



HAL
open science

Llamipa: An Incremental Discourse Parser

Kate Thompson, Akshay Chaturvedi, Julie Hunter, Nicholas Asher

► **To cite this version:**

Kate Thompson, Akshay Chaturvedi, Julie Hunter, Nicholas Asher. Llamipa: An Incremental Discourse Parser. EMNLP 2024, ACL, Nov 2024, Miami (FL), United States. pp.6418-6430. hal-04829235

HAL Id: hal-04829235

<https://hal.science/hal-04829235v1>

Submitted on 10 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC0 - Public Domain Dedication 4.0 International License

Llamipa: An Incremental Discourse Parser

Kate Thompson*[†], Akshay Chaturvedi[†], Julie Hunter*, Nicholas Asher^{†‡}

[†]IRIT, [‡]CNRS, *LINAGORA Labs
Toulouse, France

Abstract

This paper provides the first discourse parsing experiments with a large language model (LLM) finetuned on corpora annotated in the style of SDRT (Segmented Discourse Representation Theory [Asher, 1993](#); [Asher and Lascarides, 2003](#)). The result is a discourse parser, Llamipa (Llama Incremental Parser), that leverages discourse context, leading to substantial performance gains over approaches that use encoder-only models to provide local, context-sensitive representations of discourse units. Furthermore, it can process discourse data incrementally, which is essential for the eventual use of discourse information in downstream tasks.

1 Introduction

Understanding a conversation requires more than just grasping the contents of its utterances. One must also understand how different utterances *relate* to one another, for example whether an utterance answers a question or serves to explain or correct something previously said. Discourse parsing, the task of predicting such relations, often represents the discourse structure of texts or conversations as graphs in which *elementary discourse units* or EDUs (roughly, clause-level contents) are represented by nodes, and semantic relations between clauses are represented as labeled edges.

Discourse parsing remains a challenge in NLP due to three factors: (i) the complexity of discourse graphs, (ii) the frequent lack of surface cues provided by EDUs, which forces parsers to rely on deep, semantic information, (iii) the context dependent nature of the inferences required to predict the graph elements.

Recently, improved discourse parsers have incorporated neural models and encoder-only pretrained language models like BERT ([Devlin et al., 2019](#); [Bennis et al., 2023](#)) which offer

richer representations of linguistic content, and can address factor (ii). But given their limited input size, they are confined to look at one small part of a discourse at a time, and so cannot really address factor (iii).

Generative large language models (LLMs) can substantially improve discourse parsing. They not only encode large amounts of lexical and conceptual information, addressing problem (ii), but their increased context windows permit an incremental parsing strategy to address problem (iii). Such a strategy makes sense cognitively and linguistically, as prior discourse structure has an effect on inferences about where subsequent discourse units are attached (the link task), and how they are attached, i.e., which types of semantic relations link them (the link+relation task) ([Asher and Lascarides, 2003](#)). Further, an incremental approach naturally applies to downstream applications such as conversational agents, in which the information about the previously inferred discourse structure must be available and updated in real time.

We develop an incremental parsing model by fine-tuning Llama3 ([Dubey et al., 2024](#)). Llamipa (Llama Incremental Parser) uses previously inferred discourse structure to simultaneously predict the links and relation types between new discourse units. We show that Llamipa substantially outperforms the state-of-the-art on both the link and link+relation tasks.

We demonstrate Llamipa’s performance using three datasets annotated in the style of SDRT (Segmented Discourse Representation Theory [Asher, 1993](#); [Asher and Lascarides, 2003](#)). Our central dataset is the Minecraft Structured Dialogue Corpus (MSDC; [Thompson et al., 2024](#)), a corpus of situated, chat-based, task-oriented dialogues, but we also show the generality of our approach by separately fine-tuning on STAC, a corpus of situated multiparty chats from an online game

(Asher et al., 2016). We further test Llamipa fine-tuned on STAC on an out of domain dataset, Molweni (Li et al., 2020).

Not only does Llamipa outperform previous discourse parsing models, but it handles particularly complex structural features of our situated, multiparty datasets, addressing problem (i). In particular, Llamipa can identify cases in which a single EDU has two parent EDUs, a situation that arises naturally in both multiparty and situated dialogue, where a speaker can respond to multiple channels (conversational or multimodal) at once. Most current SDRT based parsers cannot handle such structures (Shi and Huang, 2019; Liu and Chen, 2021; Wang et al., 2021; Chi and Rudnicky, 2022).

In what follows, Section 2 provides background on discourse structure. Section 3 presents related work, while Section 4 details the multiparty corpora we use for training and testing. Sections 5 and 6 describe our model and results. Section 7 describes our ablation experiments, and we conclude in Section 8.

2 Background

Discourse representations for a text or conversation often take the form of tree or graph structures that are built recursively from discourse units or EDUs, where EDUs are clauses or subclausal units that serve as minimal, linguistic constituents in *discourse relations* (Marcu, 1999). These relations, such as *Explanation*, *Question-Answer Pair* or *Acknowledgement*, are represented as linking EDUs in a discourse graph.

There are two main theories that have investigated complete discourse structures for texts: Rhetorical Structure Theory (RST; Mann and Thompson, 1987) and SDRT (Asher, 1993; Asher and Lascarides, 2003). RST assigns relations between contiguous EDUs which are connected hierarchically to further EDUs in a treelike structure.

SDRT provides a dependency style, weakly connected directed acyclic graph (DAG) in which typed relations are represented as labeled arcs in the graph, and the terms of the relations are EDUs or *elementary event units* (EEUs). EEUs represent salient aspects of the nonlinguistic context in which a conversation takes place, and may be connected to the discourse DAG. SDRT also allows subgraphs, or *complex discourse units* (CDUs) to act as terms

of semantic relations. SDRT is the only framework to date that handles multiparty conversation.

Whatever the underlying paradigm, discourse parsing requires addressing three tasks. First, the segmentation task, which involves decomposing a text and (descriptions of) actions into EDUs and EEUs. The link task involves identifying pairs of EDUs and/or EEUs that are semantically related via discourse relations—that is, the arcs in the discourse graph. Finally, the link+relation task involves identifying the labels for each attached pair.

