yang, and 7 others. 2025. Seed-X: Building strong multilingual translation LLM with 7B parameters. *arXiv [cs.CL]*.

- Jeongsoo Choi, Se Jin Park, Minsu Kim, and Yong Man Ro. 2024. Av2av: Direct audio-visual speech to audio-visual speech translation with unified audio-visual speech representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27325–27337.
- Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. 2023. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*.
- Alexis Conneau, Min Ma, Simran Khanuja, Yu Zhang, Vera Axelrod, Siddharth Dalmia, Jason Riesa, Clara Rivera, and Ankur Bapna. 2023. Fleurs: Few-shot learning evaluation of universal representations of speech. In *2022 IEEE Spoken Language Technology Workshop (SLT)*, pages 798–805. IEEE.
- Marta R Costa-Jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, and 1 others. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Cheng Cui, Ting Sun, Manhui Lin, Tingquan Gao, Yubo Zhang, Jiaxuan Liu, Xueqing Wang, Zelun Zhang, Changda Zhou, Hongen Liu, and 1 others. 2025. Paddleocr 3.0 technical report. *arXiv preprint arXiv:2507.05595*.
- Maha Elbayad, Laurent Besacier, and Jakob Verbeek. 2020. Efficient wait-k models for simultaneous machine translation. In *Interspeech 2020-Conference of the International Speech Communication Association*, pages 1461–1465.
- Desmond Elliott, Stella Frank, Loïc Barrault, Fethi Bougares, and Lucia Specia. 2017. [Findings of the second shared task on multimodal machine translation and multilingual image description](#). In *Proceedings of the Second Conference on Machine Translation*, pages 215–233, Copenhagen, Denmark. Association for Computational Linguistics.
- Desmond Elliott, Stella Frank, Khalil Sima'an, and Lucia Specia. 2016. [Multi30K: Multilingual English-German image descriptions](#). In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.
- Biao Fu, Donglei Yu, Minpeng Liao, Chengxi Li, Yidong Chen, Kai Fan, and Xiaodong Shi. 2025. Efficient and adaptive simultaneous speech translation with fully unidirectional architecture. *arXiv preprint arXiv:2504.11809*.
- Matthieu Futral, Cordelia Schmid, Ivan Laptev, Benoît Sagot, and Rachel Bawden. 2023. [Tackling ambiguity with images: Improved multimodal machine translation and contrastive evaluation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5394–5413, Toronto, Canada. Association for Computational Linguistics.
- Matthieu Futral, Cordelia Schmid, Benoît Sagot, and Rachel Bawden. 2025. [Towards zero-shot multimodal machine translation](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 761–778, Albuquerque, New Mexico. Association for Computational Linguistics.
- Marco Gaido, Sara Papi, Matteo Negri, and Luisa Bentivogli. 2024a. Speech translation with speech foundation models and large language models: What is there and what is missing? *arXiv [cs.CL]*.
- Marco Gaido, Sara Papi, Matteo Negri, and Luisa Bentivogli. 2024b. Speech translation with speech foundation models and large language models: What is there and what is missing? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14760–14778.
- Arushi Goel, Sreyan Ghosh, Jaehyeon Kim, Sonal Kumar, Zhifeng Kong, Sang-Gil Lee, Chao-Han Huck Yang, Ramani Duraiswami, Dinesh Manocha, Rafael Valle, and Bryan Catanzaro. 2025. Audio flamingo 3: Advancing audio intelligence with fully open large audio language models. *arXiv [cs.SD]*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. The llama 3 herd of models. *arXiv [cs.AI]*.
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2024. [xCOMET: Transparent machine translation evaluation through fine-grained error detection](#). *Transactions of the Association for Computational Linguistics*, 12:979–995.
- Shoutao Guo, Xiang Li, Mengge Liu, Wei Chen, and Yang Feng. 2025. Streamuni: Achieving streaming speech translation with a unified large speech-language model. *arXiv preprint arXiv:2507.07803*.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- J. Iranzo-Sánchez, J. A. Silvestre-Cerdà, J. Jorge, N. Roselló, A. Giménez, A. Sanchis, J. Civera, and A. Juan. 2020. Europarl-st: A multilingual corpus for speech translation of parliamentary debates. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233.

