

ViDove: A Translation Agent System with Multimodal Context and Memory-Augmented Reasoning

Yichen Lu^{1,2*†}, Wei Dai^{1,4*†}, Jiaen Liu^{1,8*†}, Ching Wing Kwok^{1,3*},
Zongheng Wu^{1,5*}, Xudong Xiao^{1,7}, Ao Sun⁹, Sheng Fu¹,
Jianyuan Zhan⁷, Yian Wang⁶, Takatomo Saito², Sicheng Lai^{1†},

¹Pigeon AI, ²Carnegie Mellon University, ³Fudan University,

⁴University of California San Diego, ⁵University of Toronto, ⁶University of California Irvine,

⁷University of Illinois Urbana-Champaign, ⁸Institute of Science Tokyo,

⁹Hong Kong University of Science and Technology

Abstract

LLM-based translation agents have achieved highly human-like translation results and are capable of handling longer and more complex contexts with greater efficiency. However, they are typically limited to text-only inputs. In this paper, we introduce **ViDove**, a translation agent system designed for multimodal input. Inspired by the workflow of human translators, ViDove leverages visual and contextual background information to enhance the translation process. Additionally, we integrate a multimodal memory system and long-short term memory modules enriched with domain-specific knowledge, enabling the agent to perform more accurately and adaptively in real-world scenarios. As a result, ViDove achieves significantly higher translation quality in both subtitle generation and general translation tasks, with a 28% improvement in BLEU scores and a 15% improvement in SubER compared to previous state-of-the-art baselines. Moreover, we introduce **DoveBench**, a new benchmark for long-form automatic video subtitling and translation, featuring 17 hours of high-quality, human-annotated data. Our project is shown at <https://github.com/pigeonai-org/ViDove>.

1 Introduction

Recent advances in Large Language Models (LLMs) have demonstrated remarkable capabilities in Machine Translation (MT) tasks (Robinson et al., 2023; Gao et al., 2023a; Xu et al., 2024a; Zhu et al., 2024). The integration of autonomous agent frameworks with LLM-based MT has shown promising results in enhancing translation capabilities, pushing modern translation systems closer to human-level professional performance (Wu et al., 2024a; Wang et al., 2025a; Guo et al., 2024a; Peter et al., 2024). Through the incorporation of long-short memory system and multi-

agent strategies, these approaches have achieved significant improvements in both translation quality and efficiency, enabling LLM-based MT to handle document-level translation effectively (Wang et al., 2025a).

Professional human translators often rely on more than just text to ensure accurate translations. For example, in a cooking video, “fold” could mean combining ingredients or folding a napkin—visual cues like stirring motions and ingredients clarify meaning, while audio tone and emphasis convey intent and emotion (Sulubacak et al., 2019; Shen et al., 2024; et al., 2025). While some Multimodal Machine Translation (MMT) studies incorporate limited visual or audio inputs, they typically cannot handle document-level translation or take entire video as input (Lu et al., 2025; Lv et al., 2025). However, most existing LLM-based MT systems focus solely on textual input, missing out on these valuable contextual signals.

To close the gap between LLM-based translation and professional human performance, we present **ViDove**, a multimodal translation agent that integrates visual, audio, and textual inputs. Built on Retrieval-Augmented Generation (RAG) (Arslan et al., 2024; Abootorabi et al., 2025; Zhai, 2024) and recent Multimodal LLMs (MLLMs) advances (Lu et al., 2024a; Xu et al., 2025; Lu et al., 2024b; Chu et al., 2024), ViDove uses a memory system for domain-specific knowledge, multimodal context, and instruction customization (Long et al., 2024; Ding et al., 2024). Specialized agents handle these inputs to produce more accurate, human-like translations and subtitles. ViDove achieves state-of-the-art results, with a **28%** BLEU and **15%** SubER improvement over baselines. We also introduce **DoveBench**, a video automatic subtitling and translation benchmark with 17 hours of high-quality human-annotated subtitles to support future research.

The key innovations of our work include:

*Equal Contribution

†Project Lead

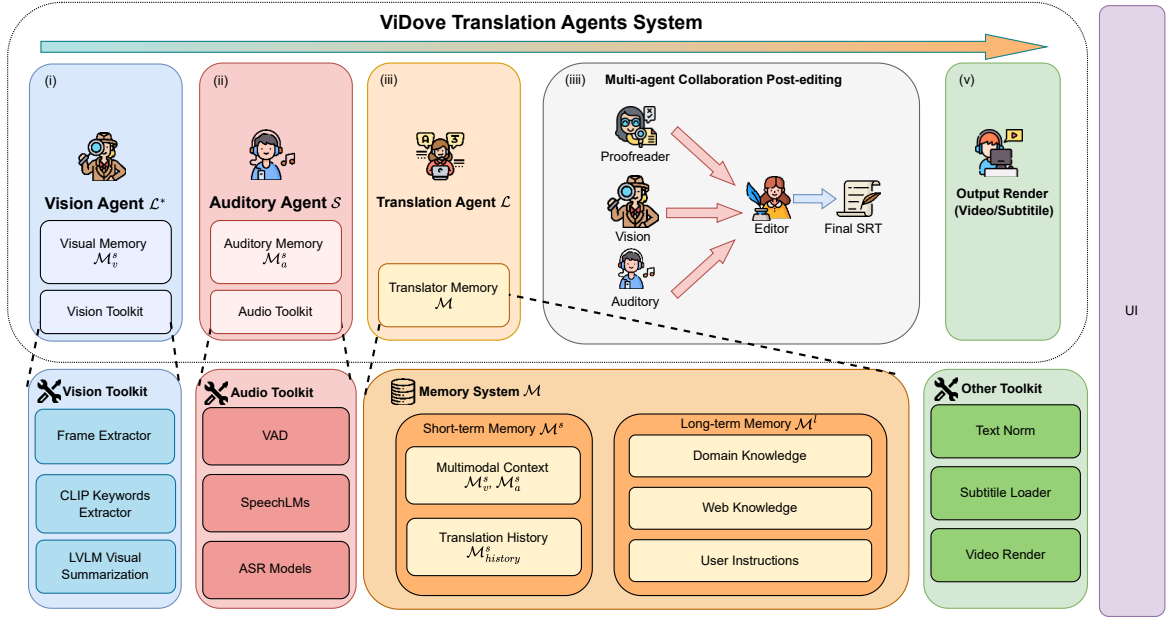


Figure 1: **Architecture of the ViDove Translation Agents System.** The system consists of five modules: (i) Vision Agent \mathcal{L}^* and (ii) Auditory Agent \mathcal{S} extract multimodal cues; (iii) Translation Agent \mathcal{L} utilizes memory $\mathcal{M} = \{\mathcal{M}^s, \mathcal{M}^l\}$ for context-aware translation; (iv) a multi-agent post-editing module refines the output via collaboration; (v) Output Render generates final subtitles and video.

- **Multimodal Multi-Agent Collaboration.** A modular translation agent system that simulates human translator workflows by integrating audio, visual, and textual modalities through specialized agents and their interactions, achieving performance comparable to or better than existing baseline systems.
- **Memory-Augmented Reasoning.** A long-short term memory system for managing multimodal and domain-specific knowledge during translation.
- **DoveBench.** A long-form video automatic subtitling and translation benchmark that reflects real-world subtitling challenges.