3 Related work

Traditionally, SDRT style parsers assume gold discourse unit segmentation (Muller et al., 2012; Afantenos et al., 2015; Perret et al., 2016; Shi and Huang, 2019; Liu and Chen, 2021; Chi and Rudnicky, 2022), and parsing follows a pipeline architecture in which discourse units are first linked, and then labels are predicted for those links. Our proposed architecture assumes gold discourse unit segmentation, but combines the attachment and labeling tasks.

Several strategies have evolved to harness the complex information required for discourse parsing within the dependency based paradigm proposed by SDRT or DLTAg (Forbes et al., 2003). Research has shown that semantic features of EDUs are crucial for solving the link+relation task (Mann and Thompson, 1987; Asher and Lascarides, 2003). Thus, older approaches to multiparty conversation (Afantenos et al., 2015; Perret et al., 2016) input manually designed features of the EDU pairs to simple ML models to predict link and link+relation label successively. These *local models* use information local to the pairs of EDUs that are to be attached, but do not exploit information about the surrounding discourse structure. In order to leverage some more global information, Afantenos et al. (2015); Perret et al. (2016) apply special decoding mechanisms with post hoc constraints on the output EDU pairs, following Muller et al. (2012). Although local models exploit little to no information about other components of the discourse structure, such a model can still be competitive (Bennis et al., 2023). We do a detailed comparison between an encoder based local parsing model and Llamipa.

Shi and Huang (2019) were the first to obtain competitive discourse parsing results using a deep

learning architecture on STAC. They attempted to capture incremental and contextual effects in their model by training a supplemental *Structured Encoder* to incrementally update attachment paths (sequences of parent-child EDUs). However, Wang et al. (2021) show that the model obtains similar scores with or without this encoder.

Liu and Chen (2021) develop an approach that uses a pre-trained RoBERTa model (Liu et al., 2019) to provide embeddings for EDUs. In addition, they use a bi-GRU cell (Cho et al., 2014) to capture contextual information; but this requires them to limit the size of the input. Their model uses two linear layers for link and relation prediction.

Wang et al. (2021) provide the Structure Self-Aware Graph Neural Network (SSA-GNN), which is a complex GNN-based architecture and model that uses both EDU and relation embeddings. The model uses a Hierarchical GRU gate for contextual EDU representations. They then apply the SSA-GNN to capture implicit structural information between EDUs, using a Structure-Aware Scaled Dot-Product Attention (Zhu et al., 2019; Wang et al., 2020) to update edge and EDU representations. They further train a teacher network to supplement the standard classification loss with an auxiliary loss. Chi and Rudnicky (2022) developed a more complex model based on Shi and Huang (2019).

None of the models by Shi and Huang (2019); Liu and Chen (2021); Guz et al. (2020); Wang et al. (2021); Chi and Rudnicky (2022) predict multiple parents for an EDU (or EEU). The only previously developed parser that predicts multiple parents during attachment is Bennis et al. (2023)’s model BERTLine. BERTLine uses BERT embeddings for EDUs and EEUs, along with a linear layer for features like speaker change and distance information, to build a local model and predicts link and link+relation in pipeline fashion. However the relation labeling uses a multitask architecture to capture informational dependencies between link and relation labeling decisions. Bennis et al. (2023) show that their model is competitive with approaches that try to build in more contextual knowledge (Liu and Chen, 2021; Wang et al., 2021). Thompson et al. (2024) uses BERTLine on the MSDC, which makes it a natural baseline for our work on situated, multiparty conversation.

Recent work investigates using LLMs to predict discourse structure by finetuning Llama2 7b, 13b,

and 70b for RST-style parsing (Maekawa et al., 2024). The reported results show improvements over the state-of-the-art for RST corpora analogous to those we find using Llamipa on SDRT corpora. Chan et al. (2024); Yung et al. (2024); Fan et al. (2024) have investigated discourse parsing through prompting with large GPT-style models to produce PDTB (Prasad et al., 2008) or other discourse annotations. These models fare worse than dedicated discourse parsers: Chan et al. (2024) test with 3-shot learning on a version of the STAC corpus (see Section 4 for more corpora details) and report a near 50 point drop in F1 for link and link+relation prediction compared to Shi and Huang (2019).

4 Data sets

MSDC. The MSDC provides full discourse structures in the style of SDRT for dialogues in the Minecraft Dialogue Corpus (MDC; Narayan-Chen et al., 2019).

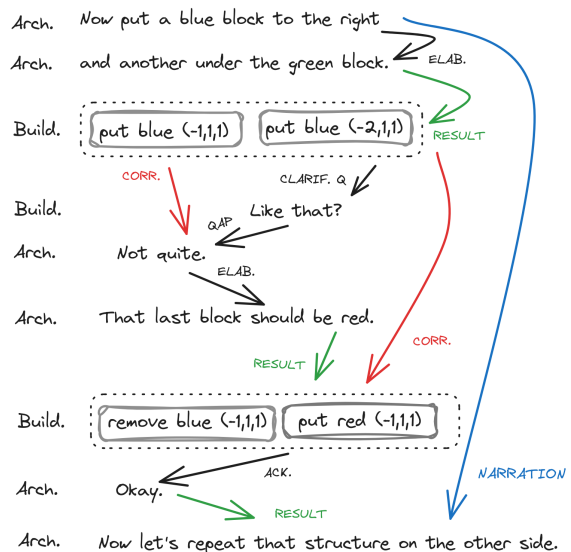


Figure 1: MSDC dialogue example showing CDUs composed of EEUs for builder action sequences and multiparent discourse units (MPDUs). The Narration on the right links “high level” instructions, which are particularly hard to predict, as explained in Section 7.