- Sai Koneru, Maike Züfle, Thai-Binh Nguyen, Seymanur Akti, Jan Niehues, and Alexander Waibel. 2025. KIT’s offline speech translation and instruction following submission for IWSLT 2025. *arXiv [cs.CL]*.
- Roman Koshkin, Katsuhito Sudoh, and Satoshi Nakamura. 2024a. [LLMs are zero-shot context-aware simultaneous translators](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1207, Miami, Florida, USA. Association for Computational Linguistics.
- Roman Koshkin, Katsuhito Sudoh, and Satoshi Nakamura. 2024b. [TransLLaMa: LLM-based simultaneous translation system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 461–476, Miami, Florida, USA. Association for Computational Linguistics.
- Jindong Li, Yali Fu, Jiahong Liu, Linxiao Cao, Wei Ji, Menglin Yang, Irwin King, and Ming-Hsuan Yang. 2025. Discrete tokenization for multimodal llms: A comprehensive survey. *arXiv preprint arXiv:2507.22920*.
- Danni Liu and Jan Niehues. 2024. Recent highlights in multilingual and multimodal speech translation. In *Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024)*, pages 235–253, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Danni Liu, Gerasimos Spanakis, and Jan Niehues. 2020. Low-latency sequence-to-sequence speech recognition and translation by partial hypothesis selection. In *21st Annual Conference of the International Speech Communication Association*, pages 3620–3624.
- Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. 2019. [STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3036, Florence, Italy. Association for Computational Linguistics.
- Xutai Ma, Javad Mohammad Dousti, Changhan Wang, Jiatao Gu, and Juan Pino. 2020. Simuleval: An evaluation toolkit for simultaneous translation. In *Proceedings of the EMNLP*.
- Sara Papi, Peter Polák, Dominik Macháček, and Ondřej Bojar. 2025a. How “real” is your real-time simultaneous speech-to-text translation system? *Trans. Assoc. Comput. Linguist.*, 13:281–313.
- Sara Papi, Marco Turchi, and Matteo Negri. 2023. Alignatt: Using attention-based audio-translation alignments as a guide for simultaneous speech translation. In *Proceedings of INTERSPEECH 2023*, pages 3974–3978.
- Sara Papi, Maike Züfle, Marco Gaido, Beatrice Savoldi, Danni Liu, Ioannis Douros, Luisa Bentivogli, and Jan Niehues. 2025b. [Mcif: Multimodal crosslingual instruction-following benchmark from scientific talks](#). Preprint, arXiv:2507.19634.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Supriti Sinhamahapatra and Jan Niehues. 2025. Do slides help? multi-modal context for automatic transcription of conference talks. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 16111–16121, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Lucia Specia, Stella Frank, Khalil Sima’an, and Desmond Elliott. 2016. [A shared task on multimodal machine translation and crosslingual image description](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553, Berlin, Germany. Association for Computational Linguistics.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). Preprint, arXiv:2503.19786.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv [cs.CL]*.
- Francesco Verdini, Pierfrancesco Melucci, Stefano Perna, Francesco Ciriaggi, Marco Gaido, Sara Papi, Szymon Mazurek, Marek Kasztelnik, Luisa Bentivogli, Sébastien Bratières, and 1 others. 2024. How to connect speech foundation models and large language models? what matters and what does not. *arXiv preprint arXiv:2409.17044*.
- Vipin Vijayan, Braeden Bowen, Scott Grigsby, Timothy Anderson, and Jeremy Gwinnup. 2024. [Adding multimodal capabilities to a text-only translation model](#).

In *Proceedings of the 16th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 14–28, Chicago, USA. Association for Machine Translation in the Americas.

