

Thai Winograd Schemas: A Benchmark for Thai Commonsense Reasoning

Phakphum Artkaew

Department of Electrical and Computer Engineering
Tandon School of Engineering
New York University
pa2497@nyu.edu

Abstract

Commonsense reasoning is one of the important aspects of natural language understanding, with several benchmarks developed to evaluate it. However, only a few of these benchmarks are available in languages other than English. Developing parallel benchmarks facilitates cross-lingual evaluation, enabling a better understanding of different languages. This research introduces a collection of Winograd Schemas in Thai, a novel dataset designed to evaluate commonsense reasoning capabilities in the context of the Thai language. Through a methodology involving native speakers, professional translators, and thorough validation, the schemas aim to closely reflect Thai language nuances, idioms, and cultural references while maintaining ambiguity and commonsense challenges. We evaluate the performance of popular large language models on this benchmark, revealing their strengths, limitations, and providing insights into the current state-of-the-art. Results indicate that while models like GPT-4 and Claude-3-Opus achieve high accuracy in English, their performance significantly drops in Thai, highlighting the need for further advancements in multilingual commonsense reasoning.

1 Introduction

Commonsense reasoning is an important challenge in artificial intelligence; however, most resources and benchmarks are in English, with only a few translations (Davis, 2023). The Winograd Schemas Challenge (WSC) has emerged as a widely adopted benchmark for evaluating the commonsense reasoning capabilities of language models (Levesque et al., 2012). WSC is featured in both GLUE and SuperGLUE (Wang et al., 2018, 2019). These benchmarks are widely used to evaluate a model’s understanding of general language. Significant progress has been made in developing commonsense reasoning models for high-resource languages like English. However, their performance

on low-resource languages remains largely unexplored. This research gap hinders the development and fair evaluation of multilingual and cross-lingual NLU systems. Linzen (2020) argues for developing multilingual training data and benchmarks to avoid English-centric model development, suggesting composite scores that average performance across languages to better evaluate systems.

There have been several one-to-one translations of the Winograd Schemas into various languages, including French (Amsili and Semineck, 2017), Portuguese (De Melo et al., 2019), Mandarin (Bernard and Han, 2020), Hebrew (Shwartz, 2024), Hungarian (Vadász and Ligeti-Nagy, 2022), and Russian (Shavrina et al., 2020). There are also various Japanese translations and adaptations (Shibata et al., 2015; Tanaka et al., 2013b,a). Before our translation, no Thai version of the Winograd Schemas existed. This lack prompted us to take the initiative.

The key contributions of our work are:

1. We introduce Thai-WS, the first benchmark dataset for evaluating the Winograd task and commonsense reasoning capabilities of language models and NLP systems in the Thai language.
2. We evaluate the performance of state-of-the-art models, including GPT-4, GPT-3.5, Claude-3-Haiku, Claude-3-Sonnet, Claude-3-Opus, Typhoon, and Command R+ on the Thai-WS dataset to assess their cross-lingual reasoning abilities.

2 Related Work

Background

The Winograd Schema Challenge, introduced by Levesque et al. (2012) in 2012 as an alternative Turing Test, consists of pairs of similar sentences with slight differences that introduce ambiguities

Example
<p>สมาชิกสภาเทศบาลเมืองปฏิเสธใบอนุญาตผู้ชุมนุมเพราะพวกเขากลัวความรุนแรง</p> <p>Samachik sapha thesaban mueang patiset bai anuyat phu chum num phro phuak khao klua khwam run raeng.</p> <p>The city councilmen refused the demonstrators a permit because they feared violence.</p>
<p>สมาชิกสภาเทศบาลเมืองปฏิเสธใบอนุญาตผู้ชุมนุมเพราะพวกเขาสนับสนุนความรุนแรง</p> <p>Samachik sapha thesaban mueang patiset bai anuyat phu chum num phro phuak khao sanapsanun khwam run raeng.</p> <p>The city councilmen refused the demonstrators a permit because they advocated violence.</p>
<p>มานะพยายามโทรศัพท์หาปิติ แต่เขาโทรไม่ติด</p> <p>Mana phayayam thoraphap ha Piti, tae khao thor mai tit.</p> <p>Paul tried to call George on the phone, but he wasn't successful.</p>
<p>มานะพยายามโทรศัพท์หาปิติ แต่เขาไม่สะดวก</p> <p>Mana phayayam thoraphap ha Piti, tae khao mai saduak.</p> <p>Paul tried to call George on the phone, but he wasn't available.</p>

Figure 1: Winograd Schema examples in Thai, with transliteration and corresponding English version.

resolved through general knowledge and logical reasoning. A classic example by Winograd (1972) illustrates the importance of context: "The city councilmen refused the demonstrators a permit because they feared/advocated violence." The choice of "feared" or "advocated" shifts the referent of "they" from the city councilmen to the demonstrators. This is because "feared" aligns with concerns typical of councilmen's reasons for refusal, whereas "advocated" suggests demonstrators' motivations, reflecting how context dictates meaning.

