

HausaNLP at SemEval-2025 Task 2: Entity-Aware Fine-tuning vs. Prompt Engineering in Entity-Aware Machine Translation

Abdulhamid Abubakar¹, Hamidatu Abdulkadir², Ibrahim Rabiu Abdullahi³,
Abubakar Auwal Khalid¹, Ahmad Mustapha Wali¹, Amina Aminu Umar¹, Maryam Bala¹,
Sani Abdullahi Sani⁴, Ibrahim Said Ahmad⁵, Shamsuddeen Hassan Muhammad⁶,
Idris Abdulmumin⁷, Vukosi Marivate⁷

¹HausaNLP, ²Kaduna State University, ³Ahmadu Bello University, ⁴University of the Witwatersrand,

⁵Northeastern University, ⁶Imperial College, ⁷Data Science for Social Impact, University of Pretoria

correspondence: abdulhamid@ab-bkr.com, idris.abdulmumin@up.ac.za

Abstract

This paper presents our findings for SemEval 2025 Task 2, a shared task on entity-aware machine translation (EA-MT). The goal of this task is to develop translation models that can accurately translate English sentences into target languages, with a particular focus on handling named entities, which often pose challenges for MT systems. The task covers 10 target languages with English as the source. In this paper, we describe the different systems we employed, detail our results, and discuss insights gained from our experiments. In our initial approach, we selected the distilled 600M-parameter variant of NLLB-200, which had been fine-tuned using the training dataset provided by the task organizers. For the second and third approaches, we opted for prompt engineering using two templates on Google’s Gemini-1.5 model. The two templates were based on zero-shot and few-shot learning techniques respectively. Ultimately, the Gemini results demonstrated superior performance compared to the fine-tuned NLLB model. Furthermore, our approach revealed that European languages had higher overall scores compared to other languages. However, it is worth noting that we observed an interesting performance from Turkish, which even outperformed some European languages.

1 Introduction

The quality of translations produced by machine translation (MT) systems has improved considerably (Abdulmumin et al., 2024). However, despite these advancements, translations into the target language still contain errors, often due to the challenges associated with translating named entities (Riktors and Miwa, 2024). Entity-aware machine translation is a type of MT that considers specific entities, such as names, locations, and organizations, to enhance translation accuracy and fluency (Conia et al., 2024). Several approaches have been

proposed to improve these models’ ability to translate named entities more effectively, accounting for the need for transliteration in some cases while generating equivalent translations in others.

In this paper, we describe our submissions to SemEval 2025 Task 2: Entity-Aware Machine Translation shared task (Conia et al., 2025). Our systems include sequence-to-sequence entity-aware supervised models trained to improve the translation of English sentences into French, German, Spanish, Italian, Japanese, and Arabic. A pre-trained NLLB (NLLB Team et al., 2022) model was fine-tuned for bilingual translation with the provided training data in each of these languages. Furthermore, we investigated the performances of a closed-source Large Language Model (LLM), Gemini (Team et al., 2024), on these languages, in addition to Chinese, Korean, Thai, and Turkish. Our results indicate that Gemini, in a zero-shot setup, achieved the best overall performance, with only a few languages showing improvements when examples from the training data were incorporated in few-shot setups.

2 Related Works

Several studies have explored approaches to improving named entity (NE) translation in machine translation (MT) systems. Xie et al. (2022) proposed an entity-aware model that employs classifiers in both the encoder and decoder to handle named entities more effectively. Other methods include hierarchical encoders with chunk-based processing (Ugawa et al., 2018), IOB-tagging for improved NE annotation (Modrzejewski et al., 2020), and a decoupled NE handling approach that enhances translation quality without modifying the NMT architecture (Mota et al., 2022). Zeng et al. (2023) developed an Extract-and-Attend strategy that first identifies and translates named entities separately before incorporating them into the trans-

lation process.