The MSDC represents conversational moves as well as builder move sequences, individuated into block placements and removals. Figure 1 shows an excerpt from an MSDC dialogue containing a few examples of discourse units with multiple parents, or MPDUs. In one such example, after placing two blue blocks, the builder asks, “like that?”. The

DU-type	STAC-L Train/Test	STAC-Sit Train/Test	MSDC Train/Test
EDU	11398/1154	11707/1163	17135/5402
EEU	0/0	11566/1156	4687/1473
MPDU	838//86	999/96	4789/1476
MPDU=3	94/9	95/6	85/35
MPDU>3	11/0	7/0	1/3

Table 1: EDU, EEU, and MPDU counts for SDRT datasets. MPDU=3 gives the number of MPDUs with 3 parents, MPDU>3 those with more than 3 parents.

architect’s reply, “not quite,” not only serves to answer the builder’s question, but also corrects the builder’s last moves. The EDU for “like that?” is an MPDU that is semantically linked to both linguistic and nonlinguistic units. Such MPDUs are an important feature of the MSDC, as they facilitate the interplay of multimodal contributions in a single discourse structure.

Another noteworthy feature of the MSDC is how nonlinguistic actions are annotated. Builder moves are usually prompted by architect instructions, but often transpire in longer sequences of block placements and removals, which sometimes overlap with other linguistic chat moves. These action sequences are grouped together in single complex discourse units (CDUs) in the MSDC, indicated by the dashed boxes in Figure 1. Bennis et al. (2023), Thompson et al. (2024) merged these highly repetitive sequences of nonlinguistic moves into a single EEU to prevent the repetitive connections between them from drowning out the more complex and interesting linguistic interactions in the corpus. We used the MSDC with this modification.

STAC. STAC¹ is a corpus of multiparty chats from the game *Settlers of Catan*, in which players trade or otherwise acquire resources in order to build roads and settlements and thereby score victory points. STAC has two versions: *STAC-L*, in which only relations between chat moves (EDUs) are annotated in the absence of the context provided by nonlinguistic events, and a situated version, *STAC-Sit*, in which both chat moves and (descriptions of) game moves (EEUs) are annotated in the presence of an integrated linguistic and nonlinguistic context.

Like MSDC, STAC contains MPDUs due to multimodal interactions, but also features MPDUs which arise from multiparty interaction, in which

one player can simultaneously respond to multiple players at once. Relevant statistics for STAC and the MSDC are found in Table 1.

STAC-Sit, like the MSDC, also contains sequences of moves containing only EEUs. Bennis et al. (2023) compressed these sequences into individual EEUs and we follow this choice. Unlike the MSDC, however, STAC contains numerous CDUs for chat moves and for chat/game interactions (as opposed to CDUs of action sequences only). Following Bennis et al. (2023) and most of the literature on SDRT parsing, we do not treat CDUs directly, but use the “flattening” procedure outlined in Muller et al. (2012) that converts a structure with CDUs into one with only EDUs and EEUs. A final simplification of the original STAC corpus involved removing EEUs that neither interact with EDUs nor figure in EEU sequences that interact with EDUs. This choice, like the one to compress EEU sequences, is made to avoid drowning out the more complex and interesting data involving linguistic EDUs and their interactions with EEUs.

Molweni. The Molweni corpus contains SDRT graphs for multiparty interactions in the Ubuntu Chat Corpus (Lowe et al., 2015), a collection of exchanges from an online forum for discussing issues related to Ubuntu. Unlike the MSDC and STAC, Molweni does not contain CDUs, EEUs, or MPDUs. We use the Molweni-clean test set, a version of Molweni re-annotated by Li et al. (2024).

5 Llamipa: An Incremental Parser

Llamipa leverages the enhanced contextual reasoning capacities of generative LLMs by introducing an *incremental discourse parsing strategy*. As it generates a discourse graph for a conversation, the LLM bases its link and relation predictions for an elementary unit ϵ based on the graph it has already constructed for units preceding ϵ .

Formally, let C be a sequence $\epsilon_0, \dots, \epsilon_n$ where ϵ_i denotes an EDU or EEU. Our LLM takes as input a set U , containing EDUs and EEUs, and a discourse graph G . At step 0, U is $\{\epsilon_0\}$, the head EDU, and G , the null discourse graph. After that, U is a set $\{\epsilon_l, \dots, \epsilon_m\}$, where $0 \leq l < m \leq n$ and $m - l \leq k$ where k is the window size, and G contains links connecting $\{\epsilon_l, \dots, \epsilon_m\}$. At each step, the model extends the previous discourse graph by integrating link and relation predictions for all of the EDUs or

¹<https://www.irit.fr/STAC/corpus.html>

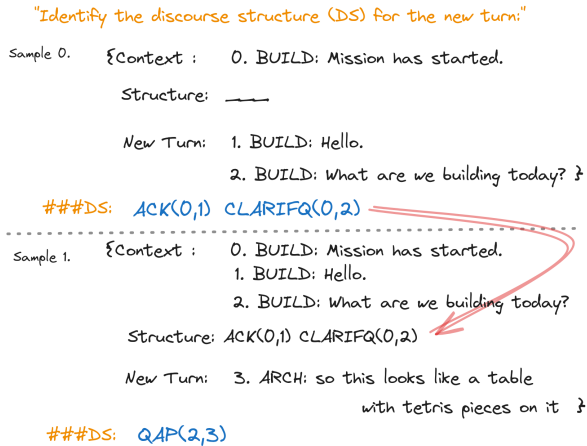


Figure 2: Depiction of dialogue increments seen during Llamipa3+p generation. During generation, the predicted relations are added to the context structure for the following increment.

EEUs in the current dialogue turn.

In the data, we represent relation instances as elementary unit index pairs connected by a relation type, e.g. $RES(0,1)$, for the instance of a *Result* relation holding between units 0 and 1. During training, each sample contains the previous EDUs and EEUs within the window size, the gold structure, which is a list of relation instances between these elementary units, and finally, the units of the current turn. The model is trained to predict the gold relation instances that connect the units of the current turn with previous ones.

As LLMs can handle longer input sequences, we use a window size $k = 15$ which allows the model to exploit a large amount of contextual information. Since discourse links follow an inverse logarithmic curve (Afantenos et al., 2015), there is little statistical support for links of distance longer than 15 in the SDRT corpora.