Other Thai Reasoning Benchmarks

To the best of our knowledge, three public benchmarks are available for Thai reasoning. First is XNLI, which is a translation of the MNLI benchmark into Thai, among 14 other languages (Conneau et al., 2018). Second is XCOPA, translated from the COPA benchmark to assess causal commonsense reasoning (Ponti et al., 2020). Third, the recently released M3Exam, is based on examination questions (Zhang et al., 2024). Unlike these benchmarks, the Winograd Schema Challenge is specifically designed to test commonsense reasoning through pronoun reference disambiguation, a task that requires resolving ambiguity using implicit world knowledge. This makes it fundamentally different and more challenging than tasks like XCOPA, which focus on causal reasoning but do not require the same level of subtle language understanding or contextual interpretation.

A Simple Method for Commonsense Reasoning

Trinh and Le (2018) introduced a simple approach to solve the Winograd Schema challenge by substituting one of the pronouns in the sentence and allowing the language model to determine which substitution has a higher probability. This method effectively frames Winograd Schemas as a binary classification task. Since then, many language models, such as GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020), have used this method to evaluate their performance on Winograd Schemas. However, due to limited access to models and computational power, the evaluation of this version of the Winograd Schema Challenge was conducted differently, using a prompt-based method. Details of this method are provided in the experimental setup section.

3 Dataset Construction

3.1 Translation

Two professional translators, who were native Thai speakers fluent in English and had experience translating from English to Thai, were hired. A one-to-one translation approach was followed to ensure that each schema was translated directly and accurately while preserving its original meaning. In a pilot translation phase, one native speaker translated the first 85 Winograd Schemas. Based on a qualitative analysis of these initial translations, guidelines were provided for a second native speaker to translate the remaining 200 schemas. In total, 285 Winograd Schemas were translated from English to Thai.

Translation guidelines were provided, instructing them to adapt names and contexts to suit the Thai language while preserving the ambiguity and nuances of the original schema. We also note that Thai pronouns are mostly similar to English but with a wider variety of formality levels. For example, in English, addressing an older individual with the pronoun "you" is generally acceptable. However, in Thai, maintaining politeness requires using a specific term of respect when referring to an older person. Therefore, we chose the translation based on the context of the sentence. The translators were also asked to mark any translated names and translations they were unsure about in red, so that the validator in the next step could pay extra attention to those instances. For example, in Figure 1, the names Paul and George were changed to Mana and Piti, respectively, adapting the names to better suit the Thai context while preserving the essence of the original content.

Some phrases could be directly translated, while others required adjustments to names and contextual adaptations to better suit the Thai language and culture. However, there were a few instances where the translators highlighted certain phrases in red, indicating that they are worth mentioning. The following includes two examples of schemas highlighted in red and one instance of a word adaptation made to better fit the nuances of the Thai language.

- (1) (i) Many people start to read Paul’s books and can’t put them down. They are gripped because Paul writes so well.
(ii) Many people start to read Paul’s books and can’t put them down. They are popular because Paul writes so well.
- (2) (i) During a game of tag, Ethan chased Luke because he was "it".
(ii) During a game of tag, Ethan ran from Luke because he was "it".
- (3) (i) Bob was playing cards with Adam and was way ahead. If Adam hadn’t had a sudden run of good luck, he would have won.
(ii) Bob was playing cards with Adam and was way ahead. If Adam hadn’t had a sudden run of good luck, he would have lost.

Most of the translation problems encountered involve examples like (1), where in English the word

"they" can refer to both people and objects. In this case, "they" can mean either "people" or "Paul’s books." However, in Thai, it is uncommon to use the same pronoun for both objects and people. Therefore, an alternative pronoun that is acceptable in Thai was chosen. In this instance, the English equivalent word to "those" was used instead.

In some cases, the meaning had to be adjusted to sound more natural in Thai by replacing idioms or figurative language. For example, in Thai, the phrase "he was" in (2) is sufficient to imply that he was the chaser in a game of tag, so the subject "it" is omitted from the sentence.