Recent advancements also include entity-aware multi-task training on pre-trained models such as T5, which improves NE translation quality and also increases the number of named entities generated in German translations of English texts (Rikters and Miwa, 2024). Yang et al. (2024) introduced the AMFF and CAMFF frameworks, which utilize attention mechanisms to improve named entity recognition (NER) by incorporating multilevel contextual features. Jauhari et al. (2024) introduced Entity-Aware Techniques (EaT) that integrate semantic parsing to help MT models recognize and accurately translate named entities. Awiszus et al. (2024) evaluated NE translation in speech translation systems, highlighting persistent challenges despite improvements in recall and precision. While these methods improve entity translation, our approach fine-tunes the pre-trained NLLB model using both the provided training data and extracted named-entity translations from Wikidata. Additionally, we evaluate the closed-source Gemini model without fine-tuning, assessing its performance in both zero-shot and few-shot setups. Our work focuses on achieving a balance between overall translation quality and ensuring the accurate translation of critical named entities.

3 Proposed Approaches

The team adopted four approaches for this task. The first two approaches included fine-tuning the NLLB pre-trained model in a bilingual setup for **eng**→**xxx** translation. The two other approaches were the use of two different prompt templates to evaluate the performance of Google’s Gemini closed-source model. The team combined the traditional fine-tuning technique and prompt-engineering strategies to assess their relative effectiveness in preserving entity integrity during bilingual translation.

3.1 Supervised Fine-tuning

This method was employed to train bilingual translation models for six languages, as training data was not available for all target languages. In the first fine-tuning approach, we fine-tuned the base NLLB model using the provided training data. The second approach involved leveraging the SpaCy (Honnibal et al., 2020) NER framework to extract named entities from the training data and then searching for their equivalent translations on Wiki-

data (Vrandečić and Krötzsch, 2014). The resulting entity pairs for each source-target language pair were used as additional training data to fine-tune the models. Rather than focusing solely on the named entities provided in the task, we opted to use all entities present in the training data. This approach aimed to improve the models’ ability to translate a broader range of entities rather than limiting ourselves to the subset specified in the task. While restricting to the provided subset might have resulted in better performance, we prioritized enhancing the model’s accuracy, even at the potential cost of competitiveness. This second approach led to the development of an entity-optimized variant of the NLLB-200 model, refined through targeted fine-tuning using the extracted named entities.

3.2 Prompt Engineering

Complementing these fine-tuning methods, two prompt-based strategies were implemented. The zero-shot approach utilized minimalist templates instructing the model to "preserve entity integrity" during translations. The few-shot variant extended this with 10 curated demonstration pairs from the training data in the task repository. The choice of these two prompts was to see the effect on performance of the Gemini model when given an example outputs with preserved entities. For target languages without training data, the template provided no examples. The examples followed a structured template showing entity preservation by providing the Wikidata-id of the entity as present in the training dataset.

4 Experiments

4.1 Dataset and pre-processing

The dataset used to fine-tune the NLLB-200 model was made available as part of the shared task (Conia et al., 2024), with the number of data points in the different splits detailed in Table 2. This repository included training data for English-to-Arabic, German, French, Spanish, Italian, and Japanese translations, which were also used for the few-shot prompting approach. Additionally, the validation and test data for all 10 languages are provided, with the test data being unlabelled.

We also extracted the NEs from Wikidata¹, resulting in 4,587 unique entities and their translations (where available); see Table 3. We used these