André G Viveiros, Patrick Fernandes, Saul Santos, Sonal Sannigrahi, Emmanouil Zaranis, Nuno M Guerreiro, Amin Farajian, Pierre Colombo, Graham Neubig, and André F T Martins. 2025. TowerVision: Understanding and improving multilinguality in vision-language models. *arXiv [cs.LG]*.

Changhan Wang, Anne Wu, and Juan Pino. 2020. *Covost 2: A massively multilingual speech-to-text translation corpus*. Preprint, arXiv:2007.10310.

Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. 2025. Qwen2.5-omni technical report. *arXiv [cs.CL]*.

Hanrong Ye, Chao-Han Huck Yang, Arushi Goel, Wei Huang, Ligeng Zhu, Yuanhang Su, Sean Lin, An-Chieh Cheng, Zhen Wan, Jinchuan Tian, and 1 others. 2025. Omnivinci: Enhancing architecture and data for omni-modal understanding llm. *arXiv preprint arXiv:2510.15870*.

Jun Zhan, Junqi Dai, Jiasheng Ye, Yunhua Zhou, Dong Zhang, Zhigeng Liu, Xin Zhang, Ruibin Yuan, Ge Zhang, Linyang Li, Hang Yan, Jie Fu, Tao Gui, Tianxiang Sun, Yu-Gang Jiang, and Xipeng Qiu. 2024. *AnyGPT: Unified multimodal LLM with discrete sequence modeling*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9637–9662, Bangkok, Thailand. Association for Computational Linguistics.

A Appendix

A.1 CoMMuTE Manual Evaluation

To further validate whether the model is genuinely disambiguating examples in the CoMMuTE benchmark, we manually annotate the **first 50 source sentences from the English-to-German direction**, each paired with two images, resulting in a subset of 150 examples. We compare OmniFusion against TowerVision on this subset. The annotation is performed by a single annotator who inspects the source sentence together with the corresponding image. To reduce potential bias, we provide the annotator with the hypothesis files together with the disambiguation labels. For each example, we evaluate whether the system correctly disambiguates the sentence under both image conditions, while ignoring overall translation quality⁵.

⁵The source and reference sentences along with model hypothesis are submitted in data upload.

We find that, out of 50 source sentences, OmniFusion correctly disambiguates 18 examples, whereas TowerVision correctly disambiguates 23. Although the absolute scores are modest, the benchmark consists of highly challenging and synthetic examples. Nonetheless, it is still better than a text-only baseline as it will not disambiguate all and have a score of 0.

Moreover, while OmniFusion underperforms TowerVision on this targeted disambiguation analysis, it is important to note that OmniFusion also supports the audio modality and was not trained on the same scale of data as TowerVision.

This suggests that, despite lower performance on this image-focused disambiguation subset, OmniFusion retains broader multimodal applicability while requiring substantially less modality-specific training data.

A.2 Inference-Time Gate Control

Inference-time Gate Force	Speech	Speech + Image
First	55.22	63.02
Middle	Halluc. + ASR	Halluc. + ASR
Last	EOS	EOS
x	89.90	90.05

Table 5: Inference-time control of the gating module by forcing the first, middle, or last MMFM layer. We report XCOMET-XL scores on the MCIF en → de test set.

To further test the gate, we force the model to use only one MMFM layer at inference time and report results in Table 5. Forcing the first layer degrades quality but does not collapse the system, suggesting that early MMFM representations remain partially aligned with the translation LLM. In contrast, forcing the middle layer leads to hallucinated transcripts and ASR-like outputs, while forcing the last layer causes the model to emit only the eos token. These failure modes show that the gate does more than select a single best layer: it balances complementary information from multiple MMFM depths.

Together, the ablations and analysis support the main claim of OmniFusion. The model requires both an effective architectural interface and an appropriate training strategy: selected MMFM layers and gated fusion provide transferable multimodal representations, while OCR and self-cascading objectives make those representations useful for translation.