Throughout the translation process, the translators and validators prioritized making the text sound natural in Thai, which required word adaptations. Instead of strictly adhering to literal translations, wording was selected to better fit Thai language norms and context. For instance, in (3), while a direct translation of “good luck” is understandable to Thai speakers, the context of Bob and Adam playing cards made the Thai equivalent of “hand up” more suitable. Although “good luck” would be acceptable, “hand up” aligns more closely with the situation and feels more natural in Thai. Thus, “hand up” was chosen to maintain contextual and linguistic appropriateness.

3.2 Validation

The translated Winograd Schemas were reviewed by three native Thai speakers, and a validator was tasked with identifying potential issues, focusing on text flagged by the translators. Based on their feedback, final adjustments and typographical corrections were made. The dataset is publicly available¹ and consists of a test set containing 285 schemas, each with corresponding choices and answers.

3.3 Human baseline

The study for the Thai Winograd Schemas human baseline was conducted in a manner similar to that of Davis et al. (2016). A total of 30 native Thai speakers participated as volunteers and were divided into two groups of 15. The Winograd Schemas were split into two parts: one group completed part A, while the other group completed part B. Unlike Davis et al. (2016), who conducted their study in person, this study was conducted virtually using Google Forms. Participants were provided

¹https://huggingface.co/datasets/pakphum/winograd_th

Model	Accuracy (English)	Accuracy (Thai)	Accuracy (Thai Exact)
Typhoon	58.60%	56.14%	53.33%
Claude-3-Haiku	64.21%	53.33%	52.28%
Claude-3-Sonnet	80.70%	66.67%	65.96%
Claude-3-Opus	92.63%	79.65%	84.21%
GPT-3.5	70.88%	53.33%	52.28%
GPT-4	94.04%	76.49%	79.65%
Command-r-plus	87.02%	61.75%	64.91%
Human	92%	88%	-

Table 1: Accuracy vs. Model in English, Thai, and Thai Exact

with a link to complete their assigned part of the Winograd Schemas and were not restricted by time limits, allowing them to pause and resume the task as needed. This approach was designed to ensure participants could take their time and complete the study effectively. The observed score was approximately 88%, representing the average accuracy score of all participants.

4 Experimental Setup

Large language models were evaluated on Thai and English Winograd Schemas to assess their Thai language understanding and enable cross-linguistic comparisons of commonsense reasoning. The Winograd Schema Challenge comes in two versions: WSC-273, which contains 273 questions, and WSC-285, an extended version with 12 additional questions, totaling 285 (Kocijan et al., 2023). In this study, the WSC-285 dataset was used.

In addition to the main Thai dataset, the models were evaluated on the Thai-exact dataset, which consists of Winograd Schemas translated into Thai using Google Translate, with hand corrections by the authors to address translation mistakes. This evaluation aimed to compare the effect of adapting Thai names and contexts, as done in the main Thai dataset, against the more literal translations in Thai-exact.

The models chosen for evaluation were Typhoon (Pipatanakul et al., 2023), Claude-3 (Anthropic, 2024a), GPT-3.5 (Brown et al., 2020) and GPT-4 (Achiam et al., 2023), and C4AI Command R+ (Cohere, 2024). Typhoon is a Thai-specific model designed to handle Thai language tasks. Both GPT-3.5 and GPT-4 represent state-of-the-art general-purpose language models, with GPT-4 offering improved reasoning and comprehension capabilities over its predecessor. Claude-3 is another cutting-edge general-purpose language model excelling in

similar tasks. C4AI Command R+, on the other hand, is a multilingual model designed to support many languages. This selection was made to evaluate performance across Thai-specific, multilingual, and general-purpose capabilities. Further details about the models can be found in Appendix B.

Implementation and evaluation

Language models were evaluated using a prompt structure approach. A system prompt and user prompt structure were utilized to assess the models. A system prompt provides context and instructions to language models before a task, specifying the model’s role, personality, tone, or other relevant information to enhance its responses (Anthropic, 2024b). The system prompt was designed to prepare the model for the Winograd Schemas task, instructing it to respond with the correct answer.

Evaluation was conducted by calculating accuracy, with only responses that exactly matched the correct answers in the schemas considered correct. A manual review was performed to address cases where correct answers were presented with additional text, such as “the answer is the city councilmen” instead of “the city councilmen.” This correction process was only required for Typhoon, as it occasionally included additional information in its responses. The models were evaluated via their respective APIs. Further details on the evaluation process and prompt setup are provided in Appendix A.