¹https://www.wikidata.org/wiki/Wikidata:Main_Page

metric	method	ar	de	es	fr	it	ja	ko	th	tr	zh
M-ETA	NLLB+NE	20.61	20.86	32.75	22.85	27.28	12.74	-	-	-	-
	Gemini 0-shot	32.66	38.15	47.92	38.77	40.31	35.10	34.67	18.80	40.82	8.53
	Gemini 10-shots+NE	34.17	38.14	48.29	35.32	39.39	34.93	33.75	18.62	41.54	8.09
COMET	NLLB+NE	87.98	88.42	91.26	85.07	89.52	88.86	-	-	-	-
	Gemini 0-shot	88.56	89.30	91.71	88.35	89.98	91.06	90.71	83.41	92.42	87.85
	Gemini 10-shots+NE	89.59	89.86	92.50	88.95	90.64	92.31	91.34	83.97	92.85	88.66
Overall	NLLB+NE	33.40	33.76	48.20	36.02	41.82	22.28	-	-	-	-
	Gemini 0-shot	47.72	53.46	62.95	53.89	55.68	50.67	50.17	30.68	56.63	15.55
	Gemini 10-shots+NE	49.47	53.55	63.45	50.56	54.92	50.68	49.29	30.48	57.40	14.83

Table 1: **M-ETA**, **COMET**, and the **Overall** scores of the evaluated approaches. The scores in **bold** font indicate our final system submission, representing our ranked system according to the task instructions.

entities when fine-tuning the NLLB model and as part of the few-shot prompting. Training the NLLB model required tokenizing the training data; this was achieved using the NLLB tokenizer. The prompt engineering approach did not require any data preprocessing.

4.2 Models and environment setup

The fine-tuning approach involved training the NLLB-200 model, specifically the distilled 600M² parameter variant. The fine-tuning was conducted using the default Hugging Face hyperparameter setup: a batch size of 32, sequence lengths of 128 for both source and target, a generation beam search width of 5, a dropout rate of 0.1, and a training duration of 10 epochs. Early stopping was applied if there was no improvement in the model’s performance after two consecutive epochs.

For the prompt engineering method, we utilized Langchain’s (Chase, 2022) ChatPromptTemplate³ and ChatGoogleGenerativeAI⁴ modules to evaluate the performance of Gemini Flash 1.5⁵. We provide the prompt templates that were used in Templates 1 and 2. The results were saved in JSON format as required from the task submission description.

4.3 Evaluation

We used the metrics provided for the shared task to evaluate our systems. These metrics are COMET

(Rei et al., 2020) and Manual Entity Translation Accuracy (M-ETA; Conia et al. 2024). COMET is a machine translation evaluation metric that leverages a pre-trained model to generate quality scores by comparing system outputs to human translations. M-ETA assesses the accuracy of entity translations in machine translation by computing the proportion of correctly translated entities against a gold standard. Untranslated source texts are scored 0. The overall score, Equation (1), is computed as the harmonic mean (F1 score) of the COMET and M-ETA metrics, ensuring that systems are rewarded for balanced performance across both rather than excelling in only one.

$$\text{Overall Score} = 2 \times \frac{\text{COMET} \times \text{M-ETA}}{\text{COMET} + \text{M-ETA}} \quad (1)$$

5 Results

Table 1 below gives a summary of the models’ performances across the proposed approaches. The NLLB column contains only results for languages that have training data.

The results indicate that while all 3 approaches achieved similar COMET scores at sentence level, the Gemini model, both in its 0- and few-shot settings, performed better at translating the named entities in the source text. This can be observed by the higher M-ETA scores obtained by Gemini compared to NLLB. In almost all the target languages where we have results, Gemini has almost twice the M-ETA scores of NLLB. At the language level, English-to-Spanish translation achieved the best performance across all the evaluated approaches on all the evaluation metrics. This is in contrast to the Chinese language translation from English, with a paltry overall score of 14.83.