We finetune Llama 2 (Touvron et al., 2023), and Llama 3 (Dubey et al., 2024) to produce Llamipa2 and Llamipa3 using QLoRA (Dettemers et al., 2023). Both the models are trained for 3 epochs. The model parameter details are in Table 6 in Appendix A. We release Llamipa3 publicly through the Hugging Face platform².

We test our models in two settings. In the first setting, the models predict the relation instances for samples containing the gold discourse structure. In the second, the model uses its predicted relation instances instead of the gold structure in the input

	MSDC		STAC Situated	
	Link	Link+Rel	Link	Link+Rel
Llamipa2+g	0.8561	0.7664	-	-
Llamipa2+p	0.8579	0.7570	-	-
Llamipa3+g	0.9004	0.8154	-	-
Llamipa3+p	0.8830	0.7951	0.8612	0.7796
BERTLine	0.7870	0.6901	0.7667	0.6788

Table 2: Performance of Llamipa2 and Llamipa3 on the MSDC test set for link prediction, evaluated using micro-averaged F1, and link+relation prediction, evaluated using a weighted average F1. We show results for Llamipa using gold structure (g) and predicted structure (p) as input. We also show link and link+relation results for Llamipa3+p finetuned on STAC-Sit and evaluated on STAC-Sit test with predicted inputs. All the results are for relation distance ≤ 10 EDUs, the cutoff used by BERTLine.

sample. That is, at each step, the structure fed to the model is just the accumulated relation instances predicted within the window up until that step (see Figure 2). This incremental setting provides a clear idea of the model’s ability to build a complete dialogue structure from scratch.

6 Results

We run finetuning experiments using the MSDC corpus, and assess Llamipa’s performance on link and link+relation prediction under the two test settings described in Section 5. We consider only the well-formed outputs during evaluation, but note that Llamipa outputs only well-formed relation instances for all experiments.

Table 2 shows that Llamipa3 significantly outperforms BERTLine and outperforms Llamipa2 both using gold structure (Llamipa3+g vs. Llamipa2+g) and predicted structure (Llamipa3+p vs. Llamipa2+p).

To test the generalizability of our approach, we also finetune Llamipa on STAC-Sit and compare the Llamipa3+p outputs to BERTLine. Llamipa3+p shows a gain of 10 F1 points on link prediction and 10.5 F1 points on link+relation prediction. As mentioned in Section 5, the predicted context setting best demonstrates the desired functionality of an incremental parser. Since Llamipa3+p give the best performance in this setting, we use it for all evaluations and henceforward refer to it using the shorthand *Llamipa*.

Performance on STAC-L. In order to compare Llamipa to a broader range of SDRT discourse

²<https://huggingface.co/linagora/Llamipa>

	STAC-L		STAC- L/Molweni-clean	
	Link	Link+Rel	Link	Link+Rel
Perret et al. (2016)	0.69	0.531	-	-
Shi and Huang (2019)	0.719	0.536	-	-
Liu and Chen (2021)	0.731	0.571	-	-
Wang et al. (2021)	0.734	0.573	-	-
Fan et al. (2022)	0.736	0.574	-	-
Chi and Rudnicky (2022)	0.744	0.596	0.645	0.380
Li et al. (2024)	0.593	0.386	0.756	0.312
Bennis et al. (2023) (BERTLine)	0.730	0.562	-	-
Fan et al. (2024) (ChatGPT)	0.599	0.252	-	-
Chan et al. (2024) (ChatGPT)	0.213	0.074	-	-
Llamipa-3+p	0.775	0.607	0.712	0.405

Table 3: A comparison of the performance of various models on STAC-L, the linguistic version of the STAC dataset, as well as out-of-domain performance on a re-annotated version of the Molweni test set provided by Li et al. (2024). Our STAC-L is a flattened version of STAC-L in the STAC repository. See Appendix B for discussion of STAC-L possible variants.

parsers, we tested it on STAC-L, as prior work has only used this linguistic version of STAC. Table 3 shows that Llamipa finetuned on STAC-L and evaluated on the STAC-L test set scores substantially higher on both link and link+relation prediction than the state-of-the-art parser (Chi and Rudnicky, 2022). We note that Llamipa has a lower performance on STAC-L than on either of the situated corpora. This reflects the findings of Asher et al. (2016, 2021), who reported that annotating the STAC linguistic data set was more difficult; annotators felt that the nonlinguistic information needed to determine the discourse structure was missing (see Section 4).

We also test Llamipa on transfer learning. Following Li et al. (2024), we finetune on STAC-L and test on their Molweni-clean test set. While Llamipa outperforms all models on STAC-L, it fares worse on link prediction than Li et al. (2024) on Molweni, though it does better on link+relation prediction; a counterintuitive result given that link+relation prediction is the harder task. Taking a closer look at the link prediction results in Li et al. (2023), we see their model performs well on links of distance 1, but drops after that ($\sim 25\%$ accuracy for distance 2, and 0 for distances greater than 5). Given that Molweni is a relatively simple corpus where most of the links are adjacent, their model excels on the link task. But for relation type prediction, Li et al. (2024) use the traditional method, which assigns relation

types after establishing the attachment structure, and using only local information. Llamipa surpasses their model on relation type prediction, which shows that an approach predicting link and relation simultaneously performs better than a local pipeline approach, even for shallower discourse structures.

Performance on MSDC. Llamipa shows impressive gains for predictions of certain relation types. Looking at F1 breakdown by relation type in Table 5, we see the largest relative increase comes from improved predictions for *Narration* and *Correction*, which Thompson et al. (2024) cite as the most difficult to predict given that they are heavily context dependent. *Narrations* are frequently quite long (with a significant proportion above distance 10), and difficult to capture by local parsers like BERTLine (Thompson et al., 2024). Table 4 shows *Narration* F1 at distances between 2 and 15 EDUs. Like other local parsers, BERTLine requires a distance cutoff that balances ignoring longer relations with mitigating the class imbalance caused by numerous unlinked candidates; for SDRT corpora, this is distance 10. As indicated in Section 5, Llamipa continues to make good predictions up to distance 15.