2 Related Works

ViDove is a multimodal translation agent framework that enables cooperative translation through multimodal grounding and memory-guided reasoning. Our work builds on two main areas: multimodal machine translation systems and multi-agent, whose integration remains underexplored.

Machine Translation: Traditional MT systems (Wu et al., 2016; Team and et al., 2022) perform well at the sentence level but struggle with limited contextual cues, particularly in multimodal scenarios. To address this, prior studies (Li et al.,

2022; Zuo et al., 2023; Lin et al., 2020; Lan et al., 2023) have introduced visual grounding, yet remain constrained to sentence-level tasks with limited context handling. LLM-based MT (Robinson et al., 2023; Hendy et al., 2023; He et al., 2024; Jiao et al., 2023) achieves strong performance by leveraging longer context, but typically treats LLMs as black-box translators rather than reasoning agents with human-like interpretive capabilities.

RAG for Machine Translation: Recent studies have increasingly leveraged RAG to enhance MT. Some works applied RAG to improve the quality and cultural grounding of MT by utilizing diverse retrieval pipelines, knowledge graphs, and multi-task fine-tuning setups (Bouthors et al., 2024; Conia et al., 2024; Wang et al., 2024; Anonymous, 2024). Concurrently, other works have specifically addressed low-resource languages, showing significant gains by augmenting prompts with retrieved bilingual dictionaries and example sentences (Merx et al., 2024; Chang et al., 2025). ViDove effectively combines the strengths of these approaches by architecting a multi-agent system where RAG serves as the core information storage and exchange protocol.

Multi-Agent Systems and Translation Agent: Recent work on multi-agent LLM systems (Li et al., 2023a; Hu et al., 2021; Guo et al., 2024b; Cheng

et al., 2024; Li et al., 2023c; Ma et al., 2024; Wang et al., 2023; Xu et al., 2024b) shows that dividing complex tasks among specialized agents with distinct roles enhances reasoning and decision-making. These agents collaborate through structured coordination, often leveraging tools (Li et al., 2023b; Ruan et al., 2023; Wu et al., 2023) and memory systems (Ding et al., 2023; Singhal et al., 2023a,b) to maintain context and task knowledge.

This paradigm has recently been explored in machine translation. For example, TransAgents (Wu et al., 2024a) improves translation quality through multi-agent critique, where agents evaluate and refine each other’s outputs. Delta (Wang et al., 2025b) ensures document-level consistency using memory modules. Both approaches enhance translation quality in domain-specific or long-form scenarios.

However, these existing approaches remain confined to the textual modality and overlook the potential of MLLMs (Lu et al., 2024a; Xu et al., 2025; Chu et al., 2024) in enhancing translation quality. To bridge this gap and bring machine translation closer to human-like performance, we propose ViDove, a novel framework that leverages recent advances in both LLM-based agent systems and MLLMs.

3 ViDove

In this section, we first introduce the characteristics of long-form video subtitle generation and translation, then describe each agent in Fig. 1 in detail.

3.1 Preliminary

For clarity, we define our primary notation as follows: \mathcal{V} denotes the input video, \mathcal{L} is the translator Agent, \mathcal{L}^* represents the visual agent, and \mathcal{S} signifies the auditory agent. Long-short term memory modules are represented by $\mathcal{M} = \{\mathcal{M}^s, \mathcal{M}^l\}$, capturing both short-term and long-term contexts. The prompt list P contains guiding prompts for analysis tasks. Video chunks are represented by \mathcal{C}_i , each consisting of visual (\mathcal{V}_i) and audio (\mathcal{A}_i) components. Transcripts are denoted by T_i , with translated transcripts represented by T_i^* . The pipeline framework can be formulated as algorithm 1.

3.2 Long-form Video Automatic Subtitling

Long-form(>10min) video automatic subtitling requires the system to process a video \mathcal{V} , which includes a synchronized audio stream \mathcal{A} . Unlike

Algorithm 1 ViDove

Require: Input video \mathcal{V} , Multi-agent translation agent \mathcal{L} as in Algorithm 2, Visual Agent \mathcal{L}^* , auditory agent \mathcal{S} , Long-short term memory $\mathcal{M} = \{\mathcal{M}^s, \mathcal{M}^l\}$, prompt list P

Ensure: Translated video, original and translated transcript tuple (V^*, T^*, T)

Initialize: $\mathcal{M}^s \leftarrow \emptyset, \mathcal{M}^l \leftarrow$ Knowledge Base \triangleright Initialize memory modules

$\{\mathcal{V}_i, \mathcal{A}_i\}_{i=1}^k \leftarrow \mathcal{C}(\mathcal{V}) \triangleright$ Chunk decomposition

for each chunk $(\mathcal{V}_i, \mathcal{A}_i) \in \mathcal{C}$ **do**

$(\mathcal{M}_v^s, cue_v) \leftarrow \mathcal{L}^*(\mathcal{V}_i, \mathcal{M}_v^s, \mathcal{M}_{domain}^l, p_{analysis}), p_{analysis} \in P \triangleright$ Update visual cues and short-term memory

$cue_a \leftarrow \mathcal{S}(\mathcal{A}_i)$

$T_i \leftarrow (cue_a, cue_v)$

$(T_i^*, \mathcal{M}_{history}^s) \leftarrow \mathcal{L}(T_i, \mathcal{M}, p_{translation})$

end for

$T^*, T \leftarrow$ **Multi-agentPostProcess** $(T^*, T, \mathcal{M}, p_{pr})$

\triangleright Final post-processing of translations

return (V^*, T^*, T)

typical sentence-level translation, this task operates at the document level and demands precise alignment of translated subtitles with the corresponding timestamps. Moreover, in contrast to standard document-level MT, the input does not contain any original text. Instead, the system must first perceive the video by “listening” to the audio—and, when necessary, “observing” the visual content—to transcribe the source language before translating it. These multimodal and multi-stage requirements—speech recognition, visual grounding, and context-aware translation—make the task considerably more complex than traditional MT, multimodal MT, or document-level MT (DocMT). Addressing this challenge calls for a holistic, agent-based approach capable of perception, reasoning, and translation.

3.3 Auditory Agent

In this section, we describe the workflow of ViDove’s auditory agent, which provides audio-based contextual information and enriched transcriptions to support the Translation Agent.

Step 1: Chunk Splitting and Timestamp Extraction. We use Pyannote (Plaquet and Bredin, 2023a) to segment the input video \mathcal{V} into k chunks, \mathcal{C} , based on speaker activity. Each chunk \mathcal{C}_i contains a corresponding set of video frames \mathcal{V}_i and an audio sequence \mathcal{A}_i .

Step 2: Auditory Information Extraction. ViDove integrates SOTA single-task models to extract key auditory features: speech transcription (Radford et al., 2022), background audio event detection (Chen et al., 2024, 2022), and speaker emotion recognition (Ma et al., 2023). In addition to these, we incorporate recent Speech Language Models (SpeechLMs) (Tang et al., 2024; Chu et al., 2024), which offer a unified approach for extracting rich audio cues. For each chunk, the extracted auditory information, denoted as cue_a , is stored in the multimodal contextual memory $\mathcal{M}_a^s \in \mathcal{M}_{multimodal}^s$ to support downstream translation.

