

Generating Text from Uniform Meaning Representation

Emma Markle
Amherst College
emarkle26@amherst.edu

Reihaneh Iranmanesh
Amherst College
riranmanesh25@amherst.edu

Shira Wein
Amherst College
swein@amherst.edu


Abstract

Uniform Meaning Representation (UMR) is a recently developed graph-based semantic representation, which expands on Abstract Meaning Representation (AMR) in a number of ways, in particular through the inclusion of document-level information and multilingual flexibility. In order to effectively adopt and leverage UMR for downstream tasks, efforts must be placed toward developing a UMR technological ecosystem. Though only a small amount of UMR annotations have been produced to date, in this work, we investigate the first approaches to producing text from multilingual UMR graphs. Exploiting the structural similarity between UMR and AMR graphs and the wide availability of AMR technologies, we introduce (1) a baseline approach which passes UMR graphs to AMR-to-text generation models, (2) a pipeline conversion of UMR to AMR, then using AMR-to-text generation models, and (3) a fine-tuning approach for both foundation models and AMR-to-text generation models with UMR data. Our best performing models achieve multilingual BERTscores of 0.825 for English and 0.882 for Chinese, a promising indication of the effectiveness of fine-tuning approaches for UMR-to-text generation even with limited UMR data.¹

1 Introduction

Uniform Meaning Representation (UMR) is a graph-based semantic representation designed to capture the core elements of meaning for a wide range of languages, accounting for some languages’ limited resources and linguistic diversity in the annotation process (Van Gysel et al., 2021). An example sentence-level UMR graph can be seen in Figure 1. The UMR project has seen recent progress but is still in the early stages, with new resources such as an online web annotation tool (Ge et al.,

¹Our checkpoints and code are available at <https://github.com/ACNLP1lab/UMR-Text-Gen>



```
(s / search-01
 :Arg0 (p / person
 :refer-person 3rd
 :refer-number Singular)
 :Arg1 (c / clue
 :refer-number Singular)
 :aspect Activity
 :modstr FullAff)
```

Figure 1: UMR graph for the sentence “He was searching for a clue” in graph form (top) and in ‘PENMAN’ notation (Kasper, 1989) (bottom).

2023) and a relatively small recently-released annotated dataset (Bonn et al., 2024).

Abstract Meaning Representation (AMR; Banarescu et al., 2013), on which UMR is based, is an annotation schema designed for English which has seen success and adoption by the broader NLP community, including incorporation into downstream applications and studies of language (Wein and Opitz, 2024); this is in large part due to the substantial efforts made towards high-quality text-to-AMR parsing and AMR-to-text generation models (Sadeddine et al., 2024).² Thus, in order to see similar success for the multilingual UMR schema, efforts towards UMR-to-text generation and text-to-UMR parsing are critical. As such, the multilingual (and document-level) benefits of UMR cannot be

²The breadth of parsing and generation work for AMR can be seen in the AMR Bibliography.

exploited by the NLP community until the tools exist to easily parse and generate.

In this work, we tackle the task of UMR-to-text generation and introduce the first models for this task. Our experimentation focuses on primarily English and Chinese, but leverages data in all of the six languages included in the UMR v1.0 dataset (detailed in Section 3.1): English, Chinese, Sana-paná, Ará-paho, Kukama, and Navajo. We introduce three approaches for this task, all of which account for the limited amount of UMR resources by exploiting the breadth of AMR-to-text generation tools, given the substantial shared features and structure between AMR and UMR. Our approaches and contributions include:

1. A baseline analysis as to how well six AMR-to-text generation models generate text from UMR out-of-the-box (Section 4).
2. A novel pipeline approach to UMR-to-text generation, which converts UMR graphs into AMRs, then uses them as input to AMR-to-text generation models (Section 5).
3. Seven fine-tuned UMR-to-text generation models, using three pretrained large language models and four pretrained AMR-to-text generation models as foundations (Section 6).

Our experiments reveal that while our baseline approach of using UMRs as input to AMR-to-text models does produce text, it often fails to handle UMR-specific terminology. Our fine-tuned models, particularly models pretrained on AMR data, achieve both high adequacy and fluency, and in fact are able to leverage the document-level information contained in UMR graphs for better output despite only generating a single sentence at a time. Additionally, the best performing models are fine-tuned exclusively on the same language UMR data, suggesting that UMR data in other languages may not be helpful in enhancing monolingual UMR-to-text generation models.

2 Uniform Meaning Representation

UMR is based on Abstract Meaning Representation (AMR), which was designed for English (Banasescu et al., 2013) but has since seen various individual language extensions and cross-lingual applications (Wein and Schneider, 2024). While AMR requires these individual language-specific adaptations to accommodate the linguistic features of the individual language (such as pronoun drop,

grammatical number, and affixes), UMR’s flexibility enables annotation of many languages within the same schema, showing its promise as a multilingual representation (Wein and Bonn, 2023).

Similarly to AMR, UMR annotations are rooted, directed graphs that capture meaning (Van Gysel et al., 2021). Unlike AMR, UMR additionally incorporates aspect and modality at the sentence-level, contains document-level information (enabling annotation of coreferential relations), and alignments between the coreferential elements.

In order to ensure that UMR could reflect meaning for many languages, the UMR schema is flexible in its annotation while ensuring consistency across languages, differentiating it from the English-based AMR. At the sentence-level, UMR accounts for linguistic diversity across languages through the use of a lattice-like annotation schema (Van Gysel et al., 2019), enabling annotators to apply a more coarse-grained or fine-grained annotation based on the features of the language being annotated. Particular care is also given to the annotation of low- or no-resource languages (Vigus et al., 2020), as “Stage 0” annotation is performed for languages that do not have existing rolesets available. Stage 0 annotation enables annotators to establish predicate-argument structures and develop rolesets while performing UMR annotation.