Further, Llamipa is able to predict many instances of the *Correction* relation, which also has a number of long distance instances in the MSDC. Table 5 shows that its performance is 73 F1 points, which is 42 points higher than and more

Distance	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Llamipa-3	0.90	0.88	0.82	0.76	0.82	0.80	0.79	0.71	0.90	0.76	0.50	0.40	1.0	n/a
BERTLine	0.70	0.61	0.43	0.32	0.38	0.02	0.21	0.28	0.08	n/a	n/a	n/a	n/a	n/a

Table 4: F1 score for Llamipa and BERTLine on Narration instances in the MSDC test with distance between 2 and 15 EDUs. There were 0 instances of Narration at distance 15 in gold. BERTLine’s input was limited to pairs whose distance was ≤ 10 .

than double that of BERTLine.

We hypothesize that Llamipa has a high F1 on *Narration* and *Correction* because it has access to rich semantic information about the discourse context: it can use the previous discourse moves as well as an explicit representation of their discourse structure. This hypothesis coheres with our understanding of how *Narration* functions in the MSDC. *Narrations* connect high-level instructions between which participants clarify, question and elaborate on those instructions, as well as correct missteps (see Figure 1). While there is some statistical similarity in their surface forms, the important cues for isolating a new instruction, and hence a *Narration* attachment, involve certain patterns of *Result* and *Acknowledgement* links in the prior context, as well as the presence of (nonlinguistic) block placements. This pattern is regular enough that one can find most *Narrations* programmatically as with Thompson et al. (2024)’s “Second Pass”. This takes predictions from a first pass of the local model for *Results* and *Acknowledgements*, then applies an algorithmic approach to predict long distance *Narration* links. We tried this second pass for Llamipa, however there was no improvement on F1 for *Narration* scores, indicating that Llamipa has already internalized this pattern during training.

Llamipa offers improvements over BERTLine for many other relation types which are typically not considered long distance. Table 5 shows Llamipa has higher scores for *Result*, *Acknowledgment*, *Elaboration*, *Question-answer Pair*, *Clarification Question*, and *Comment*. We hypothesize that predictions on these also profit from complex contextual information, something that we discuss in more detail in Section 7. Conversely, for relations where Llamipa gives slightly poorer results—*Contrast*, *Alternation*, and *Question-Elaboration*—we surmise that these types can be inferred reliably just from the local information given by the related EDUs or EEU’s.

Overall, Llamipa has several advantages. It simplifies previous attach-then-type pipeline

methods by predicting link and relation labels simultaneously. Previous research has given theoretical arguments that co-predicting attachment and relation labels should give better results (Asher and Lascarides, 2003), as well as empirical evidence for this claim (Bennis et al., 2023). And as we have demonstrated, Llamipa outperforms other parsers on all relevant corpora.

7 The role of context

So far, the hypothesis has been that Llamipa’s state-of-the-art performance, especially for longer distance relations, is due to its ability to process richer contextual information. Whether this is a function of the superior embeddings conferred by the underlying LLM, or a function of a larger context window that allows a longer dialogue history and explicit discourse structure as prediction input, is something we explore further in this section.

First ablation: embeddings vs. context size. This ablation tests whether the LLM embeddings for EDUs and EEU’s account for Llamipa’s performance. We finetune Llama3-8b over pairs of elementary units as input to predict links and relation types, resulting in a local model roughly comparable to BERTLine. Table 5 shows that, compared to BERTLine, Llama3-local still performed better without using the prior discourse context for link and link+relation prediction. However, when compared to Llamipa, Llama3-local still has considerably lower predictions overall. Llamipa’s relation predictions for many relations are also better, notably for *Narration*. This supports our conclusion that although Llama’s superior embeddings play a role, many relations, and especially relations like *Narration*, require the information available in Llamipa’s larger context window.

Second ablation: perturbing input structure. Even if Llamipa requires the input available in a larger context window, it is still an open question whether it uses the discourse structure provided in this window.

	BERTLine	Llama3-local	Llamipa3+p	Llamipa3+g	Llamipa+rand	Llamipa+∅
Result	0.85	0.86	0.90	0.91	0.83	0.72
Acknowledgement	0.81	0.85	0.86	0.86	0.81	0.72
Narration	0.50 [0.73]	0.54	0.82	0.91	0.28	0.18
Elaboration	0.75	0.77	0.77	0.77	0.75	0.68
Correction	0.31	0.64	0.73	0.80	0.49	0.52
Continuation	0.44	0.47	0.49	0.50	0.44	0.33
Question-answer Pair	0.76	0.80	0.82	0.83	0.42	0.14
Comment	0.50	0.57	0.60	0.61	0.57	0.54
Confirmation-Question	0.86	0.89	0.93	0.93	0.91	0.89
Clarification-Question	0.61	0.66	0.73	0.73	0.70	0.41
Contrast	0.80	0.79	0.75	0.74	0.72	0.68
Question-Elaboration	0.39	0.36	0.30	0.36	0.33	0.33
Alternation	0.88	0.88	0.83	0.88	0.90	0.96
Explanation	0.00	0.00	0.00	0.00	0.00	0.00
Conditional	0.58	0.00	0.00	0.00	0.00	0.00
Sequence	0.00	0.00	0.00	0.00	0.00	0.05
Link+Rel F1	0.69 [0.71]	0.73	0.80	0.81	0.65	0.56
Link F1	0.78 [0.80]	0.82	0.88	0.90	0.77	0.72

Table 5: Performance breakdown for BERTLine, Llamipa2 and Llamipa3, Llama3-local by relation type on MSDC. Relation types are listed in descending order from most to least number of instances in the corpus. The [F1] scores for Narration after the Second Pass on BERTLine as reported in Thompson et al. (2024) shown in brackets. The three columns on the right show the effects of varying the context structure during Llamipa’s generation step.

We already have some indication that Llamipa *does* use the structure when we look at Llamipa’s performance given a gold discourse structure for a context versus its predicted structure, as described in Section 5. The results of this comparison are shown in Table 5: there is a small but consistent drop between gold and predicted structure contexts in F1 performance for all relation types.