Step 3: Audio Transcription and Timestamp Refinement. ViDove’s auditory agent uses either a SpeechLM or an ASR model to transcribe the audio sequence \mathcal{A}_i , leveraging the multimodal contextual memory $[\mathcal{M}_a^s, \mathcal{M}_v^s] \in \mathcal{M}_{multimodal}^s$. This fusion enables enhanced transcription through keyword injection and agent-based reasoning.

3.4 Vision Agent

ViDove implements a visual agent to gather informative visual cues from video input \mathcal{V} . These cues are used to enhance speech recognition (Lu et al., 2024a; Wu et al., 2024b) and are passed through a memory system to support more sophisticated sentence comprehension. This visual cue allows the agent to resolve textually ambiguousness and accurately interpret domain-specific terminology during translation. The ViDove system is designed to be model-agnostic and supports multiple vision-language backends, offering deployment flexibility across local and remote environments.

The pipeline is designed to process pre-segmented video chunks

$$\mathcal{V} = \{\mathcal{V}_0, \mathcal{V}_1, \dots, \mathcal{V}_k\}, k = |\mathcal{V}|$$

, from which key frames are extracted to represent salient visual moments. These frames are then analyzed by the selected vision-language model \mathcal{L}^* , providing a high-level semantic understanding cue_v of the visual context.

$$cue_v = \mathcal{L}^*(v_i, \mathcal{M}_{vision}^s, \mathcal{M}_{domain}^l, P_{analysis})$$

Where \mathcal{M} is multi-modal memory for ViDove system.

3.5 Memory system

ViDove’s memory system, denoted as $\mathcal{M} = \{\mathcal{M}^s, \mathcal{M}^l\}$, is implemented using LLaMA-

Index (Liu, 2022). This memory system stores and organizes multimodal information, enabling the system to deliver contextually informed and consistent translations.

3.5.1 Short-term Memory (\mathcal{M}^s)

The short-term memory (\mathcal{M}^s) is tailored to the current video translation task. It holds information specific to each video chunk (\mathcal{C}_i), which includes visual (\mathcal{V}_i) and audio (\mathcal{A}_i) components. Its contents include:

Translation History: Records of prior translations within the same video, ensuring consistency in terminology and phrasing across transcripts (e.g., from T_i to T_i^*).

Visual Cues: Contextual data extracted from \mathcal{V}_i , such as scene descriptions or objects, that helps to disambiguate textual content.

Audio Cues: Auditory information extracted from \mathcal{A}_i , denoted as cue_a , including transcriptions and other speaker information, stored in the multimodal contextual memory $\mathcal{M}_a^s \in \mathcal{M}_{multimodal}^s$.

This component provides immediate context, supporting accurate and coherent translations within a single video.

3.5.2 Long-term Memory (\mathcal{M}^l)

The long-term memory (\mathcal{M}^l) is designed for adaptability across multiple translation tasks. It accumulates knowledge over time, storing:

Domain Knowledge: This type of knowledge captures specialized community language to ensure accurate video translations for a diverse, multilingual audience.

Web Knowledge: General information sourced from the web, implemented by Tavily (Tavily AI, 2025), offering broader context.

This component enhances ViDove’s flexibility, enabling it to handle diverse domains and improve performance over time. The memory system acts as a centralized repository, providing essential information that supports the translation process by ensuring consistency and contextual relevance.

3.6 Multi-agent Translation

ViDove’s translation process relies on a multi-agent system featuring three specialized agents—*Translator*, *Proofreader*, and *Editor*—that work together to produce high-quality subtitle translations. These agents collaborate by accessing the unified memory system $\mathcal{M} = \{\mathcal{M}^s, \mathcal{M}^l\}$, ensuring consistency and context throughout the workflow.

| | DoveBench | | | | BigVideo | |
|------------------|--------------|--------------|--------------|-------------|--------------|-------------|
| | BLEU(↑) | BLEURT(↑) | SubER(↓) | SubSONAR(↑) | BLEU(↑) | sCOMET(↑) |
| Gemini-2.5-Flash | 8.11 | 17.21 | 103.46 | 0.31 | 26.43 | 0.75 |
| Qwen-2.5-Omni | 14.60 | 13.83 | 108.94 | 0.39 | 10.67 | 0.58 |
| VideoCaptioner | 12.65 | 14.62 | 85.75 | 0.41 | 30.36 | 0.75 |
| Whisper + DelTA | 18.26 | 12.30 | 86.83 | 0.28 | 29.09 | 0.69 |
| ViDove | 23.51 | 19.55 | 73.38 | 0.39 | 26.05 | 0.73 |

Table 1: ViDove compared with different baseline system on DoveBench and BigVideo.

Agents and Their Roles

Translator Agent (\mathcal{L}_t): This agent generates the initial translation (T_i^*). To achieve this, it interacts with the memory system by retrieving *translation history* and *visual cues* from short-term memory (\mathcal{M}^s) to maintain intra-video consistency, while simultaneously drawing upon *domain knowledge* from long-term memory (\mathcal{M}^l) to ensure domain-specific accuracy and stylistic alignment from the outset.

Proofreader Agent (\mathcal{L}_{pr}): Proofreader agent focuses on refining the initial translation by correcting grammar, style, and terminology. It interacts with memory by accessing *domain knowledge* from long-term memory (\mathcal{M}^l) to ensure linguistic precision and adherence to specialized terminology, while referencing the *translation history* in short-term memory (\mathcal{M}^s) to maintain consistency with prior segments. The proofreader agent generates revision suggestions for the editor agent, which makes the final decision on whether to apply them. A sample interaction log is provided in Appendix A.1.1.

Editor Agent (\mathcal{L}_{ed}): The editor agent performs the final quality check and applies necessary modifications to ensure the translation quality with the full modality context. It receives suggestions from the proofreader agent and decides whether to adopt them. It also accepts user instructions, enabling more free-form and practical interactions between the user and agents. To verify contextual accuracy, the editor agent leverages the memory system by retrieving *visual cues*, *audio cues*, and *translation history* from the short-term memory (\mathcal{M}^s). It also queries *web knowledge* from the long-term memory (\mathcal{M}^l) to incorporate broader external context and ensure logical consistency.

4 DoveBench

To the best of our knowledge, there is currently no standardized open-sourced benchmark for long-form video automatic subtitling (Sec 3.2). To address this gap, we introduce **DoveBench**, an open-source benchmark designed specifically for this task. DoveBench contains approximately 17 hours of video data, each annotated with Chinese (ZH) subtitles translated by professional translators. The average video length is around 20 minutes, reflecting typical durations found in real-world scenarios. Detailed statistics of DoveBench are provided in Appendix B.

5 Experiments

In this section, we first describe the evaluation datasets, baseline systems, metrics, and ViDove’s configuration. We then present and analyze experiment results.

5.1 Datasets and Metrics

We evaluate long-form video automatic subtitling on DoveBench and MMT on BigVideo (Kang et al., 2023). To assess translation quality, we use BLEU (Freitag et al., 2020), BLEURT (Selam et al., 2020), and sCOMET (Rei et al., 2020). For subtitle quality—capturing both translation accuracy and timestamp alignment—we adopt SubER (Wilken et al., 2022) and SubSONAR (Gaido et al., 2024), specifically for evaluating on DoveBench.