While text-to-AMR parsers and AMR-to-text generation models have seen substantial progress in recent years (with Smatch scores of 86% for parsing (Lee et al., 2022; Vasylenko et al., 2023) and 50 BLEU points for generation (Cheng et al., 2022)), no efforts have yet been made towards UMR-to-text generation models, largely due to the recent development of the schema. We opt to use AMR-to-text models as a baseline for UMR-to-text generation, though the fact that UMR contains *more* semantic information than AMR (such as aspect, modality, and document-level information) makes AMR-to-text models not particularly well-suited for this task. Models trained solely on AMR are unlikely to correctly account for these elements, leading to a loss of critical semantic information. The use of AMR-to-text generation models as a baseline motivates progress on dedicated UMR-to-text generation approaches, such as ours.

Recent work has investigated automated Ará-paho UMR annotation (Buchholz et al., 2024), automatic annotation of tense and aspect for UMR (Chen et al., 2021), and conversion of AMR graphs to UMR (Bonn et al., 2023b; Post et al., 2024). Spe-

cial attention has been given to UMR annotation of multi-word expressions (Bonn et al., 2023a) and Chinese verb compounds (Sun et al., 2023).

Contributing to the UMR technological ecosystem, Chun and Xue (2024) released a text-to-UMR parser, which produces the document-level UMR graph based on the contents of the sentence, and the sentence-level UMR graph by running existing text-to-AMR parsers and then converting the AMR into a UMR. This pipeline parsing approach mirrors our baseline approach, in the reverse direction.

3 Methodology

Given that in this work we are developing the first approaches to UMR-to-text generation, here we first establish the data (Section 3.1) and evaluation methods (Sections 3.2 to 3.4) that we use for this task. Then, in the sections that follow, we introduce each of our three different approaches to text generation, which all exploit the wealth of AMR technologies in different ways.

First, our baseline approach uses six AMR-to-text generation models, passing the sentence-level UMR graphs as input (Section 4). Next, using the same six models used for our baseline, we develop a pipeline approach (Section 5) to UMR-to-text generation, which involves first converting the UMR graphs to AMR graphs, and then passing the converted AMR graphs as input to the models. Then, based on the baseline performance of the AMR-to-text generation models, we fine-tune the top performing models as well as three large language models (with no AMR input) on the UMR training set (Section 6).

3.1 Data

We use the first release of UMR data (Bonn et al., 2024), which contains annotations in English, Chinese, and four languages indigenous to the Americas: Aráphaho, Navajo, Sanapaná, and Kukama. The English data contains LORELEI news text and a description of a silent film. The Chinese data consists of wikinews sentences. The Aráphaho, Navajo, and Kukama annotations all represent narrative documents, while it is not clear what genre the Sanapaná annotations are (Bonn et al., 2024). We first take out the English UMR data that contains equivalent AMRs in the AMR3.0 dataset (Knight et al., 2020), as to avoid leakage and unfair evaluations of models which contain the corresponding AMR data, removing 66 English UMRs.

Language	Training (70%)	Dev (10%)	Test (20%)
English	100 (96)	13 (13)	30 (28)
Chinese	236 (236)	40 (40)	82 (82)
Aráphaho	256 (46)	36 (7)	114 (54)
Navajo	371 (148)	52 (20)	83 (0)
Sanapaná	433 (366)	62 (53)	107 (104)
Kukama	76 (76)	10 (10)	19 (0)
Total	1472 (968)	213 (143)	435 (268)

Table 1: Our training, development, and test splits for the UMR data at the sentence-level, with the amount of sentence-level annotations that also had document-level information displayed in parentheses. We only count data as being document-level if it contains both alignment and a document-level graph beyond (s / sentence).

Not all annotations contain both sentence-level and document-level graphs. Henceforth, given that all sentences which have document-level annotations contain both sentence- and document-level graphs, we use *sentence-level annotations* to refer to just the sentence-level graph, and *document-level annotations* to refer to the combination of a sentence-level graph and document-level graph for an individual sentence (note that the document-level graphs contain information about how the sentence in question relates to other sentences in the document). Our final data split is 70% for training, 10% for development, and 20% for testing (see Table 1).

3.2 Automatic Evaluation

We evaluate the generated text via automatic metrics as well as human evaluation.

We compare the generated text from each of our approaches against the reference sentences by using BERTscore (Zhang et al., 2020), given its previously evidenced correlation with human judgments for AMR-to-English text generation (Manning et al., 2020). We specifically use multilingual BERTscore (mBERTscore), as well as METEOR (Banerjee and Lavie, 2005) and BLEU (Papineni et al., 2002) to enable ease of comparison with future work.

3.3 Initial Indigenous Language Evaluation

Our initial experimentation reveals primarily positive quantitative indications as far as the ability of our models to produce text in the four indigenous languages, for example, mBERTscores of 0.799 for Navajo, 0.816 for Sanapaná, 0.780 for Aráphaho, and 0.673 for Kukama, as generated by an

amrlib model fine-tuned on all languages’ sentence-level data. However, given that these languages are extremely low-resource and are not likely to be well-evaluated by automatic metrics, we consult speakers of Aráphaho and Navajo to provide a human evaluation of the generated text. The Aráphaho and Navajo speakers indicate that, while these fine-tuned models do a fairly good job at imitating the script of the four indigenous languages, the output is nonsensical and ungrammatical. We thus opt against a full human evaluation or automatic evaluation for the four indigenous languages, with the understanding that even our top-performing models output gibberish for these languages. This is likely due to their morphological complexity, necessitating additional resources for coherent generation. We still leverage the indigenous language data in fine-tuning our models (Section 6).

3.4 Human Evaluation

In order to validate the quantitative results obtained by the automatic metrics, we perform a human evaluation, here focusing on English and Chinese. Six college students participated in the evaluation of the English and Chinese texts, three native speakers of English and Chinese each. Each annotator judged fluency and adequacy on a scale from 1-4. Fluency was judged first, without exposure to the reference, and then adequacy was judged in relation to the reference.