5 Result and Discussion

Table 1 presents the accuracy results for each evaluated model. Several human baselines exist for the English Winograd Schema Challenge (Kocijan et al., 2023). Our English human baseline is derived from the study conducted by Davis et al. (2016), which involved human participants evaluat-

ing Winograd Schemas in English. In [Davis et al. \(2016\)](#) research, human accuracy was observed to be approximately 92%.

Comparison of models

GPT-4 leads with the highest accuracy in English at 94.04%, while Claude-3-Opus has the best performance in Thai Exact at 84.21%. In the Thai language context, Claude-3-Opus again performs the best with 79.65%, followed closely by GPT-4 at 76.49%. Other models, such as GPT-3.5 and Claude-3-Haiku, show significantly lower accuracy, particularly in the Thai and Thai Exact columns, indicating a drop in performance when these models are applied to languages other than English. Human performance is slightly lower in Thai, but the gap is larger for language models, highlighting challenges in language adaptation.

The performance of the Thai dataset versus the Thai-Exact dataset is relatively similar, with less than a five percent difference across all models. This may suggest that adapting Thai names and nuances does not have a significant impact on model performance. Overall, the results indicate that all models perform better on the English version of the dataset compared to the Thai versions, whether adapted by humans or translated by machines. The performance drop in Thai suggests that these models struggle more with the linguistic nuances and structures of the Thai language. The performance drop may also be due to data leakage ([Elazar et al., 2021](#)), which could affect the integrity of the results.

Do larger models consistently perform better

The analysis of model parameter sizes for GPT and Claude is derived from the assumptions provided by [Coda-Forno et al. \(2024\)](#). According to their study, GPT-4 has approximately 1760 billion parameters, while Claude-3-opus has around 300 billion parameters. The results indicate that larger models tend to perform better on this task, with performance consistently improving as model size increases. Specifically, the smallest model in the analysis, Typhoon (based on LLaMA 8B), demonstrated the lowest performance. Performance improved with larger models, starting from Command-r-plus (104B) and continuing to larger models like Claude and GPT. However, in the context of Thai, larger models do not always perform better, as Claude-3-Opus outperformed the larger GPT-4. This finding suggests that while model size

is a contributing factor, other elements, such as training data quality, multilingual capabilities, or architectural design, may also significantly impact performance.

Are the mistakes similar to those made by English LLMs?

To understand whether the errors stem from language understanding or commonsense reasoning, we further analyze the models' output by examining the overlapping incorrect questions in both the English and Thai sets. By observing the percentage of consistently incorrect overlapping questions from the English set to the Thai set (i.e., the number of overlapping incorrect answer in both sets divided by the number of incorrect answer in the English set), we find that the percentage usually exceeds half of the incorrect questions in the Thai set for models such as Typhoon, GPT-3.5, Claude-3-haiku, Claude-3-sonnet, and command-r-plus. This suggests that most models struggle with commonsense reasoning in general, while the remaining percentage may be attributed to language understanding. This pattern may not hold for GPT-4 and Claude-3-opus, as the percentage of consistent incorrect questions falls below 40%, suggesting that these models may exhibit better commonsense understanding but face challenges in Thai language understanding. The full details can be found in [appendix C](#).

6 Conclusion

In conclusion, the Thai Winograd Schemas benchmark represents a noteworthy step in evaluating commonsense reasoning capabilities within the Thai language context. This novel dataset, meticulously developed and validated by native speakers and professional translators, aims to preserve linguistic and cultural nuances unique to Thai. The comprehensive evaluation of state-of-the-art language models, including GPT-4, Claude-3 variants, Typhoon, and Command R+ on both English and Thai versions of the Winograd Schema Challenge offers insights into their cross-lingual performance. The observed performance drop in Thai suggests challenges these models may face in handling low-resource languages, indicating a need for further research and development in multilingual natural language processing.

Limitations

While cultural nuances are preserved as much as possible during the translation process, it is acknowledged that complete preservation is not always achievable due to differences between Thai and English. This linguistic difference may also contribute to the slightly lower performance observed for the Thai human baseline compared to its English counterpart. Direct evaluation on large language models like GPT-4 and Claude-3 cannot be performed due to lack of access and insufficient computational power. Therefore, an alternative approach using prompt-based evaluation is adopted.

