

Terminology-Aware Translation with Constrained Decoding and Large Language Model Prompting

Nikolay Bogoychev*

Pinzhen Chen*

School of Informatics, University of Edinburgh
n.bogoych@ed.ac.uk, pinzhen.chen@ed.ac.uk

Abstract

Terminology correctness is important in the downstream application of machine translation, and a prevalent way to ensure this is to inject terminology constraints into a translation system. In our submission to the WMT 2023 terminology translation task, we adopt a translate-then-refine approach which can be domain-independent and requires minimal manual efforts. We annotate random source words with pseudo-terminology translations obtained from word alignment to first train a terminology-aware model. Further, we explore two post-processing methods. First, we use an alignment process to discover whether a terminology constraint has been violated, and if so, we re-decode with the violating word negatively constrained. Alternatively, we leverage a large language model to refine a hypothesis by providing it with terminology constraints. Results show that our terminology-aware model learns to incorporate terminologies effectively, and the large language model refinement process can further improve terminology recall.

1 Introduction

One of the major obstacles encountered by neural machine translation (NMT) systems pertains to the utilization of suitable domain-related words when translating specialized content not present in the training data. An illustrative instance of this challenge arises when translating “transformer” from English into another language, where the accurate translation depends on the context or the preference of the audience (Figure 1). A straightforward literal translation approach often leads to suboptimal outcomes, prompting human translators unfamiliar with domain-specific knowledge to resort to reference materials for terminology precision. This issue is prevalent in the translation industry, with many commercial translation service providers offering paid solutions to address it. Furthermore, it

*Equal contribution.



Translate "transformer" to Chinese?

变压器 (electric transformer)

变形金刚 (the Transformer character)

变换器 (something that changes)



Figure 1: Terminology hints can help disambiguate polysemantic words when translating with limited context.

is a popular area in machine translation research, indicated by efforts such as WMT shared tasks organization and participation focusing on terminology and domain-specific translations (Alam et al., 2021; Bawden et al., 2019, 2020, inter alia).

This year’s WMT terminology translation task features three language directions: German-to-English, Chinese-to-English, and English-to-Czech. In addition to reading in a source sentence, participating systems need to employ a provided dictionary, which contains source-target terminology word mappings, to incorporate into the target translation. For each source sentence in the test set, there are three modes of applying terminology constraints:

1. *Terminology* constraint: Dictionaries of real terminology words are provided, to be incorporated in the translations.
2. *Random* constraint: Random (but presumably correct) word mappings are obtained using a word alignment tool and provided as a pseudo-terminology dictionary.
3. *No* constraint: Source sentences can be freely translated without external information.

We interpret that the no-constraint setting allows us to measure the competing systems’ quality and understand to what degree the systems effectively utilize the provided random and terminology dictionaries. Our baseline approach is to train

a terminology-aware translation (TAT) system inspired by [Dinu et al. \(2019\)](#), where, in the training data, source words are tagged with desired translations inline on the source side. Then we propose two separate refinement strategies on top of it to aggressively encourage the appearance of terminologies:

1. We use a neural word aligner to identify terminology constraints missed by the baseline system, and use the same system to re-decode the source by negatively constraining (disallowing) previously incorrectly translated tokens.
2. We also investigate the capability of a large language model to simultaneously paraphrase an existing translation to include the desired terminology constraints via curated prompts.

Our proposed techniques can incorporate target terminology words with around 80% recall, using automatic and soft constraints in a two-step refinement process. We observe that for German-English, our terminology-aware training and negatively constrained decoding perform better, whereas, for Chinese-English and English-Czech, LLM-based refinement achieves higher scores. In terms of overall translation accuracy, we find that negatively constrained decoding could lead to a tiny drop and LLMs are able to maintain or improve quality according to a reference-free neural metric.