There were a total of four surveys: English fluency, English adequacy, Chinese fluency, and Chinese adequacy, each of which contained 25 questions. For English, we chose to exclude the 5 shortest sentences from the 30 English test sentences, which left us with a final set of 25 sentences. For Chinese, we randomly selected 25 sentences, again from the test set. Each question in the surveys displayed six sentences, five of which are from our top-performing models (across multiple approaches), as well as the reference. The instructions provided to the human raters for the English fluency and adequacy surveys can be seen in Appendix B.

We evaluated the inter-annotator agreement using the Pearson correlation coefficient (Pearson and Galton, 1895), calculating pairwise agreement for each of the four surveys (Chinese and English, fluency and adequacy for each). The average correlation coefficient for English fluency was 0.72, for English adequacy was 0.78, for Chinese fluency was 0.55, and for Chinese adequacy was 0.64. This indicates that there is a strong correlation between

the annotators’ scores, validating the judgments, while the English raters exhibited even greater agreement.³ Note also that all pair-wise correlation coefficients had a statistically significant p-value ($p < 0.05$).

4 Baseline Approach

Now that we have established our methods for evaluating UMR-to-text generation tools, we are able to start exploring various methods of generation.

First, we leverage the similarity between UMR and AMR graphs and the prevalence of AMR technologies, by using AMR-to-text generation models out-of-the-box. This serves as a baseline model, and we only use sentence-level UMR graphs (as AMR does not contain document-level information), to see how they perform as a zero-shot approach with no exposure to UMR.⁴

4.1 Models

We apply this approach to six AMR-to-text generation models:

1. amrlib⁵: sequence-to-sequence T5 model, trained on AMR3.0 (Knight et al., 2020)
2. AMRBART (Bai et al., 2022): BART-based (Lewis et al., 2020) model pretrained on English text, AMR graphs from AMR2.0 (Knight et al., 2017) and AMR3.0, along with 200k silver AMRs parsed by SPRING
- 3-4. SPRING2 and SPRING3 (Bevilacqua et al., 2021): BART-based sequence-to-sequence model trained on AMR v3.0 that simplifies AMR parsing and generation by treating them as symmetric tasks. SPRING2 and SPRING3 differ in their training datasets: SPRING2 is trained on AMR2.0 while SPRING3 is trained on AMR3.0
5. BiBL (Cheng et al., 2022): utilizes the architectural framework of SPRING to align AMR graphs and text, in order to share information across the parsing and generation tasks trained on AMR2.0 and AMR3.0
6. Smelting (Ribeiro et al., 2021): trained mT5 (Xue et al., 2021) on gold English

³This is likely due to the generated text being both more fluent and more adequate for English than Chinese (see Sections 5.2 and 6.2).

⁴We additionally experiment with a prompt-based LLM baseline, which prompts a GPT4o-mini to generate text from a UMR graph. The prompt-based experiment performs similarly to this baseline, under-performing our other approaches.

⁵amrlib GitHub Repository

AMR graphs and sentences as well as silver machine-translated sentences

AMRBART, amrlib, SPRING2, SPRING3 and BiBL are all trained to generate English text, while Smelting can produce Spanish, Italian, German, and Chinese text. Thus, we run the first five models on the English UMR graphs (to generate into English), and run Smelting on the Chinese UMR graphs (to generate into Chinese).

4.2 Results

The automatic metric results for the baseline models show a promising indication of the utility of AMR-to-text generation tools for UMR-to-text generation.

Out-of-the-box, the English results from the five baseline models (Table 2) all achieve mBERTscores of around 0.7 (ranging from 0.681-0.704). The Chinese baseline results (Table 3) are similarly high, with the baseline Smelting model achieving a mBERTscore of 0.716. The BLEU and METEOR scores are noticeably lower for both languages, to be expected of those metrics.

Our qualitative analysis of the baseline models reveals the common inclusion of UMR-specific terms, such as refer-number singular, full-affirmative, umr-unknown, and 3rd person, in the text output. Examples of this can be seen in the following generated sentence: Full-affirmative, though, the first singular thought that the second singular would not see the apron at first, is a state of ‘full affirmative.’ The only UMR graphs that were converted to highly fluent and adequate text via this baseline approach were one-word sentences such as “ok” or “anyway.” We also noted that the generated output from the UMR graphs tended to be much longer than the reference output, which was likely due to the inclusion of the UMR-specific terms mentioned previously.⁶

Based on the shorter output length and higher perceived fluency in our initial qualitative analysis (in comparison with the other baseline models, including the higher-scoring amrlib baseline), we determined that BiBL was the best baseline English model, and Smelting was our only baseline Chinese model. Thus, we included the baseline

⁶The average sentence length for each model was 7.3 words for the references, 11.7 for amrlib output, 11.8 for BiBL output, 13.6 for AMRBART output, 14.6 for SPRING3 output, and 14.7 for SPRING2, output.

Models	mBERT	BLEU	METEOR
Baseline amrlib	0.704	0.081	0.386
Baseline amrbart	0.691	0.062	0.333
Baseline SPRING3	0.681	0.050	0.298
Baseline SPRING2	0.698	0.068	0.357
<i>Baseline BiBL</i>	0.702	0.038	0.333
Pipeline amrlib	0.761	0.119	0.418
Pipeline amrbart	0.775	0.181	0.409
Pipeline SPRING3	0.764	0.133	0.392
Pipeline SPRING2	0.774	0.193	0.436
<i>Pipeline BiBL</i>	0.784	0.159	0.428
FT amrlib-Sent-EN	0.700	0.041	0.213
FT SPRING2-Sent-EN	0.825	0.355	0.584
<i>FT BiBL-Sent-EN</i>	0.825	0.289	0.570
<i>FT SPRING2-Doc-EN</i>	0.824	0.358	0.582
FT BiBL-Sent-EN&ZH	0.825	0.355	0.601
FT amrlib-Sent-All	0.751	0.057	0.383
FT BiBL-Sent-All	0.800	0.265	0.520
FT SPRING2-Sent-All	0.809	0.281	0.550
FT mBART-Doc-EN	0.750	0.104	0.411
FT mT5-Sent-EN&ZH	0.749	0.140	0.325
FT mT5-Sent-All	0.772	0.181	0.368
FT Gemma-Sent-All	0.440	0.000	0.004
<i>FT mBART-Sent-All</i>	0.767	0.146	0.423
FT mT5-Doc-All	0.774	0.201	0.417
FT mBART-Doc-All	0.740	0.111	0.324

Table 2: English automatic evaluation results for each of our three approaches. The fine-tuned (FT) models include their model name, whether the model was fine-tuned on sentence- or document-level information, and the languages included the fine-tuning data. Italicized entries were selected for the human evaluation. The fine-tuned models above the third horizontal line are AMR-to-text generation models, while the models below the line are foundation models. Due to having trained over 50 individual models, here we show a representative sample of the best performers from each approach.