²<https://huggingface.co/facebook/nllb-200-distilled-600M>

³https://python.langchain.com/api_reference/core/prompts/langchain_core.prompts.chat.ChatPromptTemplate.html

⁴https://python.langchain.com/api_reference/google_genai/chat_models/langchain_google_genai.chat_models.ChatGoogleGenerativeAI.html

⁵<https://ai.google.dev/gemini-api/docs/models/gemini#gemini-1.5-flash>

It is important to highlight the performance disparity between European and Asian languages. European languages, including Spanish, Italian, French, and German, consistently achieved higher scores across all methods, except for Japanese. Italian, for instance, achieved scores of 55.68 and 54.92 with Gemini’s approaches, while NLLB managed 22.28. In contrast, Asian languages presented lower scores, with Chinese recording the lowest scores (15.55 and 14.83) among all languages tested. Japanese, while performing moderately well with Gemini (50.57 and 50.68), showed lower scores compared to European languages in NLLB but higher than the other Asian language: Arabic. Overall, Spanish and Italian achieved better results than the others in NLLB.

As highlighted earlier, not all the target languages had training data, so this affected the testing of these languages with the NLLB model, thus resulting in no scores for Chinese, Korean, Thai, and Turkish. However, an interesting finding was Turkish’s performance with Gemini (56.63 and 57.40) despite the absence of few-shot examples, yet its performance was second to only Spanish.

A comparison of the two Gemini implementations revealed minimal differences between the zero-shot and few-shot approaches. This suggests that elaborate few-shot prompting may not be necessary to achieve optimal results in entity translation tasks.

6 Conclusion

Prompt engineering for large language models like Gemini proves effective for entity-aware machine translation. Comparative studies with fine-tuning the NLLB-200 model show a consistent performance advantage for Gemini’s zero-shot and few-shot prompting across target languages, highlighting its ability to preserve entity integrity in translation. While our results reveal nuances in performance across language families, with European languages exhibiting stronger overall scores and an intriguing performance from Turkish, the overarching trend favors prompt-based methodologies.

7 Limitations

Our evaluation was constrained by the lack of training data for Chinese, Korean, Thai, and Turkish in the task repository, preventing us from fine-tuning NLLB models for these languages. Additionally, due to limited computational resources, we were

only able to fine-tune the distilled 600M-parameter variant of NLLB-200, rather than a larger model that could potentially yield better results.

Ethics Statement

This work followed the guidelines provided in SemEval 2025.

Acknowledgements

The authors thank HausaNLP, particularly Dr. Idris Abdulmumin, Dr. Shamsudden Hassan Muhammad, and Dr. Ibrahim Said Ahmad, for their invaluable mentorship throughout the process.

References

- Idris Abdulmumin, Sthembiso Mkhwanazi, Mahlatse Mbooi, Shamsuddeen Hassan Muhammad, Ibrahim Said Ahmad, Neo Putini, Miehleketo Mathabula, Matimba Shingange, Tajuddeen Gwadabe, and Vukosi Marivate. 2024. [Correcting FLORES evaluation dataset for four African languages](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 570–578, Miami, Florida, USA. Association for Computational Linguistics.
- Maximilian Awiszus, Jan Niehues, Marco Turchi, Sebastian Stüker, and Alex Waibel. 2024. [Charles lo- cock, lowcock or lockhart? offline speech translation: Test suite for named entities](#). In *Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024)*, pages 291–297, Bangkok, Thailand (in-person and online). Association for Computational Linguistics.
- Harrison Chase. 2022. [Langchain](#).
- 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.
- Simone Conia, Min Li, Roberto Navigli, and Saloni Potdar. 2025. SemEval-2025 task 2: Entity-aware machine translation. In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*. Association for Computational Linguistics.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Sarthak Jauhari, Saisuresh Krishnakumaran, Dinkar Vasudevan, and Rahul Goel. 2024. Entity-aware joint translation of query and semantic parse.

Maciej Modrzejewski, Miriam Exel, Bianka Buschbeck, Thanh-Le Ha, and Alex Waibel. 2020. Incorporating external annotation to improve named entity translation in nmt. In *Proceedings of the 22nd annual conference of the European association for machine translation*, pages 45–51.

Pedro Mota, Vera Cabarrão, and Eduardo Farah. 2022. [Fast-paced improvements to named entity handling for neural machine translation](#). In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, pages 141–149, Ghent, Belgium. European Association for Machine Translation.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#).