2 Related Work

Previous research on terminology translation could be divided into two categories: *soft* constraint and *hard* constraint, depending on whether the resulting translation system will enforce the appearance of desired target translations. In the soft constraint setting, the convention is to train a model that is able to ingest the target terminology words inline, directly placing them after the corresponding source words in the source input ([Dinu et al., 2019](#)). Many later implementations stem from this to include new elements such as additional lemmatization ([Bergmanis and Pinnis, 2021](#)) or grammatical error correction ([Pham et al., 2021](#)) as a post-processing step in order to achieve a more fluent output. Instead of placing the target constraint words inline, some other works train a system that takes the terminology constraint as either a prefix or a suffix ([Jon et al., 2021](#); [Turcan et al., 2022](#)).

Most hard constraint work involves post-processing a translation with desired terminologies. [Post et al. \(2019\)](#) inserted untranslatable tokens (also known as placeholders) into the source, which will remain unchanged through the translation process. Then the placeholders are replaced with terminology words in the target language. This is entirely performed as a post-processing step. Such terminology replacement could also be done by keeping and replacing the source word at inference time, and it is also feasible to run target word replacement as post-processing ([Molchanov et al., 2021](#)). A hard constraint method guarantees that the chosen terminology token will appear, but often results in less fluent output, especially for morphologically rich languages because the context is not taken into consideration during replacement. It also mandates more complicated post-processing than the soft constraint approaches.

Our first post-processing proposal relies on constrained decoding, which refers to either allowing certain tokens or blocking specific tokens during inference time ([Hokamp and Liu, 2017](#)). It has been applied to terminology injection, paraphrasing, parallel sentence mining, etc ([Hasler et al., 2018](#); [Kajiwara, 2019](#); [Chen et al., 2020](#)). We opt for negatively constraining the tokens that violated the given terminology alignments by preventing them from entering the hypothesis beam in the refinement stage. These alignments are computed using word alignment tools ([Dyer et al., 2013](#); [Dou and Neubig, 2021](#)).

Another post-processing method in our study prompts an LLM to refine a translation and incorporate terminology terms simultaneously. Whilst previous studies have explored the translation capability of LLMs ([Vilar et al., 2023](#); [Zhang et al., 2023](#)), the works closely relevant to us are from [Moslem et al. \(2023\)](#) and [Ghazvininejad et al. \(2023\)](#). We adopt the paradigm from the latter, which re-words a constraint dictionary as a natural text and affixes it into a translation prompt. While they focused on rare words without directly benchmarking on terminology translation, our post-processing step can be seen as an extension of word-level controlled prompting to terminology translation with large language models. Both of our post-processing methods should be categorized as soft constraint approaches since there is no guarantee that negatively constrained decoding or an LLM will necessarily incorporate the constraints in a re-generation.

3 Terminology-Aware Training

The goal of our system implementation is to create a general-purpose terminology-aware translation system that is unsupervised and domain-agnostic, and requires the minimum effort of pre- and post-processing.

3.1 Terminology creation

Inspired by [Dinu et al. \(2019\)](#), we applied terminology constraints during training, but a key difference is that, unlike their approach, we assume that we have no access to downstream domain or terminology constraints during training, in order to build a general-purpose domain-agnostic system. Consequently, we have no curated terminology data to use. Therefore, we generate (pseudo-)terminology information using word alignments. Our workflow can be detailed as:

1. We compute the word alignment information for the entire training set using `fast_align` ([Dyer et al., 2013](#)).
2. For each sentence, we select all bijective source-target mappings as our terminology candidates. We also filter out trivial mappings where the source and target tokens are the same (e.g. numbers, names), because those mappings are simple and hence likely to be correctly translated by a translation system even without any terminology awareness.
3. In the training data, we replace `srcwordi` in the source sentence with:
`srcwordi __target__ trgwordj __done__`
where the `srcwordi` is the i -th source word inside the sentence, and `trgwordj` is the word inside the target sentence, corresponding to `srcwordi` according to word alignment information. This replacement occurs with around 10% probability for each candidate source-target pair. For a sentence that does not have an associated terminology constraint, the data is the same as normal NMT.
4. At inference time, we process the test data similarly to above, except that the source-target word mapping comes from a supplied terminology dictionary.

In practice, our translation system is trained with a mix of normal translation data and terminology-injected data. The advantage of this strategy is that

the trained models are general-purpose, so they can translate normal texts without terminology injection. Further, they have been exposed to a wide variety of constraints during training, making them robust to potentially unseen domain constraints.