5.2 Baselines and ViDove Configuration

We compare ViDove against four baseline systems. Gemini-2.5-flash (Google DeepMind and Google Research, 2025) serves as a proprietary baseline, and Qwen-2.5-Omni (Xu et al., 2025) represents an open-source alternative. Both models are single MLLMs capable of processing video input

and generating subtitle (SRT) output through carefully designed prompts. For system-level baselines, we include VideoCaptioner (Weifeng2333, 2024), an open-source cascaded pipeline for video subtitling, and DelTA (Wang et al., 2025a), a state-of-the-art text-based translation agent. Since DelTA does not support audio or video input natively, we use whisper-large-v3 (Radford et al., 2022) to first transcribe the audio for the system. For ViDove, we use Gemini-2.5-flash (Google DeepMind and Google Research, 2025) as the auditory agent, while all other agents are powered by GPT-4o (OpenAI, 2024). Note that neither the baseline models nor ViDove undergo any additional fine-tuning during the experiments. Detailed prompts for the baseline models and ViDove are provided in Appendix C.1, C.2 and A.3.

5.3 Experiment Results

Table 1 presents the evaluation results of ViDove and baseline systems on both DoveBench and BigVideo.

On DoveBench, ViDove consistently outperforms all baselines across all metrics, achieving the highest BLEU (23.51), BLEURT (19.55), and the lowest SubER (73.38). Compared to the strongest baseline, Whisper + DelTA, ViDove improves BLEU by 28.8%, BLEURT by 58.9%, and reduces SubER by 15.5%. These results indicate that ViDove not only produces more accurate translations but also aligns subtitles more precisely in time. However, despite ViDove’s leading performance, the absolute scores—especially BLEU and SubER—remain relatively modest. This highlights the intrinsic difficulty of long-form video automatic subtitling. To date, no existing system has achieved fully satisfactory results on this task, underscoring its complexity and open research nature.

In contrast, Gemini-2.5-flash and Qwen-2.5-Omni perform poorly on DoveBench, with high SubER values (103.46 and 108.94, respectively) and low BLEU scores. Although we carefully engineered prompts and applied post-processing to ensure fair evaluation, these single-model MLLMs struggle on long-form inputs. Their limited instruction-following capability, combined with a tendency to hallucinate or ignore constraints as input length increases, leads to misaligned, incomplete, or off-topic subtitles—even when the task is clearly specified.

On BigVideo, which focuses on sentence-level MMT, ViDove remains competitive. While it is

| Model | BLEU (↑) | SubER (↓) | BLEURT (↑) |
|----------------------------|--------------|--------------|--------------|
| ViDove (full) | 15.84 | 76.26 | 17.11 |
| w/o domain memory | 14.86 | 77.31 | 17.84 |
| w/o domain memory & vision | 14.56 | 77.55 | 17.50 |
| w/o Proofreader | 13.56 | 80.76 | 16.93 |

Table 2: Ablation study of ViDove under single-column setting.

not specifically optimized for short-form translation tasks, it achieves BLEU (26.05) and sCOMET (0.73) scores close to the best-performing models, such as Gemini-2.5-flash and VideoCaptioner. This demonstrates ViDove’s generalizability and robustness.

5.4 Ablation Study

To assess the contribution of different components in ViDove, we conduct an ablation study by removing modules. Our ablation study is conducted on a subset of DoveBench’s StarCraft 2 Category. As shown in Table 2, removing the domain memory significantly reduces BLEU and SubER scores, though BLEURT slightly improves—likely due to more generic paraphrasing. Excluding the proofreader agent causes the sharpest quality drop, highlighting its role in correction and consistency. While the visual module has limited impact on BLEU or BLEURT, it helps the editor correct entity-level terms (e.g., names and objects), improving factual accuracy and user experience beyond what metrics capture.

6 Conclusion

In this work, we introduced **ViDove**, a multimodal translation agent system for long-form video inputs. Our model outperforms the strongest existing baselines by up to 28.8% in BLEU and 15.5% in SubER, demonstrating significant improvements in both translation accuracy and subtitle alignment. We also release a new benchmark for the challenging task of long-form video automatic subtitling. Compared to prior work, ViDove offers a more practical and scalable solution for video automatic subtitling and translation.

Acknowledgments

We thank the FGA Subtitle Group, MetricSubs, and Star-Pigeon Group for their support in providing high-quality human-annotated datasets. This work was supported in part by funding from Pigeon AI.

References

- Mohammad Mahdi Abootorabi, Amirhosein Zobeiri, Mahdi Dehghani, Mohammadali Mohammadkhani, Bardia Mohammadi, Omid Ghahroodi, Mahdiah Soleymani Baghshah, and Ehsaneddin Asgari. 2025. [Ask in any modality: A comprehensive survey on multimodal retrieval-augmented generation](#). *Preprint*, arXiv:2502.08826.
- Anonymous. 2024. [RAG picking helps: Retrieval augmented generation for machine translation](#). In *Submitted to ACL Rolling Review - August 2024*. Under review.
- Muhammad Arslan, Hussam Ghanem, Saba Munawar, and Christophe Cruz. 2024. A survey on rag with llms. *Procedia Computer Science*, 246:3781–3790.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Maxime Bouthors, Josep Crego, and Francois Yvon. 2024. [Retrieving examples from memory for retrieval augmented neural machine translation: A systematic comparison](#). *Preprint*, arXiv:2404.02835.
- Chen-Chi Chang, Chong-Fu Li, Chu-Hsuan Lee, and Hung-Shin Lee. 2025. [Enhancing low-resource minority language translation with llms and retrieval-augmented generation for cultural nuances](#). *Preprint*, arXiv:2505.10829.
- Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, and Furu Wei. 2022. [Beats: Audio pre-training with acoustic tokenizers](#). *Preprint*, arXiv:2212.09058.
- Wenxi Chen, Yuzhe Liang, Ziyang Ma, Zhisheng Zheng, and Xie Chen. 2024. [Eat: Self-supervised pre-training with efficient audio transformer](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pages 3807–3815. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Yuheng Cheng, Ceyao Zhang, Zhengwen Zhang, Xianguang Meng, Sirui Hong, Wenhao Li, Zihao Wang, Zekai Wang, Feng Yin, Junhua Zhao, and Xiuqiang He. 2024. [Exploring large language model based intelligent agents: Definitions, methods, and prospects](#). *Preprint*, arXiv:2401.03428.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. 2024. [Qwen2-audio technical report](#). *Preprint*, arXiv:2407.10759.
- Simone Conia, Daniel Lee, Min Li, Umar Farooq Minhas, Saloni Potdar, and Yunyao Li. 2024. [Towards cross-cultural machine translation with retrieval-augmented generation from multilingual knowledge graphs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16343–16360, Miami, Florida, USA. Association for Computational Linguistics.
- Yan Ding, Xiaohan Zhang, Saeid Amiri, Nieqing Cao, Hao Yang, Andy Kaminski, Chad Esselink, and Shiqi Zhang. 2023. [Integrating action knowledge and llms for task planning and situation handling in open worlds](#). *Autonomous Robots*, 47(8):981–997.
- Yihao Ding, Kaixuan Ren, Jiabin Huang, Siwen Luo, and Soyeon Caren Han. 2024. [Pdf-mvqa: A dataset for multimodal information retrieval in pdf-based visual question answering](#). *Preprint*, arXiv:2404.12720.
- Aaron Grattafiori et al. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Alec Radford et al. 2021. [Learning transferable visual models from natural language supervision](#). *Preprint*, arXiv:2103.00020.
- Emilio Villa-Cueva et al. 2025. [Cammt: Benchmarking culturally aware multimodal machine translation](#). *Preprint*, arXiv:2505.24456.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. [BLEU might be guilty but references are not innocent](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 61–71, Online. Association for Computational Linguistics.
- Marco Gaido, Sara Papi, Matteo Negri, Mauro Cettolo, and Luisa Bentivogli. 2024. [SBAAM! Eliminating Transcript Dependency in Automatic Subtitling](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Bangkok, Thailand.
- Yuan Gao, Ruili Wang, and Feng Hou. 2023a. [How to design translation prompts for chatgpt: An empirical study](#). *Preprint*, arXiv:2304.02182.
- Zhifu Gao, Shiliang Zhang, Ian McLoughlin, and Zhijie Yan. 2023b. [Paraformer: Fast and accurate parallel transformer for non-autoregressive end-to-end speech recognition](#). *Preprint*, arXiv:2206.08317.
- Google. 2024. Google translate. <https://translate.google.com>. Accessed: 2024-07-05.
- Google DeepMind and Google Research. 2025. [Gemini 2.5 flash](#). Model card, Google AI Studio / Vertex AI. Enhanced multimodal large language model with extended capabilities in text, image, audio, and video.
- Shoutao Guo, Shaolei Zhang, Zhengrui Ma, Min Zhang, and Yang Feng. 2024a. [Agent-simt: Agent-assisted simultaneous machine translation with large language models](#). *Preprint*, arXiv:2406.06910.

- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024b. [Large language model based multi-agents: A survey of progress and challenges](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pages 8048–8057. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Zhiwei He, Tian Liang, Wenxiang Jiao, Zhuosheng Zhang, Yujiu Yang, Rui Wang, Zhaopeng Tu, Shuming Shi, and Xing Wang. 2024. [Exploring human-like translation strategy with large language models](#). *Transactions of the Association for Computational Linguistics*, 12:229–246.
- Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. [How good are gpt models at machine translation? a comprehensive evaluation](#). *Preprint*, arXiv:2302.09210.
- Junyan Hu, Parijat Bhowmick, Inmo Jang, Farshad Arvin, and Alexander Lanzon. 2021. [A decentralized cluster formation containment framework for multirobot systems](#). *IEEE Transactions on Robotics*, 37(6):1936–1955.
- Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023. [Is chat-gpt a good translator? yes with gpt-4 as the engine](#). *Preprint*, arXiv:2301.08745.
- Liyan Kang, Luyang Huang, Ningxin Peng, Peihao Zhu, Zewei Sun, Shanbo Cheng, Mingxuan Wang, Degen Huang, and Jinsong Su. 2023. [BigVideo: A large-scale video subtitle translation dataset for multimodal machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8456–8473, Toronto, Canada. Association for Computational Linguistics.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. [Segment anything](#). *Preprint*, arXiv:2304.02643.
- Zhibin Lan, Jiawei Yu, Xiang Li, Wen Zhang, Jian Luan, Bin Wang, Degen Huang, and Jinsong Su. 2023. [Exploring better text image translation with multimodal codebook](#). *Preprint*, arXiv:2305.17415.
- Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma, and JingBo Zhu. 2022. [On vision features in multimodal machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6327–6337, Dublin, Ireland. Association for Computational Linguistics.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023a. [Camel: Communicative agents for "mind" exploration of large language model society](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023b. [Api-bank: A comprehensive benchmark for tool-augmented llms](#). *Preprint*, arXiv:2304.08244.
- Wenhao Li, Dan Qiao, Baoxiang Wang, Xiangfeng Wang, Bo Jin, and Hongyuan Zha. 2023c. [Semantically aligned task decomposition in multi-agent reinforcement learning](#). *Preprint*, arXiv:2305.10865.
- Huan Lin, Fandong Meng, Jinsong Su, Yongjing Yin, Zhengyuan Yang, Yubin Ge, Jie Zhou, and Jiebo Luo. 2020. [Dynamic context-guided capsule network for multimodal machine translation](#). In *Proceedings of the 28th ACM International Conference on Multimedia, MM '20*. ACM.
- Jerry Liu. 2022. [LlamaIndex](#).
- Xinwei Long, Jiali Zeng, Fandong Meng, Zhiyuan Ma, Kaiyan Zhang, Bowen Zhou, and Jie Zhou. 2024. [Generative multi-modal knowledge retrieval with large language models](#). *Preprint*, arXiv:2401.08206.
- Chenyu Lu, Shiliang Sun, Jing Zhao, Nan Zhang, Tengfei Song, and Hao Yang. 2025. [Multimodal machine translation with visual scene graph pruning](#). *Preprint*, arXiv:2505.19507.
- Yichen Lu, Jiaqi Song, Xuankai Chang, Hengwei Bian, Soumi Maiti, and Shinji Watanabe. 2024a. [Syneslm: A unified approach for audio-visual speech recognition and translation via language model and synthetic data](#). *Preprint*, arXiv:2408.00624.
- Yichen Lu, Jiaqi Song, Chao-Han Huck Yang, and Shinji Watanabe. 2024b. [FastAdaSP: Multitask-adapted efficient inference for large speech language model](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 440–451, Miami, Florida, US. Association for Computational Linguistics.
- Jinze Lv, Jian Chen, Zi Long, Xianghua Fu, and Yin Chen. 2025. [Topicvd: A topic-based dataset of video-guided multimodal machine translation for documentaries](#). *Preprint*, arXiv:2505.05714.
- Weiyu Ma, Qirui Mi, Yongcheng Zeng, Xue Yan, Yuqiao Wu, Runji Lin, Haifeng Zhang, and Jun Wang. 2024. [Large language models play starcraft ii: benchmarks and a chain of summarization approach](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 133386–133442. Curran Associates, Inc.
- Ziyang Ma, Zhisheng Zheng, Jiaxin Ye, Jinchao Li, Zhifu Gao, Shiliang Zhang, and Xie Chen. 2023. [emotion2vec: Self-supervised pre-training for speech emotion representation](#). *Preprint*, arXiv:2312.15185.