To investigate this further, we perform a second ablation where we perturb the input structure for each test sample by randomizing both the relation type and EDU index for each relation instance in the structure. When we run Llamipa on this randomized input (denoted as Llamipa+rand in Table 5) we see a drop in performance from Llamipa3+p for link and link+relation for most relation types, with a substantial drop for *Narration*, *Correction*, and *Question-answer Pair*.

Third ablation: removing input structure. In this ablation, we removed all structure from Llamipa’s input context, leaving only a sequence of elementary units. The performance of Llamipa on this data set (in the column Llamipa+∅ of Table 5) shows an even larger drop in both link and link+relation predictions compared to Llamipa+rand. Similar to Llamipa+rand, we see a substantial drop for the relation types *Narration*, *Correction*, and *Question-Answer Pair*. Additionally, the performance dropped drastically for *Clarification Questions*.

Fourth ablation: targeted perturbations. The

results of the previous ablations support the hypothesis that the structure provided in the context window *does* have an effect on relation instance predictions. This final ablation investigates whether Llamipa predicts context dependent relations as well as it does because it leverages (or “reasons” about) the relation instances in the context structure in a manner consistent with our semantic intuitions. Here, we focus on instances of *Question-answer pair (QAP)* and *Correction*.

QAP attachments provide a straightforward example of context dependence. In the MSDC, it is common that if a prediction $QAP(y,z)$ is made for a dialogue turn z then $CLARIFQ(x,y)$, $CONFQ(x,y)$, or $QELAB(y,z)$ is in its context. That is, the questions that are answered in the *QAP* relations tend to be questions asked to clarify, confirm or elaborate upon the preceding context. To test whether Llamipa’s predictions conform to this pattern, we isolate the test samples where Llamipa correctly predicts *QAP* (Table 5), and remove all question relation types from the structure given in the context of these samples. We re-run the generation step on this amended sample set, and see a 50 point drop in F1, from 0.87 to 0.37, for *QAP* predictions, and no change for all other relation types (see Figure 3).

Correction relations are among the most difficult to identify in the MSDC, as discussed in Section 6. Although the language used for *Corrections* usually features some form of negation—*No*, *Not like*

<pre>{Context : 0. BUILD: Mission has started. 1. BUILD: Hello. 2. BUILD: What are we building today? Structure: ACK(0,1) CLARIFQ(0,2) New Turn: 3. ARCH: so this looks like a table with tetris pieces on it. } ###DS: ? QAP(3,4)</pre>	<pre>{Context : ... 50. BUILD: put purple (-2,3,0) 51. BUILD: Like that? 52. ARCH: One more towards you. ACK 53. ARCH: It should be hovering. } Structure: ... CONFQ(50,51) QAP(51,52) CORR(50,52) ELAB(52,53) New Turn: 54. BUILD: put purple (-3,3,0) remove purple (-2,3,0) } ###DS: RES(53,54) ? CORR(50,54)</pre>
---	--

Figure 3: Depiction of targeted perturbation ablations described in Section 7. Left: We remove the *Clarification-question* from the structure of samples that should predict *Question-answer pair relations*, and see a 50 point drop in F1 for *Question-answer pair* prediction. Right: We change the first *Correction* in the structure to *Acknowledgement* in samples that should predict a second *Correction* in response to the first.

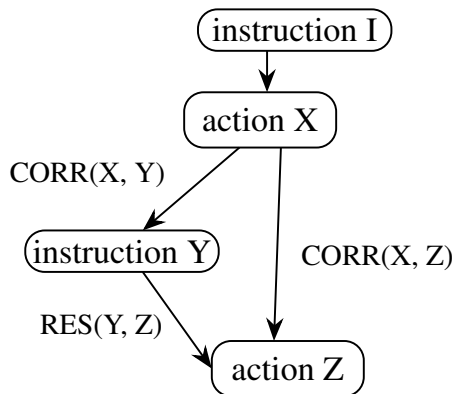


Figure 4: Correction triangles are composed of a *Correction* to an action sequence, followed by a new action sequence that is the *Result* of the corrective move and a *Correction* of the first action sequence.

that, Oops!—deciding where it attaches requires making inferences about how the content of the corrective move applies to a more global information state. Although this is beyond the capability of a discourse parser like Llamipa, there are still structural regularities in the MSDC that can help with retrieving *Corrections*.

For instance, most *Corrections* come about when the Builder makes the wrong moves in response to an instruction, and the Architect then gives a further instruction, which results in moves that address this by fixing the last wrong moves. In the discourse structure, these interactions create recurring patterns, “correction triangles” (Figure 4). These triangles provide a way to check Llamipa’s sensitivity to $CORR(x,y)$ type relations in the context structure in cases where it must predict $RES(y,z)$ $CORR(x,z)$, the last two relations in a triangle. As with the *Question-answer Pair* test, we single out each test sample in which Llamipa correctly predicts the second *CORR*; but instead of

removing $CORR(x,y)$, we change it to $ACK(x,y)$ (*Acknowledgement*(x,y)). This leaves the structure intact, but gives it a contrary sense. We re-run the generation step and see a 20 point drop in F1, from 0.71 to 0.52, for *Correction* predictions, and no change for all other relation types.

8 Conclusions

Llamipa is an incremental SDRT style discourse parser that uses its prior discourse structure predictions to predict links and relation labels for the discourse units in new conversation turns. As such, it can be naturally assimilated into downstream applications like conversational agents.

After finetuning both Llama2-13b and Llama3-8b models to produce Llamipa2 and Llamipa3, we show that Llamipa3 has substantially higher F1 scores than other parsers that use SDRT data sets on the link and link+relation tasks. We confirm that its performance is due to access to larger swaths of context, which are not available to parsers that use a local model. We further find evidence to support the hypothesis that Llamipa uses the explicit representations of structure in its context to predict new relation instances.

Notably, Llamipa excels on data sets that feature conversation containing both linguistic and nonlinguistic moves (actions, in the MSDC) that help interpret them. This indicates that a parser like Llamipa can furnish important discourse information to builder action models; preliminary demonstrations of this can be found in Chaturvedi et al. (2024).