Models	mBERT	BLEU	METEOR
<i>Baseline Smelting</i>	0.716	0.000	0.029
<i>Pipeline Smelting</i>	0.767	0.000	0.019
FT Smelting-Sent-ZH	0.689	0.000	0.027
<i>FT BiBL-Sent-ZH</i>	0.881	0.247	0.593
FT BiBL-Doc-ZH	0.881	0.247	0.593
<i>FT SPRING2-Doc-ZH</i>	0.882	0.231	0.586
FT BiBL-Sent-EN&ZH	0.879	0.250	0.602
FT BiBL-Sent-All	0.878	0.235	0.592
FT mT5-Doc-ZH	0.815	0.122	0.375
FT mBART-Sent-All	0.837	0.146	0.467
<i>FT mT5-Sent-All</i>	0.853	0.172	0.476

Table 3: Chinese automatic evaluation results for each of our three generation approaches. Again we show a representative sample of the best performers from each approach.

BiBL and Smelting models in the human evaluation survey. The human evaluation scores (Table 4 and Table 5) are low with regard to both fluency and adequacy, in spite of the fairly high mBERTscore values. Baseline BiBL for English has a fluency score of 1.44 (scale of 1-4) and an adequacy score

of 1.37, while baseline Smelting for Chinese receives a fluency score of 1.69 and an adequacy score of 1.48. The English and Chinese reference sentences receive a fluency score of 3.35 and 3.27, respectively, showing that even the ground-truth sentences are not perceived as perfectly fluent by the annotators, perhaps due to the narrative structure of the UMR data. These low scores highlight the inadequacy of the baseline approach, accentuating the need for more advanced techniques to effectively generate text from UMR graphs.

5 Pipeline Approach

Our next approach again leverages existing AMR-to-text generation models, but in this case first converting the UMR graphs into AMR graphs. We use all the same models as described in Section 4.

5.1 Conversion Process

In order to convert UMRs into AMRs, we design a rule-based conversion process. During conversion, some elements of UMR do not appear in AMR (such as aspect and mode) and are thus simply removed. Other changes include converting split roles, renaming roles, and introducing additional roles. This process is informed by Bonn et al. (2023b) as well as the AMR (Banarescu et al., 2019) and UMR guidelines (UMR, 2022). The “person” concepts (:refer-number or refer-person) and pronouns generally are handled in a more complicated way in UMR than AMR, so in order to convert these concepts into their AMR counterparts, we changed all “person” concepts to a single node of the equivalent English pronoun. We convert third-person pronouns, both singular and plural, to “they,” opting for a gender-neutral approach. Additionally, if a graph refers to the first person but did not include a refer-number tag, we default to using the singular pronoun “I.” The Sanapaná data also includes a number of instances of :wiki, which we remove. In total, this process results in 4 kinds of UMR roles being removed and over 65 kinds of roles converted to their equivalent AMR role.⁷

In order to verify the validity of the conversion process itself, we leverage the 66 held-out English UMRs which have equivalent AMRs in the AMR3.0 dataset (detailed in Section 3.1), and compare the generated graphs to their equivalent

⁷We also map :material to :consist-of, :concessive-condition to :condition, and have-91 to have-03.

Models	Fluency	Adequacy	Sum
Reference	3.35	/	3.35
Baseline BiBL	1.44	1.37	2.81
Pipeline BiBL	3.59	2.71	6.30
FT BiBL-Sent-EN	3.37	3.19	6.56
FT SPRING2-Doc-EN	3.44	3.29	6.73
FT mBART-Sent-All	2.77	2.28	5.05

Table 4: English human evaluation results for the most promising models of each approach.

Models	Fluency	Adequacy	Sum
Reference	3.27	/	3.27
Baseline Smelting	1.69	1.48	3.17
Pipeline Smelting	2.19	1.69	3.88
FT BiBL-Sent-ZH	1.76	2.79	4.55
FT SPRING2-Doc-ZH	1.92	2.81	4.73
FT mT5-Sent-All	2.28	2.29	4.57

Table 5: Chinese human evaluation results for the most promising models of each approach.

human-annotated AMR graphs using SMATCH. The average Smatch score for our converter is 0.63, indicating a fairly accurate conversion process. One of the ways in which our rule-based approach to conversion struggles to meet AMR norms is that substantial additional information is contained in the UMR, which is not removed during conversion, such as in the case of the sentence “Pleasure” (see Figure 2).

Original AMR:
(p / pleasure)

Original UMR:
(s29s / say-01
:ARG0 (s29p / person)
:ARG1 (s29h / have-experience-91
:ARG1 s29p
:ARG2 (s29p3 / pleasure)
:ARG3 (s29t / thing)
:aspect state)
:ARG2 (s29p2 / person)
:aspect performance)

Converted AMR:
(s29s / say-01
:ARG0 (s29p / person)
:ARG1 (s29h / have-experience-91
:ARG1 s29p
:ARG2 (s29p3 / pleasure)
:ARG3 (s29t / thing))
:ARG2 (s29p2 / person))

Generated text from BiBL of the converted AMR: “People said it was a pleasurable experience.”

Figure 2: AMR and UMR graphs for the sentence “Pleasure,” compared with the AMR graph produced by our UMR-to-AMR pipeline.

5.2 Results

Our pipeline approach outperforms the baseline approach for both English and Chinese. For English, BiBL achieves the highest mBERTscore, 0.784, and for Chinese, Smelting achieves an mBERTscore of 0.767.