Overall, our method is very similar to [Bergmanis and Pinnis \(2021\)](#)'s work, except that we use whole words but not lemmas to ease pre-processing. We presume that the language model will be able to adjust the terminologies accordingly, especially for morphologically rich languages on the target side. This enables our method to be trivially transferable across languages.

Finally, our systems could easily be turned into hard-constrained by replacing the source word with the desired target terminology word. This could be feasible because our terminology-aware training installs the copying behaviour in the neural translation model, although in this mode the model would produce markedly less fluent output.

3.2 Model architecture

We trained Transformer-style machine translation models ([Vaswani et al., 2017](#)) using the Marian NMT toolkit ([Junczys-Dowmunt et al., 2018](#)). We used the Transformer-Big preset which is a 6 encoder, 6 decoder architecture with 1024 hidden size, and 4096 feedforward size.¹

3.3 Data

The terminology task uses the same data as the constrained condition in the WMT23 general translation task. We carefully cleaned, filtered, and deduplicated the available WMT training sets provided by the organisers, as well as the available back-translation data. After preprocessing we were left with the following:

- German-to-English (de-en): 199M lines of parallel data and 29.5M lines of back-translated data.
- Chinese-to-English (zh-en): 21.8M lines of parallel data and 15.6M lines of back-translated data.
- Czech-to-English (cs-en): 61.8M lines of parallel data and 57M lines of back-translated data.

¹<https://github.com/marian-nmt/marian/blob/master/src/common/aliases.cpp#L114>

Query	Prompt template
Translation	Source: $\{\text{source}\}$ Please give me a translation in $\{\text{lang}\}$ without any explanation.
Refinement	Source: $\{\text{source}\}$ Translation: $\{\text{translation}\}$ Please give me a better $\{\text{lang}\}$ translation without any explanation. “ $\{\text{srcword}_0\}$ ” should be translated as “ $\{\text{trgword}_0\}$ ”; “ $\{\text{srcword}_1\}$ ” should be translated as “ $\{\text{trgword}_1\}$ ”; ... “ $\{\text{srcword}_k\}$ ” should be translated as “ $\{\text{trgword}_k\}$ ”. (<i>with $k \geq 0$</i>)

Table 1: Large language model prompt templates for unconstrained and constrained translation.

3.4 General quality

The quality of our models without terminology translation is shown in Table 2, where we report BLEU (Papineni et al., 2002) and COMET_{DA}² (Rei et al., 2020) scores on test sets from the WMT22 general translation task. We note that terminology augmentation during training could result in a slight quality drop.

	BLEU	COMET _{DA}
de-en	31.3	0.8334
en-cs	39.5	0.8715
zh-en	20.3	0.7559

Table 2: Performance of our terminology-aware translation systems in the WMT22 general translation task.

4 Post-Translation Terminology Injection

Despite training our model with terminology awareness, there is no mechanism to ensure that the desired terminology constraint will appear on the target side. The neural network decoding behaviour is not entirely predictable, especially given the assumption of no additional domain adaptation. Below, we present two distinct strategies to try *harder* to promote the terminology constraints, via automatic post-editing through constrained beam search and large language models.

4.1 Negatively constrained decoding

While it is easy enough to notice when a target terminology term is not generated as per a given constraint, it is not trivial to understand which word

²_{wmt22-comet-da}. This is a reference-based metric which requires the source input, hypothesis, and reference.

has been produced in place of the desired term. In order to do this, we make use of *awesome-align*, a neural multilingual word aligner (Dou and Neubig, 2021), with the following procedure:

1. For each source-translation pair, we check if all required terminology terms appear on the target side. If they do, then we stop processing more rules.
2. Then, we use *awesome-align* to compute word alignments and detect the word(s) that have been generated in place of the desired terms according to the provided terminology constraints.
3. We decode the source sentence again, penalising the words that violated the terminology constraint, by forbidding the decoder from generating them at each generation step, unless they carry more than 95% of the probability mass at a certain step.