A.3 Attention Fusion

Attention-based layer fusion. As an alternative to the gated fusion module, we also consider a token-wise attention-based fusion mechanism over the selected MMFM layers. We evaluate this attention-based fusion using the first, middle, and last layers, and further analyze its behavior when applied to other selected layer sets. Below, we formalize the method using three layers as an example, although the formulation can be extended to an arbitrary number of selected layers.

For each multimodal token z_t , we first stack the three selected layer representations into a length-three layer sequence:

$$H_t = [h_t^{(\text{first})}; h_t^{(\text{mid})}; h_t^{(\text{last})}], \quad (8)$$

$$H_t \in \mathbb{R}^{3 \times D_{\text{MMFM}}}, \quad t = 1, \dots, L_m.$$

Let $\mathcal{J} = \{\text{first}, \text{mid}, \text{last}\}$ denote the set of selected layers. Then $h_t^{(j)}$ denotes the representation of token z_t from layer j , where $j \in \mathcal{J}$.

We compute a token-wise query vector from the mean of the selected layer representations:

$$\bar{h}_t = \frac{1}{|\mathcal{J}|} \sum_{j \in \mathcal{J}} h_t^{(j)}, \quad (9)$$

$$q_t = W_Q \bar{h}_t + b_Q,$$

$$q_t \in \mathbb{R}^{D_{\text{MMFM}}},$$

where $W_Q \in \mathbb{R}^{D_{\text{MMFM}} \times D_{\text{MMFM}}}$ and $b_Q \in \mathbb{R}^{D_{\text{MMFM}}}$ are learnable parameters.

The attention score for each selected layer is computed by a dot product between the query and the corresponding layer representation:

$$s_t^{(j)} = q_t^\top h_t^{(j)}, \quad j \in \mathcal{J}. \quad (10)$$

The scores are normalized across the selected layers:

$$\beta_t^{(j)} = \frac{\exp(s_t^{(j)})}{\sum_{\ell \in \mathcal{J}} \exp(s_t^{(\ell)})}, \quad j \in \mathcal{J}, \quad (11)$$

$$\beta_t = [\beta_t^{(j)}]_{j \in \mathcal{J}} \in \mathbb{R}^{|\mathcal{J}|}.$$

The final attention-fused MMFM representation is obtained as a weighted sum over the selected layer representations:

$$r_t = \sum_{j \in \mathcal{J}} \beta_t^{(j)} h_t^{(j)}, \quad (12)$$

$$r_t \in \mathbb{R}^{D_{\text{MMFM}}}, \quad t = 1, \dots, L_m.$$

The fused representation r_t is then projected into the translation-model embedding space using the same projection MLP:

$$v_t = \text{MLP}(r_t), \quad (13)$$

$$v_t \in \mathbb{R}^{D_{\text{NMT}}}, \quad t = 1, \dots, L_m.$$

A.4 Prompt Building

System Prompt for all modalities

You are Qwen, a virtual human developed by the Qwen Team, Alibaba Group, capable of perceiving auditory and visual inputs, as well as generating text and speech.

User Prompt (ST with/without Image)

Transcribe the audio using the image for context, and perform OCR on the image for correct spelling of keywords and names.

User Prompt (TIT)

Translate the following sentence into {tgt_lang} using the image for context. refers to the image. Use the image to determine formality, gender, keywords, and other details important for disambiguation: {ex[src]}

A.5 Layer Analysis, Zero-shot ability and APE - AudioTower

In our preliminary experiments, we focus exclusively on the audio modality and investigate fusing Qwen-Audio 7B (Chu et al., 2023) with Tower-Instruct 7B v0.2 (Alves et al., 2024). We train the fused model for ST using the same setup described earlier, but without gating, and we independently fuse the first, middle, and last layers of Qwen-Audio to assess their individual contributions. We then compare this fused model against several baselines. First, we evaluate both the out-of-the-box and fine-tuned versions of Qwen-Audio (Qwen2 Audio FT), trained end-to-end on the same ST data. We also include a standard cascaded setup using off-the-shelf ASR and MT models. Additionally, we evaluate an Automatic Post-Editing (APE) configuration in which Tower corrects translations generated by Qwen-Audio.