Acknowledgments

Thank you to the two translators, Chanikarn Inthongpan and Korakoch Rienmek, who worked on translating these Winograd schemas. It was challenging to find individuals skilled enough for this specific task, which demands a deep understanding of language. I also thank the validators Sakrapee Namsak and Chonnawit Khumchoo for their careful review and feedback that greatly contributed to the quality and accuracy of the final dataset. Finally, thank you to Professor Ernest Davis, Professor Chinmay Hegde, and Vid Kocijan for their helpful feedback and support throughout this process.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Pascal Amsili and Olga Seminck. 2017. A google-proof collection of french winograd schemas. In *The 2nd Workshop on Coreference Resolution Beyond OntoNotes (CORBON 2017)*, co-located with *EACL 2017*, pages 24–29.
- Anthropic. 2024a. Introducing the next generation of claude. <https://www.anthropic.com/news/claude-3-family>. Accessed: 2024-05-17.
- Anthropic. 2024b. System prompts documentation. <https://docs.anthropic.com/en/docs/system-prompts>. Accessed: 2024-05-17.
- Timothée Bernard and Ting Han. 2020. *Mandarino-grad: A Chinese collection of Winograd schemas*. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 21–26, Marseille, France. European Language Resources Association.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Julian Coda-Forno, Marcel Binz, Jane X Wang, and Eric Schulz. 2024. *CogBench: a large language model walks into a psychology lab*. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 9076–9108. PMLR.
- Cohere. 2024. Command r+ with microsoft azure. <https://cohere.com/blog/command-r-plus-microsoft-azure>. Accessed: 2024-12-13.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Ernest Davis. 2023. Benchmarks for automated commonsense reasoning: A survey. *ACM Computing Surveys*, 56(4):1–41.
- Ernest Davis, Leora Morgenstern, and Charles Ortiz. 2016. Human tests of materials for the winograd schema challenge 2016. URL: <https://cs.nyu.edu/faculty/davise/papers/WS2016SubjectTests.pdf>.
- Gabriela De Melo, Vinicius Imaizumi, and Fábio Cozman. 2019. Winograd schemas in portuguese. In *Anais do XVI Encontro Nacional de Inteligência Artificial e Computacional*, pages 787–798. SBC.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021. *Back to square one: Artifact detection, training and commonsense disentanglement in the Winograd schema*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10486–10500, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Vid Kocijan, Ernest Davis, Thomas Lukasiewicz, Gary Marcus, and Leora Morgenstern. 2023. The defeat of the winograd schema challenge. *Artificial Intelligence*, page 103971.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*.

- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Kunat Pipatanakul, Phatrasek Jirabovonvisut, Potsawee Manakul, Sittipong Sripaisarnmongkol, Ruangsak Patomwong, Pathomporn Chokchainant, and Kasima Tharnpipitchai. 2023. Typhoon: Thai large language models. *arXiv preprint arXiv:2312.13951*.
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020. [XCOPA: A multilingual dataset for causal commonsense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Tatiana Shavrina, Alena Fenogenova, Anton Emelyanov, Denis Shevelev, Ekaterina Artemova, Valentin Malykh, Vladislav Mikhailov, Maria Tikhonova, Andrey Chertok, and Andrey Evlampiev. 2020. Russiansuperglue: A russian language understanding evaluation benchmark. *arXiv preprint arXiv:2010.15925*.
- Tomohide Shibata, Shotaro Kohama, and Sadao Kurohashi. 2015. Nihongo winograd schema challenge no kouchiku to bunseki. *Proceedings of NLP2015*, pages 493–496.
- Vered Shwartz. 2024. Winograd schema challenge datasets. https://www.cs.ubc.ca/~vshwartz/resources/winograd_he.html. Accessed: 2024-05-17.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. Asynchronous pipeline for processing huge corpora on medium to low resource infrastructures. In *7th Workshop on the Challenges in the Management of Large Corpora (CMLC-7)*. Leibniz-Institut für Deutsche Sprache.
- Soichiro Tanaka, Rafal Rzepka, and Shiho Katajima. 2013a. Translation preserving english names. http://arakilab.media.eng.hokudai.ac.jp/~kabura/collection_katakana.html. Translation preserving English names, PDF and HTML formats.
- Soichiro Tanaka, Rafal Rzepka, and Shiho Katajima. 2013b. Translations into japanese. http://arakilab.media.eng.hokudai.ac.jp/~kabura/collection_ja.html. Translation changing English names to Japanese, PDF and HTML formats.
- Trieu H Trinh and Quoc V Le. 2018. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*.
- Noémi Vadász and Noémi Ligeti-Nagy. 2022. Winograd schemata and other datasets for anaphora resolution in hungarian. *Acta Linguistica Academica*, 69(4):564–580.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Terry Winograd. 1972. Understanding natural language. *Cognitive psychology*, 3(1):1–191.
- Wenxuan Zhang, Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2024. M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models. *Advances in Neural Information Processing Systems*, 36.