In practice, this procedure can be repeated infinitely, until all terminology constraints are fulfilled, but we decided to limit it to only one iteration, to keep this a realistic production scenario in terms of computational budget.

4.2 Large language models

Recent years saw the rise of large language models (LLMs), which have a strong capability in various NLP tasks. In this paper, we investigate the effectiveness of using a large language model to generate terminology terms during translation by adding constraints to Chen et al. (2023)’s translation refinement prompts. We use two distinct prompts: free translation and translation refinement queries. The translation query sends a source sentence and

Mode	Model	Refine	de→en		zh→en		en→cs	
			Recall	COMET _{QE}	Recall	COMET _{QE}	Recall	COMET _{QE}
<i>terminology</i> constraints	TAT	-	82.30	.0797	49.98	-.0896	73.75	.0601
	TAT	NCD	82.01	.0775	50.42	-.0903	73.26	.0588
	TAT	LLM	64.35	.1197	83.06	.0185	76.00	.0866
	LLM	-	41.86	.1244	46.63	.0191	48.14	.0913
	LLM	LLM	70.48	.1180	81.01	.0201	78.94	.0882
<i>no</i> constraint [†]	TAT	-	39.82	.1085	13.64	-.1163	48.11	.0712
	TAT	LLM	39.59	.1251	42.76	.0203	47.31	.0955
	LLM	-	41.86	.1244	46.63	.0191	48.14	.0913
	LLM	LLM	39.65	.1258	46.72	.0228	46.22	.0943
<i>random</i> constraints	TAT	-	76.17	.0716	81.55	-.1105	57.10	.0502
	TAT	NCD	75.79	.0698	82.03	-.1123	56.42	.0465
	TAT	LLM	61.46	.1206	63.17	.0175	70.97	.0875
	LLM	-	38.70	.1244	52.49	.0191	39.34	.0913
	LLM	LLM	66.74	.1188	67.10	.0196	73.37	.0867
<i>no</i> constraint [‡]	TAT	-	35.60	.1085	36.18	-.1163	37.35	.0712
	TAT	LLM	37.58	.1251	49.48	.0203	39.03	.0955
	LLM	-	38.70	.1244	52.49	.0191	39.34	.0913
	LLM	LLM	37.62	.1258	49.00	.0228	38.42	.0943

[†]Recall computed against terminology constraints.

[‡]Recall computed against random constraints.

Table 3: Terminology recall and translation quality measured by COMET_{QE} of our systems on the *blind test* set. TAT: terminology-aware translation; NCD: negatively constrained decoding; LLM: large language model.

requests a translation in the target language without any other information. On the other hand, the refinement query feeds back an unconstrained translation together with terminology constraints to request a new translation. This essentially forms an LLM version of the constrained beam search discussed in Section 4.1. The constraints are enforced through natural language instructions in the prompts, under the situation where the softmax distribution from an LLM is not accessible by users.

The LLM we use is OpenAI’s GPT-3.5.³ It is a closed-source commercial system, where the model weights and the inference states are not available to users. The model has a context window of 4096 which is sufficient to cover an instruction, a source sentence, several terminology constraints, as well as the target translation. It is public to all users at a relatively cheap cost. In our settings, each translation is carried out in a new query session.

In Table 1 we outline the two prompt templates we used. During querying, the placeholder variables are substituted with corresponding string val-

ues. For the refinement query, when a terminology dictionary is supplied, the source and target words are fed to the LLM via the prompt (Ghazvininejad et al., 2023); if there is no terminology dictionary, the query simply asks for a refined translation. The two-step experiment with LLMs can be summarized as follows:

1. We obtain an initial unconstrained translation, which may or may not fulfil all the terminology constraints. It can come from either the LLM itself or the terminology-aware translation model built in Section 3.1.
2. We query the LLM with the constrained translation prompt to obtain a refined translation with terminology incorporated in the prompt.