Algorithm 1 Online Prompt building for examples in the training data to enable alignment and self-cascading inference mode.

Require: Example ex

// SeedX prompt (source text, language tag and reference)

1: $src, tgt, img, aud, ocr, lang \leftarrow \text{EXTRACT}(ex)$

2: $suffix \leftarrow \langle \text{LANG}(lang) \rangle$

3: **if** aud **and not** img **then**

▷ Speech only

4: **if** $\text{RANDOM} < 0.1$ **then**

5: $inp \leftarrow \langle en \rangle$

▷ English-Audio only in our dataset

6: $lbl \leftarrow \text{ASR}(src) + \langle eos \rangle + suffix + tgt + \langle eos \rangle$

▷ ASR Self-cascade

7: **else**

8: $inp \leftarrow suffix$

9: $lbl \leftarrow tgt + \langle eos \rangle$

10: **end if**

11: **else if** img **and not** aud **then**

▷ Image only

12: **if** $\text{RANDOM} < 0.05$ **then**

13: $\text{SWAP}(src, tgt)$

14: $suffix \leftarrow \langle \text{English} \rangle$

15: **end if**

16: $src_m \leftarrow \text{MASK}(src, prob_words = 0.2, mask_chance = 0.2)$ ▷ Mask 20% source words for 20% of the time to pay attention to Image

17: $inp \leftarrow \langle \text{SRC} \rangle src_m + \langle eos \rangle + suffix$

18: $lbl \leftarrow tgt + \langle eos \rangle$

19: **else if** img **and** aud **then**

▷ Audio + Image

20: **if** $\text{RANDOM} < 0.8$ **then**

21: **if** $\text{RANDOM} < 0.1$ **then**

22: $inp \leftarrow \langle en \rangle$

23: $lbl \leftarrow \text{ASR}(src) + \langle eos \rangle + suffix + tgt + \langle eos \rangle$

24: **else**

25: $inp \leftarrow suffix$

26: $lbl \leftarrow tgt + \langle eos \rangle$

27: **end if**

28: **else**

29: $inp \leftarrow \langle \text{OCR} \rangle$

▷ OCR Self-cascade

30: $lbl \leftarrow ocr + \langle eos \rangle + suffix + tgt + \langle eos \rangle$

31: **end if**

32: **else**

33: **error** “no modalities”

34: **end if**

// Omni prompt (Multimodal prompt builder)

35: **if** aud **and** img **then**

36: $user \leftarrow \text{“Transcribe with OCR and image context:”} \parallel img \parallel aud$

37: **else if** aud **then**

38: $user \leftarrow \text{“Transcribe audio:”} \parallel aud$

39: **else**

40: $user \leftarrow \text{“Translate with image:”} \parallel img$

41: **end if**

42: $\text{STORE}(user, inp, lbl)$

Model	Mode	Hypothesis	Comet-22 (↑)					
			en → de	en→ zh	zh → de	es → it	ja → ru	fr → ko
Baselines								
Qwen2 Audio	E2E	—	81.24	78.20	73.59	83.38	70.70	71.26
Qwen2 Audio FT	E2E	—	84.82	85.57	80.22	84.17	67.69	77.09
Qwen2 Audio + Tower	Cascade	—	85.85	85.97	82.94	85.36	69.26	84.78
Tower APE	Cascade APE	Qwen2 Audio + Tower	85.96	85.77	82.59	85.21	47.43	70.67
Off-the-shelf ST models								
Tower-Spire	E2E	—	82.77	81.63	—	—	—	—
Tower-Spire	Self-Cascade	—	85.17	85.73	—	—	—	—
SeamlessM4T v2	E2E	—	85.63	80.05	70.22	77.78	70.58	78.91
E2E models								
AudioTower-Zero	E2E	—	85.13	83.97	76.62	83.98	68.71	81.64
AudioTower-Mid	E2E	—	86.02	85.03	80.62	84.86	74.47	83.26
AudioTower-Last	E2E	—	82.97	81.03	54.12	49.57	50.19	52.31
E2E + Speech-Aware APE model								
AudioTower-Mid	Cascade APE	Qwen2 Audio + Tower	86.34	85.99	83.42	85.96	78.26	84.59
AudioTower-Mid	E2E	Empty	85.77	84.79	80.39	85.27	74.67	82.71

Table 6: Offline ST results on FLEURS test set. Speech-Aware APE models use the hypothesis from a system and perform post-editing. Statistically significant top-performing results per language pair are highlighted in **bold**. Underlined languages are not seen during the fusing process as source/target.