A Prompt Evaluation

To ensure consistency and reproducibility in evaluating each model, specific settings were implemented alongside the prompt structure, as illustrated in Figure 2.

The prompt follows a consistent pattern to ensure clarity and replicability across evaluations:

1. **Sentence with Pronoun:** The main sentence is provided, with the ambiguous pronoun enclosed in asterisks (**), followed by a newline character (\backslash n).
2. **Snippet:** A shortened snippet containing the pronoun and its immediate context is included after the sentence, labeled as “Snippet: ”, followed by a newline character (\backslash n).
3. **Options:** The candidate antecedents for the pronoun are listed after “Options:”, separated by newline characters (\backslash n), with no newline after the final option.

This process was executed through the model API, where each model was prompted with the designed prompts, and the answers were obtained directly from the model's output. This ensured that all evaluations followed the same method and settings across models. The code used to reproduce all of our experimental results is available at the following GitHub repository: <https://github.com/PhakphumAdev/Thai-Winograd>.

B Models Used in the Study

Typhoon

The Thai large language model, named Typhoon, was initially built on the Mistral-7B architecture (Pipatanakul et al., 2023). Typhoon was pre-trained on the MC4 (Raffel et al., 2020) and OSCAR (Suárez et al., 2019) datasets, which include approximately 3 TB of Thai text. In this experiment, the specific model tested was Typhoon-Instruct, as documented on the Open Typhoon website². This version of Typhoon, tested in September 2024, is based on the LLaMA 3 8B architecture (Dubey et al., 2024). It was released on September 5, 2024, and is distinct from versions built on the Mistral-7B architecture.

Claude

Anthropic's large language model (Anthropic, 2024a), designed to be a helpful and honest assistant, has its model architecture details kept private. In this experiment, all the Claude-3 models were tested to compare performance. The specific models tested were claude-3-haiku-20240307, claude-3-sonnet-20240229, and claude-3-opus-20240229.

GPT

GPT or Generative Pre-trained Transformer is one of the most popular and powerful large language models available. While the specific details of GPT-4's architecture are not publicly disclosed, it can be assumed that it represents an advancement over GPT-3. GPT-3 is an autoregressive, decoder-only model with 175 billion parameters (Brown et al., 2020). Although the technical report does not specify its performance in Thai specifically, it asserts that GPT-4 surpasses GPT-3 in the MMLU task for the Thai language (Achiam et al., 2023). In this experiment, gpt-4-0613 and gpt-3.5-turbo-0125 were tested.

²<https://docs.opentyphoon.ai>

System Prompt:

"You will be provided with a sentence and a snippet containing a pronoun enclosed in asterisks (**). Your task is to determine the correct referent of the pronoun from the given options. Respond only with one of the provided choices, exactly as it is written. For example, if the options are 'The city councilmen' and 'The demonstrators', respond only with 'The city councilmen' or 'The demonstrators'."

User Prompt:

"The city councilmen refused the demonstrators a permit because **they** feared violence.\nSnippet: **they** feared violence\nOptions:\nThe city councilmen\nThe demonstrators"

Accepted Answer:

"The city councilmen"

Unaccepted Answer:

"The answer is The city councilmen"

Figure 2: An example of the prompt evaluation method, detailing the system prompt, user prompt, and expected answer. Only exact matches like "The city councilmen" were considered correct, while responses such as "The answer is The city councilmen" were not accepted, ensuring consistent and reproducible evaluations.

C4AI Command R+

C4AI Command R+ is a 104-billion parameter model with advanced capabilities, including Retrieval Augmented Generation (RAG) and multi-step tool use for automating complex tasks (Cohere, 2024). Optimized for reasoning, summarization, and question answering, it excels across various use cases. In this experiment, command-r-plus-08-2024 were tested.

C Consistency of errors in English LLMs

Model	% for incorrect overlapping answers (Thai)	% for incorrect overlapping answers (English)
Typhoon	56.80%	60.17%
Claude-3-Haiku	51.88%	67.65%
Claude-3-Sonnet	33.68%	58.18%
Claude-3-Opus	15.52%	42.86%
GPT-3.5	34.59%	55.42%
GPT-4	8.96%	35.29%
Command-r-plus	22.94%	67.57%

Table 2: Consistency of errors