5 Results and Discussions

We present our *blind test* results in Table 3, which include both terminology recall and COMET_{QE} scores computed by us.⁴ We used COMET_{QE} in particular because it does not require references

³gpt-3.5-turbo-0613, a snapshot of the GPT-3.5 model on 13 June 2023

⁴wmt21-comet-da-qe

which are not accessible to us. We assess the effectiveness of our methods by comparing the terminology recall of our systems with and without applying terminology constraints, in both *random* and *real terminology* scenarios.

5.1 Translation quality

In terms of translation quality reflected in COMET_{QE} , we observe that the LLM rows attain superior results, which is not surprising considering that we use an unconstrained commercial model GPT-3.5. By comparing TAT with TAT+NCD, or comparing LLM with LLM+LLM under a constrained scenario, we conclude that applying terminology constraints usually lead to a sacrifice in translation quality regardless of the language direction or the systems involved. Nonetheless, as a contrasting experiment with no constraint, LLM+LLM achieves a slightly better COMET_{QE} score than using an LLM to translate without refinement.

Our model performed poorly on the zh-en task in terms of COMET_{QE} scores. We suspect that this is because of the domain mismatch between the translation data from the general domain and the Chinese terminology test set. Upon manual inspection, we found that the latter includes web novels and literal writing which are likely to be under-represented in the generic training data.

5.2 Terminology recall

Focusing on terminology generation, compared with TAT or LLM in unconstrained settings, TAT marks 30-40 higher recall of terminology terms in the constrained *terminology* and *random* settings. This indicates that our terminology-aware training is effective in teaching translation models to follow customized source-target word alignments.

Next, as a post-processing step, negatively constrained decoding seems to be disappointing in practice. TAT+NCD often produces worse results than TAT alone in terms of both quality and terminology recall, except for zh-en with *random* constraints. We hypothesize that this could be due to two problems: (1) word alignment errors could propagate into this process, and (2) by applying NCD, we might capture a missed terminology term but at the cost of mis-translating other words. Our constraining procedure might be improved by performing shortlisting, namely positively constrained decoding, as opposed to negatively limiting the beam search in an iterative approach.

We find the results promising when using LLMs for terminology injection. Looking at LLM+LLM versus LLM alone in various constrained conditions, terminology recall improves significantly with very little drop in overall quality. Also by comparing TAT+LLM with TAT alone, we observe that TAT and LLMs each have their own merits depending on the language direction. In terms of recall, TAT wins in de-en, TAT+LLM wins in zh-en, and they are close in en-cs. However, TAT+LLM is way ahead if measured by COMET_{QE} . However, we must note that an LLM costs significantly more resources than a dedicated translation model at both training and inference time.

6 Conclusion and Future Work

We participated in all tracks of the WMT 2023 terminology shared task with a terminology-aware translation baseline, and two distinct refinement procedures using negatively constrained beam search and large language models separately. The results we produced gave us insights into the pros and cons of our systems. In future work, we could explicitly enforce the generation of the terminology token by identifying the appropriate time step and manipulating the probability distribution after softmax computation, even in an open-source large language model. This is not entirely trivial due to the presence of subwords but could be achievable.

Acknowledgement

This project has received funding from UK Research and Innovation (UKRI) under the UK government’s Horizon Europe funding guarantee [grant numbers 10052546 and 10039436].

Some computations described in this research were performed using the Baskerville Tier 2 HPC service (<https://www.baskerville.ac.uk/>). Baskerville was funded by the EPSRC and UKRI through the World Class Labs scheme (EP/T022221/1) and the Digital Research Infrastructure programme (EP/W032244/1) and is operated by Advanced Research Computing at the University of Birmingham.

References

Md Mahfuz Ibn Alam, Ivana Kvapilíková, Antonios Anastasopoulos, Laurent Besacier, Georgiana Dinu, Marcello Federico, Matthias Gallé, Kweonwoo Jung, Philipp Koehn, and Vassilina Nikoulina. 2021. [Find-](#)