Dataset	# Examples	Language	Modality	Task
M3AV	122k	<i>en</i>	Speech, Image	ASR, OCR
M3AV - Aug	122k	<i>zh, fr, de, it, ja, ko, pt, ru, es, vi, ar, cs, hr, da, nl, fi, hu, id, ms, nb, no, pl, ro, tr</i>	Speech, Image	ST, OCR
Europarl	260k*	<i>de, es, fr, it, nl, pl, pt, ro</i>	Speech	ST
Covost2		<i>de, ja, tr, zh</i>	Speech	ST
Multi30k	87k	<i>cs, de, fr</i>	Image	TIT

Table 7: Overview of data sources collected for E2E training. * We sample 260k examples uniformly per language and dataset (260k / 24), with German appearing more frequently due to its presence in both datasets. All the speech and text data are only available with English as the source language. Language code mapping is presented in Table 13.

To contextualize performance, we further compare against other off-the-shelf ST systems such as Tower-Spire (Ambilduke et al., 2025) and SeamlessM4T (Barrault et al., 2023), which provides insight into multilingual generalization. Beyond the main fused model, we also train an APE variant of our fusion architecture that post-edits the outputs of cascaded systems. Unlike traditional APE, this Speech-Aware APE allows the model to use the source audio when correcting errors, enabling more informed revisions. For both APE variants, we generate synthetic APE tuples by running Qwen-Audio on the training data and fine-tune with LoRA.

We report results on the FLEURS test set in Table 6, including evaluations on zero-shot language

Model	Time(s)	Image	XCOMET-XL
Omni + SeedX FT	2.96	w/o	89.75
OmniFusion	1.98	w/o	88.09
OmniFusion SC	3.78	w/o	88.44
Omni + SeedX FT	3.85	w/	90.18
OmniFusion	3.15	w/	89.90
OmniFusion SC	4.82	w/	90.05

Table 8: Quality vs latency tradeoffs in the offline setting for English $\rightarrow$ German on the MCIF test set. SC indicates Self-Cascade mode.

directions not present on either the source or target side during training. First, while fine-tuning improves Qwen2 Audio, it still falls short of the out-of-the-box cascaded setup across all languages, underscoring the strength of translation-centric LLMs. Tower-based APE is competitive for seen languages but struggles in harder zero-shot directions, such as *fr* $\rightarrow$ *kr*. Off-the-shelf end-to-end ST models lag significantly behind cascaded systems as well.

Our proposed fused model (AudioTower) with middle-layer fusion achieves the strongest overall performance and even surpasses the cascaded baseline in several languages. Notably, both the first and middle layers show strong transfer to unseen languages, whereas the last layer collapses entirely—supporting the hypothesis that upper layers encode highly language-specific features, making them less suitable for cross-lingual contextual embeddings.

Finally, the APE variant of the fused model delivers the best results overall by combining the

Model	Mode	Hypothesis	Comet-22 ($\uparrow$)					
			en $\rightarrow$ de	en $\rightarrow$ zh	zh $\rightarrow$ de	es $\rightarrow$ it	ja $\rightarrow$ ru	fr $\rightarrow$ kr
Qwen2 Audio + Tower	Cascade	—	85.85	85.97	82.94	85.36	69.26	84.78
Whisper + Tower	Cascade	—	86.71	86.31	82.23	86.54	83.56	85.29
AudioTower-Mid	Cascade APE	Qwen2 Audio + Tower	86.34	85.99	83.42	85.96	78.26	84.59
AudioTower-Mid	Cascade APE	Whisper + Tower	87.09	86.34	82.81	86.60	84.38	84.95
Character-Error-Rate ($\downarrow$)								
Whisper	ASR	—	3.14	3.15	25.37	2.11	7.67	3.28