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.

Matiss Riktors and Makoto Miwa. 2024. [Entity-aware multi-task training helps rare word machine translation](#). In *Proceedings of the 17th International Natural Language Generation Conference*, pages 47–54, Tokyo, Japan. Association for Computational Linguistics.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, et al. 2024. [Gemini: A family of highly capable multimodal models](#).

Arata Ugawa, Akihiro Tamura, Takashi Ninomiya, Hiroya Takamura, and Manabu Okumura. 2018. [Neural machine translation incorporating named entity](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3240–3250, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

language	train	validation	test
Arabic (ar)	7,220	722	4,547
German (de)	4,087	731	5,876
Spanish (es)	5,160	739	5,338
French (fr)	5,531	724	5,465
Italian (it)	3,739	730	5,098
Japanese (ja)	7,225	723	5,108
Korean (ko)	-	745	5,082
Thai (th)	-	710	3,447
Turkish (tr)	-	732	4,473
Chinese (zh)	-	722	5,182

Table 2: Shared task data statistics.

Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.

Shufang Xie, Yingce Xia, Lijun Wu, Yiqing Huang, Yang Fan, and Tao Qin. 2022. End-to-end entity-aware neural machine translation. *Mach. Learn.*, 111(3):1181–1203.

Zhiwei Yang, Jing Ma, Hechang Chen, Jiawei Zhang, and Yi Chang. 2024. [Context-aware attentive multilevel feature fusion for named entity recognition](#). *IEEE Transactions on Neural Networks and Learning Systems*, 35(1):973–984.

Zixin Zeng, Rui Wang, Yichong Leng, Junliang Guo, Xu Tan, Tao Qin, and Tie-yan Liu. 2023. Extract and attend: Improving entity translation in neural machine translation. *arXiv preprint arXiv:2306.02242*.

A Appendix

We provide the dataset and extracted named entity statistics and also the prompt templates used in the experiments.

Entity type	count										
	all	ar	de	es	fr	it	ja	ko	th	tr	zh
PERSON	1,507	908	1,083	1,081	1,114	1,069	1,021	941	673	933	1,014
ORG	1,082	648	785	788	788	767	745	680	470	660	734
GPE	522	291	360	363	369	351	333	307	201	292	330
DATE	379	221	264	269	268	263	246	229	160	231	248
WORK_OF_ART	282	171	209	205	214	207	194	175	122	178	193
EVENT	187	105	135	137	138	134	121	105	73	108	120
LOC	183	102	127	128	128	125	118	101	69	105	121
NORP	169	103	124	124	125	122	116	105	67	98	116
FAC	135	78	94	95	96	94	88	80	43	80	87
PRODUCT	51	37	42	43	43	42	40	37	28	37	39
LAW	31	21	23	24	23	23	22	21	16	22	22
QUANTITY	28	14	19	19	20	20	18	14	10	13	18
MONEY	15	5	8	8	8	8	6	5	5	6	7
TIME	10	5	7	7	7	7	7	4	3	5	6
PERCENT	4	1	2	3	3	2	3	2	1	1	2
LANGUAGE	2	1	2	2	2	2	1	1	1	1	1

Table 3: Entities extracted from training data and their translations obtained from Wikidata.

Prompts 1: Template 1

```
PromptTemplate(
  input_variables=["sentence", "tgt", "ne", "examples"],
  template="""Instruction:
  Translate the following text from english to {tgt}, ensuring that all
  named-entities are accurately translated with no additional explanations. Use
  the provided translation examples and translated named-entities (if provided)
  for consistency. Do not send the English text back in the response, generate
  only the translation and nothing more.
  Named entities:
  {ne}
  Examples:
  {examples}
  Now generate the {tgt} translation of the following english text: {sentence}"""
)
```