This approach proves to be very effective in reducing the amount of UMR terms that appear in the output, leading to more comprehensible sentences and generally clear output. For example, BiBL’s pipeline conversion produces sentences such as: “The scene is opened by you seeing a tree,” “they dumped all their pears into a basket,” and “they wear apron-like thing with a huge pocket.” These examples of generated text illustrate the pipeline’s strength at producing fluent text; however, as seen in the examples, the pipeline is not able to accurately generate the applicable pronoun and thus always defaults to “they.”

The perceived fluency of this approach is validated by the human evaluation survey. For English, this approach leads to highly fluent output, with a score of 3.59 and an adequacy score of 2.71, a substantial improvement over the baseline approach. The Chinese also has increased scores of 2.19 for fluency and 1.69 for adequacy. It is evident that this approach is more successful for English, but we do observe quantitative and qualitative improvements for both languages.

While the pipeline approach leverages AMR graphs and the well-developed ecosystem of AMR technologies, it is important to note that UMR contains more linguistics information than AMR. Converting from UMR to AMR therefore inevitably leads to information loss, which can affect the accuracy of generated text with respect to UMR-specific content, such as tense, aspect, and referential features.

An illustrative example concerns pronouns. In UMR, pronouns are richly annotated with attributes such as `:ref-person` (1st, 2nd, 3rd) and `:ref-number` (singular, dual, paucal, plural), which capture verbal cross-referencing, overt nominal number, and explicit and implicit arguments, supporting nuanced reference across sentences. In contrast, AMR typically represents pronouns using only surface forms such as “he,” “she,” or “they,” discarding much of the underlying information. As a result, converting UMR graphs to AMR inherently loses information encoded in UMR, reducing alignment with the original graph and obscuring

distinctions in number, person, or implicit arguments.

This mismatch helps explain why our method of first converting UMR to AMR is unlikely to perform optimally: the intermediate AMR representation cannot carry all the distinctions encoded in UMR, leading to systematic gaps in the generated output.

6 Fine-tuning Processes

Our final approach involves fine-tuning models with UMR data. We fine-tune two kinds of models: AMR-to-text generation models, and foundation models, which have not been exposed to any AMR previously. In this approach we incorporate the document-level graphs as well as the sentence-level graphs and fine-tune on different combinations of the language data.

6.1 Methods

We determine which AMR-to-text generation models to fine-tune based on the results from Sections 4 and 5, picking the best-performing models for the baseline and pipeline approaches. Accordingly, the first model we fine-tune is `amrlib`, which also follows Wein et al. (2022) (which fine-tunes the `t5wtense` AMR-to-English generation model for Spanish text generation). Next, we fine-tune `SPRING2` and `Smelting`, given their quantitative performances for English and Chinese, respectively. Finally, we fine-tune `BiBL` due to its high scores, relative brevity, and perceived fluency (as discussed in Section 4.2).

Then, in order to investigate whether the inclusion of AMR data in the fine-tuning process aids model performance, we fine-tune large language models with no prior exposure to AMR (as opposed to fine-tuning the AMR-to-text generation models).⁸ These models are `mT5` (Xue et al., 2021), `mBART` (Liu et al., 2020), and `Gemma 2B` (Gemma Team, 2024).

We then fine-tune each model using the UMR data. In order to determine whether document-level information and UMR data of other languages benefit UMR-to-text generation, we create eight datasets by forming all combinations of sentence or document-level data across English, Chinese, and all languages. This culminates in 8 fine-tuned

⁸Granted, it is possible that these models have been exposed to AMR graphs in pretraining, though prior work has indicated that foundation models are unable to parse AMR graphs out-of-the-box (Ettinger et al., 2023).

versions of each of the 7 models; we highlight the key findings in Table 2.⁹

For reproducibility, details of our hyperparameters follow. We fine-tune BiBL and SPRING, for 30 epochs with a constant learning rate of 0.0001, incorporating gradient clipping at 2.5 and dropout regularization at 0.25. We fine-tune mBART-large-50 and mT5 models for 15 epochs. Similarly, we fine-tune Gemma2-2b using the same optimization parameters.

6.2 Results

While the fine-tuning results show more variance than the baseline and pipeline scores, the fine-tuned models consistently outperform those approaches in terms of automatic metrics, fluency, and adequacy. For English, BiBL and SPRING2 trained on sentence-level English data achieved the highest mBERTSCORE of 0.825. As for Chinese, SPRING2 trained on Chinese document-level data achieved the highest mBERTSCORE of 0.882. Furthermore, these high automatic scores are corroborated by human evaluation results, which show high levels of fluency and adequacy scores for both languages. For English, SPRING2 trained on English document-level data, received the second-highest English fluency rating along with the highest adequacy rating. For Chinese, the best fluency score came from mT5 trained on all languages’ sentence-level data, while SPRING2 led in adequacy when trained on Chinese document-level data.

These results point to several broad patterns: (1) the benefit of document-level data, (2) the relevance of AMR for UMR technologies, and (3) the strength of our fine-tuning approaches even with limited UMR annotations. First, the inclusion of document-level data appears beneficial, especially for SPRING2. This suggests that document-level information helps in the task of generation, even when only producing individual sentences.

Second, our fine-tuned AMR-to-text models outperform the fine-tuned foundation models such as mT5 and mBART, indicating that models pre-trained on AMR carry over information that is beneficial to the generation of UMR graphs, despite UMR’s broader scope of information.

Third, the AMR-to-text fine-tuned models generally perform better at a specific language when they are only trained on that language’s data versus

being trained on all languages. For example, English generation quality drops slightly when trained on multilingual corpora compared to English-only data. This may be due to the “curse of multilinguality:” exposure to additional languages can degrade monolingual abilities (Conneau et al., 2020). Chinese models in particular exhibit more variability and, on average, perform slightly below their English counterparts across automatic and human metrics. This may reflect lower-quality annotations or differences in UMR complexity.