Table 9: Analysis of system combination by using Whisper to generate the hypothesis for post-editing on FLEURS test set. Statistically significant top-performing results per language pair are highlighted in **bold**. Underlined languages are not seen during the fusing process as source/target.

strong hypotheses of the cascaded pipeline with the ability to exploit the raw audio when correcting errors. We also test this APE model with empty hypotheses and observe performance comparable to AudioTower-Mid, indicating its flexibility for multiple use cases.

Given the strong APE results, we additionally explore whether the model can correct outputs from different ASR systems. Using Whisper (Radford et al., 2023) for ASR, we evaluate the fused APE model on Whisper+Tower outputs and report results in Table 9. When Whisper+Tower produce the strongest cascaded baseline, using them as hypotheses likewise yields the highest APE performance—highlighting that Speech-Aware APE can be beneficial not only for E2E setups but also for offline cascaded scenarios.

A.6 Latency Reproduction

We use the SimulST toolkit to calculate Average Lagging for both computation-aware and computation-unaware settings. For a given chunk size, we iteratively feed the system audio segments of that duration and apply the local agreement algorithm. At each step, we commit any finalized words and record both the committed output and the amount of audio consumed.

In the E2E setup, local agreement is applied directly to the translation output. In the cascade setup, however, we first wait for the ASR output to stabilize before applying local agreement to the MT stage. All experiments were run with batch size 1 on an NVIDIA RTX A6000 GPU to ensure reproducibility.

A.7 Decoding

All the inferences are done by setting max_new_tokens to 256 and beam size with num_beams set to 5. The same decoding settings

were also followed when evaluating other models. All the inferences ran on 1 NVIDIA RTX A6000 GPU.

Category	Model Name	Image	XCOMET-XL	Minor	Major	Critical	Total
Omni FT (Direct ST)	LoRA FT	✗	85.90	2047	166	173	2386
Cascade Systems (ASR → MT) Omni + SeedX	No FT	✗	89.66	2172	132	68	2372
	No FT	✓	89.49	2115	133	76	2324
	LoRA FT	✗	89.75	2209	128	67	2404
	LoRA FT	✓	90.18	2135	118	56	2309
End-to-End (Direct ST) OmniFusion	Mid Fusion	✗	88.44	2126	157	74	2357
	Mid Fusion	✓	90.03	2256	104	68	2428
	Gated Fusion	✗	88.09	2133	181	87	2401
	Gated Fusion	✓	89.90	2217	106	60	2383
	Gated Fusion + Self-Cascade	✓	90.05	2191	94	62	2347

Table 10: Offline ST results en→de on the MCIF test. ✓ indicates models using image input. Bold indicates best XCOMET-XL and lowest error counts.

Category	Model Name	Image	XCOMET-XL	Minor	Major	Critical	Total
Omni FT (Direct ST)	LoRA FT	✗	80.88	727	1018	54	1799
Cascade Systems (ASR → MT) Omni + SeedX	No FT	✗	86.03	850	932	51	1833
	No FT	✓	86.11	862	891	67	1820
	LoRA FT	✗	86.76	884	855	54	1793
	LoRA FT	✓	86.85	868	855	55	1778
End-to-End (Direct ST) OmniFusion	Mid Fusion	✗	82.80	754	1065	93	1912
	Mid Fusion	✓	85.75	899	903	64	1866
	Gated Fusion	✗	84.06	745	1085	56	1886
	Gated Fusion	✓	86.79	914	847	48	1809
	Gated Fusion + Self-Cascade	✓	87.20	879	838	44	1761

Table 11: Offline ST results en→it on the MCIF test. ✓ indicates models using image input. Bold indicates best XCOMET-XL and lowest error counts.