Limitations

We did not exhaustively test all models on the situated data sets. Given the number of MPDUs especially in the MSDC, however, we felt that their performance would not compare to Llamipa’s. We plan to do more extensive testing in the future to compare models.

Llamipa, like previous parsers, presumes that the EDUs/EEUs are already segmented. However, existing work on automatic segmentation suggests that this task is relatively tractable on diverse data sets (Muller et al., 2019) including even corpora of spoken conversation (Gravellier et al., 2021; Prevot et al., 2023)

Like all supervised models, Llamipa is limited to making predictions on data sets that are similar to its training data. Given the heavy annotation effort required to build data sets annotated for discourse structure, this arguably puts a bottleneck on the generalizability of our approach to data sets in languages other than English and for conversation styles, such as multiparty spoken conversation, that are significantly different from the chat-based dialogues in the MSDC and STAC. Future work will explore efficient ways of harnessing extant data sources through, for instance, automatic translation, data augmentation and/or heuristics, to facilitate the transfer to other languages and domains.

Ethics Statement

Our work here has been to improve the capacities of AI models to understand the semantic structure of discourse and conversation. We see no direct ethical concerns that arise from this work, though we hope that a better grasp of the semantic structure of conversation may enable the output of models to be more pragmatically grounded.

Acknowledgements

For financial support, we thank the National Interdisciplinary Artificial Intelligence Institute ANITI (Artificial and Natural Intelligence Toulouse Institute), funded by the French ‘Investing for the Future–PIA3’ program under the Grant agreement ANR-19-PI3A-000. This project has also been funded by the France 2030 program and is funded by the European Union - Next Generation EU as part of the France Relance. We also thank the projects COCOBOTS (ANR-21-FAI2-0005) and DISCUTER (ANR-21-ASIA-0005), and

the COCOPIL “Graine” project of the Région Occitanie of France. This work was granted access to the HPC resources of CALMIP supercomputing center under the allocation 2016-P23060.

References

- Afantenos, S., Kow, E., Asher, N., and Perret, J. (2015). Discourse parsing for multi-party chat dialogues. In *Empirical Methods in Natural Language Processing*, pages 928–937. Association for Computational Linguistics.
- Asher, N. (1993). *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers.
- Asher, N., Hunter, J., Morey, M., Benamara, F., and Afantenos, S. (2016). Discourse structure and dialogue acts in multiparty dialogue: The STAC corpus. In *Language Resources and Evaluation Conference*, pages 2721–2727.
- Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Cambridge University Press.
- Asher, N., Paul, S., and Russell, C. (2021). Fair and adequate explanations. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pages 79–97. Springer.
- Bennis, Z., Hunter, J., and Asher, N. (2023). A simple but effective model for attachment in discourse parsing with multi-task learning for relation labeling. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3404–3409.
- Chan, C., Jiayang, C., Wang, W., Jiang, Y., Fang, T., Liu, X., and Song, Y. (2024). Exploring the Potential of ChatGPT on Sentence Level Relations: A Focus on Temporal, Causal, and Discourse Relations. In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 684–721.
- Chaturvedi, A., Thompson, K., and Asher, N. (2024). NeBuLa: A discourse aware Minecraft Builder. *arXiv preprint arXiv:2406.18164*.
- Chi, T.-C. and Rudnicky, A. (2022). Structured Dialogue Discourse Parsing. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 325–335.
- Cho, K., van Merriënboer, B., Bahdanau, D., and Bengio, Y. (2014). On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In Wu, D., Carpuat, M., Carreras, X., and Vecchi, E. M., editors, *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.

- Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. (2023). QLoRA: Efficient finetuning of quantized LLMs. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S., editors, *Advances in Neural Information Processing Systems*, volume 36, pages 10088–10115. Curran Associates, Inc.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J., Doran, C., and Solorio, T., editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. (2024). The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Fan, Y., Jiang, F., Li, P., and Li, H. (2024). Uncovering the potential of ChatGPT for discourse analysis in dialogue: An empirical study. In Calzolari, N., Kan, M.-Y., Hoste, V., Lenci, A., Sakti, S., and Xue, N., editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16998–17010, Torino, Italia. ELRA and ICCL.
- Fan, Y., Li, P., Kong, F., and Zhu, Q. (2022). A distance-aware multi-task framework for conversational discourse parsing. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 912–921.
- Forbes, K., Miltsakaki, E., Prasad, R., Sarkar, A., Joshi, A., and Webber, B. (2003). D-LTAG system: Discourse parsing with a lexicalized tree-adjoining grammar. *Journal of Logic, Language and Information*, 12:261–279.
- Gravellier, L., Hunter, J., Muller, P., Pellegrini, T., and Ferrané, I. (2021). Weakly supervised discourse segmentation for multiparty oral conversations. In *2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)*, pages 1381–1392. ACL: Association for Computational Linguistics.
- Guz, G., Huber, P., and Carenini, G. (2020). Unleashing the Power of Neural Discourse Parsers - A Context and Structure Aware Approach Using Large Scale Pretraining. In Scott, D., Bel, N., and Zong, C., editors, *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3794–3805, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Li, C., Braud, C., Amblard, M., and Carenini, G. (2024). Discourse Relation Prediction and Discourse Parsing in Dialogues with Minimal Supervision. In *Proceedings of the 5th Workshop on Computational Approaches to Discourse (CODI 2024)*, pages 161–176.
- Li, C., Huber, P., Xiao, W., Amblard, M., Braud, C., and Carenini, G. (2023). Discourse Structure Extraction from Pre-Trained and Fine-Tuned Language Models in Dialogues. In Vlachos, A. and Augenstein, I., editors, *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2562–2579, Dubrovnik, Croatia. Association for Computational Linguistics.
- Li, J., Liu, M., Kan, M.-Y., Zheng, Z., Wang, Z., Lei, W., Liu, T., and Qin, B. (2020). Molweni: A Challenge Multiparty Dialogues-based Machine Reading Comprehension Dataset with Discourse Structure. pages 2642–2652.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach.
- Liu, Z. and Chen, N. (2021). Improving Multi-Party Dialogue Discourse Parsing via Domain Integration. In *Proceedings of the 2nd Workshop on Computational Approaches to Discourse*, pages 122–127, Punta Cana, Dominican Republic and Online. Association for Computational Linguistics.
- Lowe, R., Pow, N., Serban, I., and Pineau, J. (2015). The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In Koller, A., Skantze, G., Jurcicek, F., Araki, M., and Rose, C. P., editors, *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic. Association for Computational Linguistics.
- Maekawa, A., Hirao, T., Kamigaito, H., and Okumura, M. (2024). Can we obtain significant success in RST discourse parsing by using Large Language Models? In Graham, Y. and Purver, M., editors, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2803–2815, St. Julian’s, Malta. Association for Computational Linguistics.
- Mann, W. and Thompson, S. (1987). Rhetorical Structure Theory: A Framework for the Analysis of Texts. *International Pragmatics Association Papers in Pragmatics*, 1:79–105.
- Marcu, D. (1999). A Decision-Based Approach to Rhetorical Parsing. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 365–372, College Park, Maryland, USA. Association for Computational Linguistics.
- Muller, P., Afantenos, S., Denis, P., and Asher, N. (2012). Constrained Decoding for Text-Level Discourse Parsing. In *Proceedings of COLING 2012*, pages 1883–1900, Mumbai, India. The COLING 2012 Organizing Committee.