- Raphaël Merx, Aso Mahmudi, Katrina Langford, Leo Alberto de Araujo, and Ekaterina Vylomova. 2024. [Low-resource machine translation through retrieval-augmented llm prompting: A study on the mambai language](#). *Preprint*, arXiv:2404.04809.
- et al. OpenAI. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Anishka Peter, Mai Dang, Michael Liu, Joaquin Dominguez, and Nibhrat Lohia. 2024. Multi-agent translation team (matt): Enhancing low-resource language translation through multi-agent workflow. *SMU Data Science Review*, 8(3):3.
- Alexis Plaquet and Hervé Bredin. 2023a. Powerset multi-class cross entropy loss for neural speaker diarization. In *Proc. INTERSPEECH 2023*.
- Alexis Plaquet and Hervé Bredin. 2023b. [Powerset multi-class cross entropy loss for neural speaker diarization](#). In *INTERSPEECH 2023*, interspeech_2023. ISCA.
- Qwen, :, and An Yang et al. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#). *Preprint*, arXiv:2212.04356.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Nathaniel Robinson, Perez Ogayo, David R. Mortensen, and Graham Neubig. 2023. [ChatGPT MT: Competitive for high- \(but not low-\) resource languages](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 392–418, Singapore. Association for Computational Linguistics.
- Jingqing Ruan, Yihong Chen, Bin Zhang, Zhiwei Xu, Tianpeng Bao, Guoqing Du, Shiwei Shi, Hangyu Mao, Ziyue Li, Xingyu Zeng, and Rui Zhao. 2023. [Tptu: Large language model-based ai agents for task planning and tool usage](#). *Preprint*, arXiv:2308.03427.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. [Bleurt: Learning robust metrics for text generation](#). In *Proceedings of ACL*.
- Huangjun Shen, Liangying Shao, Wenbo Li, Zhibin Lan, Zhanyu Liu, and Jinsong Su. 2024. [A survey on multi-modal machine translation: Tasks, methods and challenges](#). *Preprint*, arXiv:2405.12669.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2023a. [Large language models encode clinical knowledge](#). *Nature*, 620(7972):172–180.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaeckermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agüera y Arcas, Nenad Tomasev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. 2023b. [Towards expert-level medical question answering with large language models](#). *Preprint*, arXiv:2305.09617.
- Umut Sulubacak, Ozan Caglayan, Stig-Arne Grönroos, Aku Rouhe, Desmond Elliott, Lucia Specia, and Jörg Tiedemann. 2019. [Multimodal machine translation through visuals and speech](#). *Preprint*, arXiv:1911.12798.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2024. [Salmonn: Towards generic hearing abilities for large language models](#). *Preprint*, arXiv:2310.13289.
- Tavily AI. 2025. [Tavily](#). Accessed: 2025-07-05.
- NLLB Team and et al. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.
- Suramya Tomar. 2006. Converting video formats with ffmpeg. *Linux journal*, 2006(146):10.
- Jiaan Wang, Fandong Meng, Yingxue Zhang, and Jie Zhou. 2024. [Retrieval-augmented machine translation with unstructured knowledge](#). *Preprint*, arXiv:2412.04342.
- Shenzhi Wang, Chang Liu, Zilong Zheng, Siyuan Qi, Shuo Chen, Qisen Yang, Andrew Zhao, Chaofei Wang, Shiji Song, and Gao Huang. 2023. [Avalon’s game of thoughts: Battle against deception through recursive contemplation](#). *Preprint*, arXiv:2310.01320.
- Yutong Wang, Jiali Zeng, Xuebo Liu, Derek F. Wong, Fandong Meng, Jie Zhou, and Min Zhang. 2025a. [Delta: An online document-level translation agent based on multi-level memory](#). *Preprint*, arXiv:2410.08143.
- Yutong Wang, Jiali Zeng, Xuebo Liu, Derek F. Wong, Fandong Meng, Jie Zhou, and Min Zhang.

- 2025b. [Delta: An online document-level translation agent based on multi-level memory](#). *Preprint*, arXiv:2410.08143.
- Weifeng2333. 2024. Videocaptioner: An open-source cascaded system for video subtitling. <https://github.com/WEIFENG2333/VideoCaptioner>. Accessed: 2025-07-04.
- Patrick Wilken, Panayota Georgakopoulou, and Evgeny Matusov. 2022. [SubER - a metric for automatic evaluation of subtitle quality](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 1–10, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Minghao Wu, Jiahao Xu, and Longyue Wang. 2024a. [TransAgents: Build your translation company with language agents](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 131–141, Miami, Florida, USA. Association for Computational Linguistics.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. 2023. [Autogen: Enabling next-gen llm applications via multi-agent conversation](#). *Preprint*, arXiv:2308.08155.
- Yihan Wu, Yichen Lu, Yifan Peng, Xihua Wang, Ruihua Song, and Shinji Watanabe. 2024b. [Enhancing audio-visual speech recognition through bifocal preference optimization](#). *Preprint*, arXiv:2412.19005.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *Preprint*, arXiv:1609.08144.
- Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Marianna Nezhurina, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. 2024c. [Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation](#). *Preprint*, arXiv:2211.06687.
- Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024a. [Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation](#). *Preprint*, arXiv:2401.08417.
- 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](#). *Preprint*, arXiv:2503.20215.
- Peng Xu, Andrea Madotto, Chien-Sheng Wu, Ji Ho Park, and Pascale Fung. 2018. [Emo2vec: Learning generalized emotion representation by multi-task training](#). *Preprint*, arXiv:1809.04505.
- Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. 2024b. [Language agents with reinforcement learning for strategic play in the werewolf game](#). *Preprint*, arXiv:2310.18940.
- Wenjia Zhai. 2024. [Self-adaptive multimodal retrieval-augmented generation](#). *Preprint*, arXiv:2410.11321.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. [Sigmoid loss for language image pre-training](#). *Preprint*, arXiv:2303.15343.
- Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2024. [Multilingual machine translation with large language models: Empirical results and analysis](#). *Preprint*, arXiv:2304.04675.
- Yuxin Zuo, Bei Li, Chuanhao Lv, Tong Zheng, Tong Xiao, and JingBo Zhu. 2023. [Incorporating probing signals into multimodal machine translation via visual question-answering pairs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14689–14701, Singapore. Association for Computational Linguistics.

A ViDove Details

A.1 Multi-agent Translation System

| | | |
|---------------------------------------|----------|--|
| Auditory Agent | SpeechLM | SALMONN (Tang et al., 2024), Gemini-2.5-flash (Google DeepMind and Google Research, 2025) Qwen2-Audio (Chu et al., 2024), Qwen-2.5-Omni (Xu et al., 2025), |
| | ASR | Whisper-series (Radford et al., 2022) Paraformer(Gao et al., 2023b) |
| | Others | Emo2Vec(Xu et al., 2018), BEATs(Chen et al., 2022), EATs(Chen et al., 2024) CLAP(Wu et al., 2024c) Pyannote(Plaquet and Bredin, 2023b) |
| Visual Agent | VLMs | GPT4-o(OpenAI, 2024), Qwen2.5-VL(Bai et al., 2025) |
| | Others | CLIP(et al., 2021), SigCLIP(Zhai et al., 2023), SegAnything(Kirillov et al., 2023) |
| Translation Agent & Post-editing Team | LLMs | GPT-series(OpenAI, 2024), LLaMA-series(et al., 2024), Qwen-series(Qwen et al., 2025) |
| | Others | Google-translate(Google, 2024), NLLB(Team and et al., 2022) |
| Memory System | Web | Tavily(Tavily AI, 2025) |
| | Local | Llama-index(Liu, 2022) |
| Other Tools | - | FFmpeg(Tomar, 2006) |
| Evaluation Metrics | - | BLEU (Freitag et al., 2020), COMET (Rei et al., 2020), BLEURT (Sellam et al., 2020), SubER (Wilken et al., 2022), SubSONAR (Sulubacak et al., 2019) |

Table 3: Base models and tools used in ViDove

Algorithm 3 Multi-agent Translation Pipeline

Require: Transcript chunk T_i , Memory $\mathcal{M} = \{\mathcal{M}^s, \mathcal{M}^l\}$, LLM Translator \mathcal{L}_t , Proofreader \mathcal{L}_{pr} , Editor \mathcal{L}_{ed}