Category	Model Name	Image	XCOMET-XL	Minor	Major	Critical	Total
Omni FT (Direct ST)	LoRA FT	✗	80.13	463	1372	44	1879
Cascade Systems (ASR → MT) Omni + SeedX	No FT	✗	81.94	615	1255	104	1974
	No FT	✓	82.23	597	1254	81	1932
	LoRA FT	✗	82.74	624	1271	104	1999
	LoRA FT	✓	82.74	604	1326	56	1986
End-to-End (Direct ST) OmniFusion	Mid Fusion	✗	79.47	462	1343	50	1855
	Mid Fusion	✓	82.15	614	1240	53	1907
	Gated Fusion	✗	79.78	480	1298	45	1823
	Gated Fusion	✓	82.03	636	1255	63	1954
	Gated Fusion + Self-Cascade	✓	82.47	623	1225	60	1908

Table 12: Offline ST results en→zh on the MCIF test. ✓ indicates models using image input. Bold indicates best XCOMET-XL and lowest error counts.

Language	Code
Arabic	<i>ar</i>
Chinese	<i>zh</i>
Croatian	<i>hr</i>
Czech	<i>cs</i>
Danish	<i>da</i>
Dutch	<i>nl</i>
Finnish	<i>fi</i>
French	<i>fr</i>
German	<i>de</i>
Hungarian	<i>hu</i>
Indonesian	<i>id</i>
Italian	<i>it</i>
Japanese	<i>ja</i>
Korean	<i>ko</i>
Malay	<i>ms</i>
Norwegian	<i>no</i>
Norwegian Bokmål	<i>nb</i>
Polish	<i>pl</i>
Portuguese	<i>pt</i>
Romanian	<i>ro</i>
Russian	<i>ru</i>
Spanish	<i>es</i>
Turkish	<i>tr</i>
Urdu	<i>ur</i>
Vietnamese	<i>vi</i>

Table 13: Language codes denoted in our paper following [Cheng et al. \(2025\)](#).

Hyperparameter	Value
Gradient checkpointing	True
Gradient checkpointing kwargs	{'use_reentrant': False}
Save steps	1000
Evaluation steps	1000
Save strategy	steps
Evaluation strategy	steps
Logging steps	5
Logging strategy	steps
Learning rate	1e-4
Train batch size (per device)	4
Eval batch size (per device)	4
Gradient accumulation steps	2
Weight decay	0.01
Save total limit	3
Max steps	20000
bf16	True
peft type	LORA
qalora_group_size	16
r	16
target_modules	q,k,v,up,down,gate projection
Push to hub	False
Metric for best model	loss
Remove unused columns	False
Label names	["labels"]
Dataloader num workers	1
DDP find unused parameters	True

Table 14: Huggingface trainer hyperparameters used for the model on 4 Nvidia A100 48GB GPUs with DDP training mode. Unspecified hyperparameters are set to default

Image	Cascaded FT Output	OmniFusion Output	Reference
	Hallo, willkommen zu unserer Präsentation von Deplain, einem neuen Korpus für die Identifizierung deutscher Texte auf Dokumenten- und Satzebene.	Hallo, willkommen zu unserer Präsentation von DEPLAIN, einem neuen Korpus für die deutsche Texterkennung auf Dokumentebene und auf Satzebene.	Herzlich willkommen zu unserer Präsentation von DEPLAIN, einem neuen Korpus für die deutsche Texterkennung auf Dokument- und Satzebene.

Table 15: Example illustrating where OmniFusion is better than cascaded variant. Example taken from MCIF segment 1.

Image	SeedX Output	OmniFusion Output	Reference
	AUSGANG	AUSFAHRT	AUSFAHRT
	AUSGANG	AUSGANG	AUSGANG

Table 16: Example illustrating OmniFusion disambiguates source sentence with image context. The word 'EXIT' when translated to German can be 'AUSFAHRT' (as vehicle exit) or 'AUSGANG' (as person exit) and needs the image for disambiguation.