Through qualitative analysis, we discover that different models had specific tendencies during generation. For example, BiBL reproduces ellipses and pauses that can be found in the training data due to the data being narrative annotations. This suggests that BiBL closely mirrors the training input structure, even when it is potentially reducing fluency. Furthermore, SPRING2 produces smoother and more fluent output, while still reflecting the aforementioned narrative cues. While fine-tuned LLMs such as mT5 and mBART showed improvements after fine-tuning, they were consistently outperformed by AMR-to-text generation models. This reveals that the inclusion of AMR data in the training process improves model performance as opposed to solely UMR data.

To illustrate these trends, below are examples of output from mBART and SPRING2 from the same graphs. Our mBART model fine-tuned on sentence-level data of all languages’ UMR data produces “A–nd.. one of the pears.. drops down the floor,” and “and he’s dumping in the basket with all pears.” On the other hand, SPRING2 trained on document-level English UMR data produces the much more fluent “One of the pears.. drops down to the floor,” and “and he dumps all his pears into a basket.” Compared to mBART, SPRING2 produces more structured, grammatically correct, and fluent output, likely due to previous training on AMR data.

In a linguistic error analysis, we find that model failures fall into a few consistent observable categories. First, pronoun consistency degrades across models, with BiBL and mBART incorrectly switching between pronouns when referring to the same entity. Furthermore, grammatical and lexical errors appear in otherwise fluent outputs, most visibly in mBART (such as “one of the pears plucks into the tree”). Lastly, models often changed how events were originally framed. For example, in a sentence where the speaker expresses some uncertainty (like

⁹Recall that the document-level annotations subsume individual sentence-level annotations as well.

“I don’t think you see the apron at first”), our models frequently reframe these as confident statements. This indicates that while models can extract the general story from the UMR graphs, they struggle to preserve the speaker’s tone and intent.

Overall, our fine-tuned AMR-to-text models outperform all other approaches in this work, with the most successful models trained solely on the target language with exposure to both sentence-level and document-level data.

7 Conclusion

In this work, we develop tools to generate text from Uniform Meaning Representation graphs, introducing three distinct approaches that make use of existing Abstract Meaning Representation technologies to varying degrees.

Our results indicate that the fine-tuned AMR-to-text generation models are stronger for both English and Chinese. Further, the best performing models are fine-tuned exclusively on the same language UMR data, suggesting that UMR data in other languages may not be helpful in enhancing monolingual UMR-to-text generation models. We also find that, while we only generate a single sentence at a time, the document-level information improves model performance. Finally, we find that leveraging AMR data also benefits UMR-to-text generation in a fine-tuned setting, as the AMR-to-text generation models perform better text generation than the foundation models when both are fine-tuned on UMR data.

Future work may explore advancements toward producing text in other languages from UMR, perhaps leveraging additional external resources. The annotation and release of additional UMR data will support this effort, as well as lead to improved performance for English and Chinese generation.

Limitations

One of our limitations is the small size of the UMRv1.0 dataset, which constrains our model fine-tuning capabilities—especially for low-resource languages where training data is scarce. Additionally, we lack robust automatic evaluation metrics for the four indigenous languages, with initial reviews by native speakers indicating the generated outputs do not carry meaning. Due to inherent randomness in LLM generation, the LLMs used in this work exhibit slight variability across different runs, which may result in occasionally producing varied

responses and thus different BLEU, METEOR, and mBERT scores.

Our approaches use sequence-to-sequence architectures with linearized graphs to leverage existing AMR technologies given limited UMR data. Future work might explore additional architectures such as graph neural networks which directly encode graph structure, or other forms of encoding the UMR without linearization.

Acknowledgments

Thank you to the adjudicators for our human evaluation: Crawford Dawson, Yilin Huang, Hank Hsu, Ryan Ji, Megan Li, and Yichen Liu. Thank you also to Jiyuan Ji and Javier Gutierrez Bach for comments and experimental feedback, and to Lukas Denk and Andrew Cowell for providing judgments on some of the Navajo and Arápho sentences, respectively. This work is supported by the Amherst College HPC, which is funded by NSF Award 2117377.

References

- 2022. Uniform Meaning Representation (UMR) 0.9 specification. <https://github.com/umr4nlp/umr-guidelines/blob/master/guidelines.md>.
- Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022. Graph pre-training for AMR parsing and generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6001–6015, Dublin, Ireland. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Dis-course*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2019. Abstract Meaning Representation (AMR) 1.2.6 specification. <https://github.com/amrisi/amr-guidelines/blob/master/amr.md>.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

- Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021. One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline. In *Proceedings of AAAI*.
- Julia Bonn, Matthew J. Buchholz, Jayeol Chun, Andrew Cowell, William Croft, Lukas Denk, Sijia Ge, Jan Hajič, Kenneth Lai, James H. Martin, Skatje Myers, Alexis Palmer, Martha Palmer, Claire Benet Post, James Pustejovsky, Kristine Stenzel, Haibo Sun, Zdeňka Urešová, Rosa Vallejos, Jens E. L. Van Gysel, Meagan Vigus, Nianwen Xue, and Jin Zhao. 2024. [Building a broad infrastructure for uniform meaning representations](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2537–2547, Torino, Italy. ELRA and ICCL.
- Julia Bonn, Andrew Cowell, Jan Hajič, Alexis Palmer, Martha Palmer, James Pustejovsky, Haibo Sun, Zdenka Uresova, Shira Wein, Nianwen Xue, and Jin Zhao. 2023a. [UMR annotation of multiword expressions](#). In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 99–109, Nancy, France. Association for Computational Linguistics.
- Julia Bonn, Skatje Myers, Jens E. L. Van Gysel, Lukas Denk, Meagan Vigus, Jin Zhao, Andrew Cowell, William Croft, Jan Hajič, James H. Martin, Alexis Palmer, Martha Palmer, James Pustejovsky, Zdenka Urešová, Rosa Vallejos, and Nianwen Xue. 2023b. [Mapping AMR to UMR: Resources for adapting existing corpora for cross-lingual compatibility](#). In *Proceedings of the 21st International Workshop on Treebanks and Linguistic Theories (TLT, GURT/SyntaxFest 2023)*, pages 74–95, Washington, D.C. Association for Computational Linguistics.
- Matthew J. Buchholz, Julia Bonn, Claire Benet Post, Andrew Cowell, and Alexis Palmer. 2024. [Bootstrapping UMR annotations for Arapaho from language documentation resources](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2447–2457, Torino, Italy. ELRA and ICCL.
- Daniel Chen, Martha Palmer, and Meagan Vigus. 2021. [AutoAspect: Automatic annotation of tense and aspect for uniform meaning representations](#). In *Proceedings of the Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 36–45, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ziming Cheng, Zuchao Li, and Hai Zhao. 2022. [BiBL: AMR parsing and generation with bidirectional Bayesian learning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5461–5475, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Jayeol Chun and Nianwen Xue. 2024. [Uniform meaning representation parsing as a pipelined approach](#). In *Proceedings of TextGraphs-17: Graph-based Methods for Natural Language Processing*, pages 40–52, Bangkok, Thailand. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Allyson Ettinger, Jena Hwang, Valentina Pyatkin, Chandra Bhagavatula, and Yejin Choi. 2023. [“you are an expert linguistic annotator”: Limits of LLMs as analyzers of Abstract Meaning Representation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8250–8263, Singapore. Association for Computational Linguistics.
- Sijia Ge, Jin Zhao, Kristin Wright-bettner, Skatje Myers, Nianwen Xue, and Martha Palmer. 2023. [UMR-writer 2.0: Incorporating a new keyboard interface and workflow into UMR-writer](#). In *Proceedings of the 17th Linguistic Annotation Workshop (LAW-XVII)*, pages 211–219, Toronto, Canada. Association for Computational Linguistics.
- Gemma Team. 2024. [Gemma 2: Improving open language models at a practical size](#). *arXiv preprint arXiv:2408.00118*.
- Robert T. Kasper. 1989. [A flexible interface for linking applications to Penman’s sentence generator](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Philadelphia, Pennsylvania, February 21-23, 1989*.
- Kevin Knight, Bianca Badarau, Laura Banarescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O’Gorman, and Nathan Schneider. 2017. [Abstract Meaning Representation \(AMR\) Annotation Release 2.0](#). Technical Report LDC2017T10, Linguistic Data Consortium, Philadelphia, PA.
- Kevin Knight, Bianca Badarau, Laura Baranescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O’Gorman, et al. 2020. [Abstract Meaning Representation \(AMR\) Annotation Release 3.0](#). Technical Report LDC2020T02, Linguistic Data Consortium, Philadelphia, PA.
- Young-Suk Lee, Ramón Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian, and Salim Roukos. 2022. [Maximum Bayes Smatch ensemble distillation for AMR parsing](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human*

- Language Technologies*, pages 5379–5392, Seattle, United States. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *arXiv preprint arXiv:2001.08210*.
- Emma Manning, Shira Wein, and Nathan Schneider. 2020. [A human evaluation of AMR-to-English generation systems](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4773–4786, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Karl Pearson and Francis Galton. 1895. [VII. note on regression and inheritance in the case of two parents](#). *Proceedings of the Royal Society of London*, 58(347-352):240–242.
- Claire Benet Post, Marie C. McGregor, Maria Leonor Pacheco, and Alexis Palmer. 2024. [Accelerating UMR adoption: Neuro-symbolic conversion from AMR-to-UMR with low supervision](#). In *Proceedings of the Fifth International Workshop on Designing Meaning Representations @ LREC-COLING 2024*, pages 140–150, Torino, Italia. ELRA and ICCL.
- Leonardo F. R. Ribeiro, Jonas Pfeiffer, Yue Zhang, and Iryna Gurevych. 2021. [Smelting gold and silver for improved multilingual AMR-to-Text generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 742–750, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zacchary Sadeddine, Juri Opitz, and Fabian Suchanek. 2024. [A survey of meaning representations – from theory to practical utility](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2877–2892, Mexico City, Mexico. Association for Computational Linguistics.
- Haibo Sun, Yifan Zhu, Jin Zhao, and Nianwen Xue. 2023. [UMR annotation of Chinese verb compounds and related constructions](#). In *Proceedings of the First International Workshop on Construction Grammars and NLP (CxGs+NLP, GURT/SyntaxFest 2023)*, pages 75–84, Washington, D.C. Association for Computational Linguistics.
- Jens E. L. Van Gysel, Meagan Vigus, Pavlina Kalm, Sook-kyung Lee, Michael Regan, and William Croft. 2019. [Cross-linguistic semantic annotation: Reconciling the language-specific and the universal](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 1–14, Florence, Italy. Association for Computational Linguistics.
- Jens EL Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O’Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, et al. 2021. Designing a uniform meaning representation for natural language processing. *KI-Künstliche Intelligenz*, 35(3):343–360.
- Pavlo Vasylenko, Pere Lluís Huguet Cabot, Abelardo Carlos Martínez Lorenzo, and Roberto Navigli. 2023. [Incorporating graph information in transformer-based AMR parsing](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1995–2011, Toronto, Canada. Association for Computational Linguistics.
- Meagan Vigus, Jens E. L. Van Gysel, Tim O’Gorman, Andrew Cowell, Rosa Vallejos, and William Croft. 2020. [Cross-lingual annotation: a road map for low- and no-resource languages](#). In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 30–40, Barcelona Spain (online). Association for Computational Linguistics.
- Shira Wein and Julia Bonn. 2023. [Comparing UMR and cross-lingual adaptations of AMR](#). In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 23–33, Nancy, France. Association for Computational Linguistics.
- Shira Wein, Lucia Donatelli, Ethan Ricker, Calvin Engstrom, Alex Nelson, Leonie Harter, and Nathan Schneider. 2022. [Spanish Abstract Meaning Representation: Annotation of a general corpus](#). In *Northern European Journal of Language Technology, Volume 8*, Copenhagen, Denmark. Northern European Association of Language Technology.
- Shira Wein and Juri Opitz. 2024. [A survey of AMR applications](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6856–6875, Miami, Florida, USA. Association for Computational Linguistics.
- Shira Wein and Nathan Schneider. 2024. [Assessing the Cross-linguistic Utility of Abstract Meaning Representation](#). *Computational Linguistics*, pages 1–55.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and

Colin Raffel. 2021. *mt5: A massively multilingual pre-trained text-to-text transformer*. *arXiv preprint arXiv:2010.11934*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. *BERTScore: Evaluating text generation with BERT*. In *Proc. of ICLR*, Online.