- ings of the WMT shared task on machine translation using terminologies. In *Proceedings of WMT*.
- Rachel Bawden, Kevin Bretonnel Cohen, Cristian Grozea, Antonio Jimeno Yepes, Madeleine Kittner, Martin Krallinger, Nancy Mah, Aurelie Neveol, Mariana Neves, Felipe Soares, Amy Siu, Karin Verspoor, and Maika Vicente Navarro. 2019. Findings of the WMT 2019 biomedical translation shared task: Evaluation for MEDLINE abstracts and biomedical terminologies. In *Proceedings of WMT*.
- Rachel Bawden, Giorgio Maria Di Nunzio, Cristian Grozea, Inigo Jauregi Unanue, Antonio Jimeno Yepes, Nancy Mah, David Martinez, Aurélie Névéal, Mariana Neves, Maite Oronoz, Olatz Perez-de Viñaspre, Massimo Piccardi, Roland Roller, Amy Siu, Philippe Thomas, Federica Vezzani, Maika Vicente Navarro, Dina Wiemann, and Lana Yeganova. 2020. Findings of the WMT 2020 biomedical translation shared task: Basque, Italian and Russian as new additional languages. In *Proceedings of WMT*.
- Toms Bergmanis and Mārcis Pinnis. 2021. Facilitating terminology translation with target lemma annotations. In *Proceedings of EACL*.
- Pinzhen Chen, Nikolay Bogoychev, Kenneth Heafield, and Faheem Kirefu. 2020. Parallel sentence mining by constrained decoding. In *Proceedings of ACL*.
- Pinzhen Chen, Zhicheng Guo, Barry Haddow, and Kenneth Heafield. 2023. Iterative translation refinement with large language models. *arXiv preprint*.
- Georgiana Dinu, Prashant Mathur, Marcello Federico, and Yaser Al-Onaizan. 2019. Training neural machine translation to apply terminology constraints. In *Proceedings of ACL*.
- Zi-Yi Dou and Graham Neubig. 2021. Word alignment by fine-tuning embeddings on parallel corpora. In *Proceedings of EACL*.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of NAACL-HLT*.
- Marjan Ghazvininejad, Hila Gonen, and Luke Zettlemoyer. 2023. Dictionary-based phrase-level prompting of large language models for machine translation. *arXiv preprint*.
- Eva Hasler, Adrià de Gispert, Gonzalo Iglesias, and Bill Byrne. 2018. Neural machine translation decoding with terminology constraints. In *Proceedings of NAACL-HLT*.
- Chris Hokamp and Qun Liu. 2017. Lexically constrained decoding for sequence generation using grid beam search. In *Proceedings of ACL*.
- Josef Jon, Michal Novák, João Paulo Aires, Dusan Varis, and Ondřej Bojar. 2021. CUNI systems for WMT21: Terminology translation shared task. In *Proceedings of WMT*.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. Marian: Fast neural machine translation in C++. In *Proceedings of ACL*.
- Tomoyuki Kajiwara. 2019. Negative lexically constrained decoding for paraphrase generation. In *Proceedings of ACL*.
- Alexander Molchanov, Vladislav Kovalenko, and Fedor Bykov. 2021. PROMT systems for WMT21 terminology translation task. In *Proceedings of WMT*.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. Adaptive machine translation with large language models. In *Proceedings of EAMT*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of ACL*.
- Minh Quang Pham, Josep Crego, Antoine Senellart, Dan Berrebbi, and Jean Senellart. 2021. SYSTRAN @ WMT 2021: Terminology task. In *Proceedings of WMT*.
- Matt Post, Shuoyang Ding, Marianna Martindale, and Winston Wu. 2019. An exploration of placeholdering in neural machine translation. In *Proceedings of MT Summit*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proceedings of EMNLP*.
- Elsbeth Turcan, David Wan, Faisal Ladhak, Petra Galuscakova, Sukanta Sen, Svetlana Tchistiakova, Weijia Xu, Marine Carpuat, Kenneth Heafield, Douglas Oard, and Kathleen McKeown. 2022. Constrained re-generation for cross-lingual query-focused extractive summarization. In *Proceedings of COLING*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.
- David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2023. Prompting PaLM for translation: Assessing strategies and performance. In *Proceedings of ACL*.
- Biao Zhang, Barry Haddow, and Alexandra Birch. 2023. Prompting large language model for machine translation: A case study. In *Proceedings of ICML*.