- Muller, P., Braud, C., and Morey, M. (2019). ToNy: Contextual Embeddings for Accurate Multilingual Discourse Segmentation for Full Documents. In *Discourse Relation Parsing and Treebanking*, pages 115–124.
- Narayan-Chen, A., Jayannavar, P., and Hockenmaier, J. (2019). Collaborative Dialogue in Minecraft. In Korhonen, A., Traum, D., and Màrquez, L., editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5405–5415, Florence, Italy. Association for Computational Linguistics.
- Perret, J., Afantenos, S., Asher, N., and Morey, M. (2016). Integer Linear Programming for Discourse Parsing. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 99–109, San Diego, California. ACL.
- Prasad, R., Dinesh, N., Lee, A., Miltsakaki, E., Robaldo, L., Joshi, A. K., and Webber, B. L. (2008). The Penn Discourse TreeBank 2.0. In *LREC*. Citeseer.
- Prevot, L., Hunter, J., and Muller, P. (2023). Comparing Methods for Segmenting Elementary Discourse Units in a French Conversational Corpus. In *24th Nordic Conference on Computational Linguistics (NoDaLiDa 2023)*. ACL: Association for Computational Linguistics; University of Tartu Library.
- Shi, Z. and Huang, M. (2019). A Deep Sequential Model for Discourse Parsing on Multi-Party Dialogues. In *AAAI*.
- Thompson, K., Hunter, J., and Asher, N. (2024). Discourse Structure for the Minecraft Corpus. In Calzolari, N., Kan, M.-Y., Hoste, V., Lenci, A., Sakti, S., and Xue, N., editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4957–4967, Torino, Italia. ELRA and ICCL.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. (2023). Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv preprint arXiv:2307.09288*.
- Wang, A., Song, L., Jiang, H., Lai, S., Yao, J., Zhang, M., and Su, J. (2021). A Structure Self-Aware Model for Discourse Parsing on Multi-Party Dialogues. In Zhou, Z.-H., editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3943–3949. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Wang, T., Wan, X., and Jin, H. (2020). AMR-To-Text Generation with Graph Transformer. *Transactions of the Association for Computational Linguistics*, 8:19–33.
- Yung, F., Ahmad, M., Scholman, M., and Demberg, V. (2024). Prompting Implicit Discourse Relation Annotation. In Henning, S. and Stede, M., editors, *Proceedings of The 18th Linguistic Annotation Workshop (LAW-XVIII)*, pages 150–165, St. Julians, Malta. Association for Computational Linguistics.
- Zhu, J., Li, J., Zhu, M., Qian, L., Zhang, M., and Zhou, G. (2019). Modeling Graph Structure in Transformer for Better AMR-to-Text Generation. In Inui, K., Jiang, J., Ng, V., and Wan, X., editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5459–5468, Hong Kong, China. Association for Computational Linguistics.

A Appendix: Model Parameters

version of STAC-L available on the STAC corpus website cited above.

GPUs	
4 NVIDIA Volta V100	
Hyperparameters	
Training epochs	3
batch size	4
optimizer	Adam
learning rate	2e-4
learning rate scheduler	linear warm-up and cosine annealing
warm-up ratio	0.03
gradient clipping	0.3
lora r	64
lora (alpha)	16
lora dropout ratio	0.1
lora target modules	Only Attention Blocks (q_proj, v_proj)
quantization for Llama3	4-bit NormalFloat

Table 6: Details on computing resources and hyperparameters for finetuning Llamipa.

Table 6 gives the hyperparameters we use for finetuning Llamipa-2 and Llamipa-3 along with the computing resources. We adapt the finetuning code from the following repository³.

B Appendix: STAC Linguistic data

We derive STAC-L from the STAC data available on the website (<https://www.irit.fr/STAC/corpus.html>) by adopting standard procedures for flattening CDUs. However, we find that our data set did not quite match that of Bennis et al. (2023) in terms of numbers of EDUs. Shi and Huang (2019) also report a different number of EDUs from what we have *and* from what Bennis et al. (2023) report. When we finetune and test Llamipa on the STAC-L data available in the BERTLine repository⁴, we had lower scores than what we report in Table 3. Llamipa-3+p had an F1 for link of 75.4 and a link+relation F1 of 57.7 on this version of STAC linguistic. On the transfer learning from STAC linguistic to the Molwenic test set, Llamipa-3+p had an F1 of 68.5 for link and 38.7 for link+relation. We make our

³https://github.com/mlabonne/llm-course/blob/main/Fine_tune_Llama_2_in_Google_Colab.ipynb

⁴<https://github.com/zineb198/LineBert>