A Full Fine-tuning Automatic Metric Results

Models	mBERT	BLEU	METEOR
<i>amrlib</i>	0.751	0.057	0.383
<i>SPRING2</i>	0.809	0.281	0.550
<i>BiBL</i>	0.800	0.265	0.520
<i>Smelting</i>	0.589	0.000	0.070
<i>mT5</i>	0.772	0.181	0.368
<i>Gemma2</i>	0.440	0.000	0.004
<i>mBART</i>	0.767	0.146	0.423

Table 6: **English** automatic evaluation results for the fine-tuning approach *trained on all languages*.

Models	mBERT	BLEU	METEOR
<i>amrlib</i>	0.780	0.144	0.400
<i>SPRING2</i>	0.821	0.312	0.596
<i>BiBL</i>	0.825	0.355	0.601
<i>Smelting</i>	0.612	0.000	0.072
<i>mT5</i>	0.749	0.140	0.325
<i>Gemma2</i>	0.506	0.001	0.031
<i>mBART</i>	0.744	0.079	0.341

Table 7: **English** automatic evaluation results for the fine-tuning approach *trained on English and Chinese*.

Models	mBERT	BLEU	METEOR
<i>amrlib</i>	0.700	0.041	0.213
<i>SPRING2</i>	0.825	0.355	0.584
<i>BiBL</i>	0.825	0.289	0.570
<i>Smelting</i>	0.565	0.000	0.055
<i>mT5</i>	0.570	0.000	0.005
<i>Gemma2</i>	0.506	0.001	0.031
<i>mBART</i>	0.744	0.111	0.356

Table 8: **English** automatic evaluation results for the fine-tuning approach *trained on English*.

Models	mBERT	BLEU	METEOR
<i>amrlib</i>	0.033	0.000	0.001
<i>SPRING2</i>	0.870	0.188	0.553
<i>BiBL</i>	0.878	0.235	0.592
<i>Smelting</i>	0.706	0.000	0.024
<i>mT5</i>	0.853	0.172	0.476
<i>Gemma2</i>	0.449	0.000	0.002
<i>mBART</i>	0.837	0.146	0.467

Table 9: **Chinese** automatic evaluation results for the fine-tuning approach *trained on all languages*.

Models	mBERT	BLEU	METEOR
<i>amrlib</i>	0.027	0.000	0.002
<i>SPRING2</i>	0.882	0.246	0.587
<i>BiBL</i>	0.879	0.250	0.602
<i>Smelting</i>	0.685	0.000	0.028
<i>mT5</i>	0.826	0.147	0.421
<i>Gemma2</i>	0.510	0.001	0.055
<i>mBART</i>	0.834	0.150	0.469

Table 10: **Chinese** automatic evaluation results for the fine-tuning approach *trained on English and Chinese*.

Models	mBERT	BLEU	METEOR
<i>amrlib</i>	0.039	0.000	0.003
<i>SPRING2</i>	0.880	0.229	0.572
<i>BiBL</i>	0.881	0.247	0.593
<i>Smelting</i>	0.689	0.000	0.027
<i>mT5</i>	0.826	0.131	0.404
<i>Gemma2</i>	0.520	0.002	0.064
<i>mBART</i>	0.837	0.148	0.463

Table 11: **Chinese** automatic evaluation results for the fine-tuning approach *trained on Chinese*.

B Instructions for Human Evaluation

How **fluent** is this text as an example of English? Is it well-formed grammatically with correct spelling and punctuation? Are the terms appropriately used according to common convention? Is the text generally interpretable by a native speaker of English? **Note:** A number of these sentences are transcripts of speech, and thus pauses are indicated in a number of ways such as ellipsis or dashes neither of which should be counted against fluency.

For all of the items that follow, select one of these four levels of fluency:

1. **Nonsense:** Not understandable.
2. **Poor:** Many or serious mistakes which make the text hard to understand.
3. **Good:** Few or minor mistakes. The text is mostly understandable.
4. **Flawless:** Perfectly formed English with no mistakes.

Figure 3: English fluency instructions

How much of the **meaning** from the reference text (text located at the top of each page) is included in the text options?

Note: Grammatical or spelling mistakes should not be considered here. This is not a question of fluency.

Note: A number of these sentences are transcripts of speech, and thus pauses are indicated in a number of ways such as ellipsis or dashes neither of which should be counted against fluency.

For all of the items that follow, select one of these four levels of adequacy / meaning preservation:

1. **None:** The text is completely unrelated to the reference.
2. **Little:** Some of the meaning is preserved, but much of the meaning has been lost or much additional meaning has been added.
3. **Most:** Most of the meaning from the reference is preserved, with a little information missing or added in the text.
4. **All:** All of the meaning is conveyed.

Figure 4: English adequacy instructions

C Hyperparameters

Hyperparameter	BiBL/SPRING2	Smelting	Gemma2
Learning Rate	1×10^{-4}	1×10^{-4}	1×10^{-4}
Epochs	30	15	15
Batch Size (per device)	8	8	8
Gradient Accumulation	16	2	2
Optimizer	–	AdaFactor	AdaFactor
Max Source Length	1024	1024	1024
Max Target Length	1024	1024	2048
Beam Size	1	5	5
Warmup Steps	100	–	100
Weight Decay	0.004	–	0.004
Mixed Precision	–	–	FP16

Table 12: Training hyperparameters for fine-tuning BiBL, SPRING2 (identical to BiBL), Smelting and Gemma2 models.