Prompts 2: Template 2

```
PromptTemplate(
  input_variables=["sentence", "tgt"],
  template="""Instruction:
  Translate the following text from english to {tgt}, ensuring that all named-
  entities are accurately translated with no additional explanations. Do not send
  the English text back in the response, generate only the translation and nothing
  more.
  Now generate the {tgt} translation of the following english text: {sentence}"""
)
```

Text	Prediction
What kind of artwork is The Signal-Man?	¿Qué tipo de obra de arte es El Señalero?
When was The Signal-Man first published?	¿Cuándo se publicó por primera vez El Guarda-Señales?
Who is the author of The Signal-Man?	¿Quién es el autor de El Señalero?
Who are some of the main actors in the movie About Joan?	¿Quiénes son algunos de los principales actores en la película Sobre Joan?
Who directed the 2022 film About Joan?	¿Quién dirigió la película de 2022 Sobre Joan?
Can you provide a brief plot summary of About Joan?	Joan, una joven con una vida aparentemente normal, descubre que es la reencarnación de Juana de Arco. Su destino la lleva a una batalla moderna contra una poderosa corporación.
Who are the main actors in The Big Hit?	¿Quiénes son los actores principales de Golpe maestro?
Who directed the movie The Big Hit?	¿Quién dirigió la película Golpe maestro?
What is the genre of the film The Big Hit?	¿Cuál es el género de la película Golpe maestro?
What is the significance of Where are you, Dear General? in North Korean culture?	¿Cuál es la significancia de ¿Dónde está usted, querido general? en la cultura norcoreana?

Table 4: Sample translation from English to Spanish (Gemini 0-Shot)

Text	Prediction
What kind of artwork is The Signal-Man?	¿Qué tipo de obra de arte es El Señalero?
When was The Signal-Man first published?	¿Cuándo se publicó por primera vez El hombre de la señal?
Who is the author of The Signal-Man?	¿Quién es el autor de El hombre de la señal?
Who are some of the main actors in the movie About Joan?	¿Quiénes son algunos de los actores principales de la película Sobre Joan?
Who directed the 2022 film About Joan?	¿Quién dirigió la película Sobre Joan de 2022?
Can you provide a brief plot summary of About Joan?	¿Puedes proporcionar un breve resumen de la trama de Sobre Joan?
Who are the main actors in The Big Hit?	¿Quiénes son los actores principales de Golpe maestro?
Who directed the movie The Big Hit?	¿Quién dirigió la película Golpe maestro?
What is the genre of the film The Big Hit?	¿Cuál es el género de la película Golpe maestro?
What is the significance of Where are you, Dear General? in North Korean culture?	¿Cuál es la importancia de ¿Dónde estás, querido general? en la cultura norcoreana?

Table 5: Sample translation from English to Spanish (Gemini few-Shot)

Text	Prediction
What kind of artwork is The Signal-Man?	¿Qué tipo de obra de arte es El hombre de la señal?
When was The Signal-Man first published?	¿Cuándo se publicó por primera vez El hombre de la señal?
Who is the author of The Signal-Man?	¿Quién es el autor de El hombre de la señal?
Who are some of the main actors in the movie About Joan?	¿Quiénes son algunos de los actores principales de la película sobre Joan?
Who directed the 2022 film About Joan?	¿Quién dirigió la película de 2022 sobre Joan?
Can you provide a brief plot summary of About Joan?	¿Puedes proporcionar un breve resumen de la trama de About Joan?
Who are the main actors in The Big Hit?	¿Quiénes son los actores principales en El Gran Hit?
Who directed the movie The Big Hit?	¿Quién dirigió la película El Gran Hit?
What is the genre of the film The Big Hit?	¿Cuál es el género de la película El Gran Hit?
What is the significance of Where are you, Dear General? in North Korean culture?	¿Cuál es el significado de ¿Dónde estás, querido general? en la cultura norcoreana?

Table 6: Sample translation from English to Spanish (Finedtuned NLLB-200)