Ensure: Translated transcript chunk T_i^* , updated short-term memory $\mathcal{M}_{history}^s$, and translation prompt

```

 $p_{translation} \leftarrow \text{retrieve}(\mathcal{M}_{history}^s, T_i)$ 
 $context_i \leftarrow \text{retrieve}(\mathcal{M}_{context}^s)$ 
 $domain\_guide \leftarrow \text{query}(\mathcal{M}_{domain}^l, T_i)$ 
 $p_{translation} \leftarrow (p_{translation}, context_i, domain\_guide)$ 
 $T_i^* \leftarrow \mathcal{L}_t(T_i, p_{translation})$ 
 $T_{i,pr}^* \leftarrow \mathcal{L}_{pr}(T_i^*, \mathcal{M}^s, \mathcal{M}^l)$   $\triangleright$  Proofreader checks grammar, style, terminology
 $T_{i,ed}^* \leftarrow \mathcal{L}_{ed}(T_{i,pr}^*, \mathcal{M}^s, \mathcal{M}^l)$   $\triangleright$  Editor ensures logical and contextual consistency
 $\mathcal{M}_{history}^s \leftarrow (\mathcal{M}_{history}^s, T_i, T_{i,ed}^*)$ 
return  $(T_{i,ed}^*, \mathcal{M}_{history}^s)$ 

```

A.1.1 Proofreader Agent

To further demonstrate the role of the proofreader agent in our pipeline, we present a series of log messages captured during a real session. The proofreader monitors intermediate translations and provides suggestions to correct terminology, sentence structure, and domain-specific references. We have provided examples(4) illustrating its intervention.

| |
|--|
| [Segment 130] PASS |
| [Segment 131] The term "pilum" in the source text seems to be a mistake or unclear. It might be intended to refer to "pylon" based on the term context provided. Consider verifying this with the editor. |
| [Segment 132] The translation of "sporter" as "孢子" is incorrect. Based on the term context, "sporter" might be a misinterpretation of "spore crawler," which should be translated as "孢子爬虫." Verify with the editor if "sporter" is indeed meant to be "spore crawler." |
| [Segment 133] The translation is missing a verb or context to make it a complete sentence. Consider adding context or a verb to improve fluency, such as "显然现在是Nagra的时刻, 伙计。" |
| [Segment 134] The translation of "Spire" as "空军基地" is incorrect. According to the term context, "Spire" should be translated as "飞龙塔." Adjust the translation to reflect this terminology. |

Table 4: proofreader log

The proofreader agent checks both the source and the translation, considering context information, which allows it to identify potential misinterpretations and suggest targeted corrections. In the above use case, it successfully detects anomalies in Segments 131 and 132. Segment 134 also showcases its ability to correct common terminology errors based on domain knowledge.

A.2 ViDove Demo Page

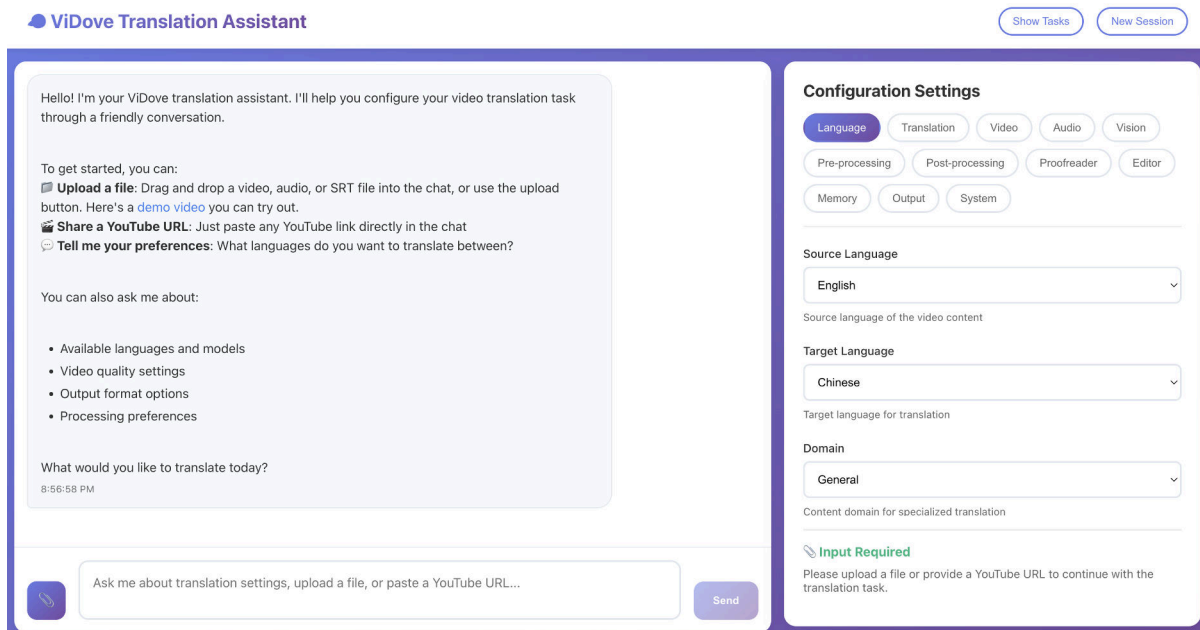


Figure 2: User interface of ViDove

A.3 Prompt for ViDove Agents

ViDove Translator Agent

“ You are a professional translator. your job is to translate texts in domain of {domain} from {source language} to {target language}
you will be provided with a segment in source language parsed by line, where your translation text should keep the original meaning and the number of lines.
Keep every \n in the translated text in the corresponding place, and make sure to keep the same number of lines in the translated text.
You must break the translated sentence into multiple lines accordingly if original text breaks a complete sentence into different lines.
You should only output the translated text line by line without any other notation.
You current task is to translate the script in the domain of {domain} from {source language} to {target language}
Here are some supporting information including previous translation history, context documenting, supporting documents from internet and video clips description that might help you translate the text. Please refer to them if necessary. Previous translation history: \n
{Translation Histories}
if you detect any word is in the following context, please use it as a reference for current translation
{Context documents retrieved from knowledge base}
Here are some supporting documents that might help you translate the text, refer to them if necessary.:
{ supporting documents from web search }
Here are some descriptions of video clips that might help you translate the text, refer to them if necessary. :
{video clips descriptions}
Now please translate the following text from {source language} to {target language}
{text to be translated}
Your translation: ”

ViDove Editor Agent

“ You are an Editor ensuring overall translation quality and coherence, aligning the translation with the original video content in domain {domain}, you must ensure the term and style are aligned with the domain’s language.
Segment index: {idx} Source text: {source}
Translated text: {translation}
Here is a provided suggestion for each segment, which may or may not useful for your revision, you may use the suggestion only if necessary (for example, term correctness). Note that the suggestion may not be accurate, the proofreader has less information comparing to you, so you need to double check before making revision. The proofreader may return "UNCLEAR" if they are not sure about the translation, they will specify the location and you need to check with other information provided to you to solve for unclear. If there is no suggestions, you may ignore this part, but still check with other modality context and long-term memory for correctness and coherence. Suggestion:
{suggestion if suggestion else "No suggestion provided."}
Your edit will also follow the following instruction if provided: User instruction: {user instruction if user introduction else "No user instruction provided."}
— Multimodal Context (Short-Term Memory) — Visual cues: You may use visual cues from the video to improve translation or make corrections, the source text might not be accurate, you need to check with the video context if provided: {visual context}
Audio cues: {audio context}
Translation context: You will be provided with the previous and next 5 segments’ translations, which may help you understand the context and make corrections: Previous translation history (up to 5 segments): {Previous translation history} Past translation history (up to 5 segments): {Past translation history}
— Long-Term Memory — Long-term memory provides broader context and domain-specific knowledge, you may use it to improve translation or make corrections: {long term memory}
Notice: 1. Corrections or adjustments to better align text with the video context. 2. Suggestions for improving coherence across segments. 3. Logical consistency and any broader context adjustments. 4. Ensure the translation is accurate and aligned with the domain {domain}. 5. Ensure translation is smooth and fluent across segments. 6. To ensure the fluency in {target language}, you do not have to ensure translation be word by word accurate, but be sure to convey the same information.
— Important — Directly return the revised content only. ”

ViDove Proofreader Agent

“ You are a translation proofreader. Below are {number of segments} subtitle segments. Some are full sentences, some are fragments. Give **specific advice** for each one, but do not treat each segment separately you need information across segment.

Return suggestions in this format: Segment 0: [your comment here] Segment 1: [your comment here] ...

DO NOT return JSON. DO NOT rewrite the translation. Just return suggestion texts.

— {segments}

Short-term memory: {short term memory}

Term context: {local context}

Web memory context: {web search context}

Focus on: 1. Translation accuracy while sticking to domain {domain} (missing or incorrect meanings) 2. Fluency (grammar, spelling, repetition. Only if it affects understanding) and ensure the translation is smooth and fluent across segments. 3. Terminology (Use term context to edit idioms, ensure every sentence is translated into domain-specific language) 4. If you have no suggestions, return "PASS" for that segment. 5. Source text isn't 100% accurate. If you have doubt about the source text, return "UNCLEAR" and specify the location, editor will check the issue. 6. Only make suggestions if you believe revision is necessary. ”

A.4 Translation Sample



Figure 3: **Vidove Video Cues Summarization for Translation Sample:** The image portrays a surreal and dramatic illustration featuring a large, exaggerated face of a bearded man on the left, expressing concern or thoughtfulness. On the right, a dark, patterned background showcases large white flames, with small skeletal figures interacting with them in the foreground. The scene is rendered in a monochromatic color scheme.

| | |
|----------------|---|
| ORIGINAL TEXT | What happens when the devil walks among us? |
| GROUND TRUTH | 若魔鬼化身凡人混迹其中，世界会变成什么样？ |
| VIDOVE | 当魔鬼行走在人间时会发生什么？ |
| VIDEOCAPTIONER | 当魔鬼在我们中间行走会发生什么？ |

Table 5: Case study for translation quality. **Blue** highlights translation deviations to VIDOVE, and **red** highlights deviations to VIDEOCAPTIONER.

B DoveBench

B.1 DoveBench Stats

DoveBench is a benchmark dataset designed to evaluate video translation and subtitling models. It contains a total of 50 videos, amounting to 17.23 hours of content and featuring 16,968 subtitle entries. The dataset’s total text includes 189,157 words.

The dataset is composed of two distinct categories sourced from fan sub groups:

- **CS:** This category includes 23 Counter-Strike related videos from the "fazeclan galaxy archive" fan sub group. The videos in this section have an average duration of approximately 13 minutes (777.7 seconds).
- **SC2:** This category consists of 27 videos in the StarCraft 2 domain from the "StarPigeon Fan sub group". These videos are generally longer, with an average duration of over 27 minutes (1635.3 seconds).

A key feature of DoveBench is the inclusion of detailed ground truth, which is essential for evaluation. We provide human-annotated ground-truth subtitles for all videos. The dataset’s comprehensive statistics on character counts, word counts, duration, and subtitle distribution make it a valuable resource for assessing the performance of video translation systems.

| Category Statistics | | | Overall Dataset Statistics | |
|------------------------|---------------------|----------------------|-------------------------------|---------------|
| Statistic | CS | SC2 | Statistic | Overall |
| NUMBER OF VIDEOS | 23 | 27 | NUMBER OF VIDEOS | 50 |
| TOTAL DURATION | 4.97 h (298.1 min) | 12.27 h (735.9 min) | TOTAL DURATION | 17.23 hours |
| textscAverage Duration | 12.96 min (777.7 s) | 27.26 min (1635.3 s) | AVERAGE DURATION | 20.68 minutes |
| | | | TOTAL SUBTITLE LINES | 16,968 |
| | | | AVG. SUBTITLE LINES PER VIDEO | 346.3 |
| | | | TOTAL WORDS | 189,157 |
| | | | AVG. WORDS PER VIDEO | 3,860 |

Table 6: Detailed statistics of the DoveBench dataset

C Baseline Systems Details

C.1 Gemini Prompt for DoveBench

Gemini Prompt (English Version)

“ You are a professional transcription and translation assistant.
Please transcribe this audio/video file and translate it into Simplified Chinese. Carefully follow the instructions below:
1. Each segment should: - Contain a natural sentence or phrase in Simplified Chinese, not too long. - Have a valid start and end time in the format ‘h:mm:ss,ms’ (e.g., "0:00:01,229"). - Ensure the start time is less than the end time, and that each segment’s start time equals the previous segment’s end time (no overlap or gap). - If uncertain, round timestamps to the nearest 10 milliseconds.
2. Translation Guidelines: - First, accurately understand the original audio content. - Translate into natural, fluent Simplified Chinese. - Retain the meaning and tone of the original speech. - Keep proper nouns and technical terms accurate. - Maintain sentence boundaries suitable for subtitle readability.
3. Notes: - Proper nouns and technical terms — remain accurate. - Sentence boundaries — avoid breaking at unnatural pauses. - Chinese grammar and natural fluency.
Please provide the transcription and translation in the specified structured format.
The output language must be Chinese. ”

C.2 Qwen-2.5-Omni

C.2.1 Prompts for DoveBench and BigVideo

Qwen 2.5 Omni Prompt (Chinese Version)

“翻译提供的视频中的说话内容到中文。只需要输出翻译内容原文，不要输出任何解释。”

C.2.2 Prompt Design and Video Processing Strategy

This section outlines the strategies employed for prompt design and video processing to optimize Qwen 2.5 Omni’s performance on the DoveBench and BigVideo datasets.

Prompt Language: Given the superior performance of Chinese prompts in enhancing Qwen 2.5 Omni’s instruction-following capabilities during preliminary evaluations, all experiments on both the DoveBench and BigVideo datasets exclusively utilized Chinese prompts.

Prompt Complexity: Specifically, this approach was adopted because initial trials with prompts designed similarly to those employed for Gemini demonstrated that Qwen exhibited a deficiency in instruction following when presented with complex instructions. This observation led to the selection of the aforementioned prompts, as they consistently yielded superior results.

Video Processing Strategy: Furthermore, due to Qwen 2.5 Omni’s constrained contextual understanding when processing video data, most videos within the datasets—even those only one to two minutes in duration—could not be processed directly. To address this limitation, videos were manually segmented into approximately ten-second clips. Each segment was then individually fed into the model, and the processed outputs for each segment were subsequently concatenated to form a complete SRT file.