Multilingual AMR Discourse Coherence: Representation of a Social Movement on Social Media

Anonymous ACL submission

Abstract

Combining AMR with social sciences and discourse coherence is a new interdisciplinary research focus that has potential to explore the meaning of online public discourse through parsed sentence structure of the Abstract Meaning Representation (AMR) graph. The objective of this research is to provide a better sense of how the anti-femicide movement in Mexico is circulated on social media by understanding the narrative structure of post comments. This work attempts to reproduce a similar research framework as Pournaki and Willaert's (2024) and their political discourse narratives on social media text comments about the anti-femicide movement in Mexico. It uses a three step process to translate the data into all English, parse the data into an AMR graph, and analyze that graph to understand the narrative surrounding the movement. It found that the process is almost generalizable, but only provides legitimate findings with specific data types.

1 Introduction

Consuming news about the anti-femicide social movement in Mexico online is difficult to follow and lacks a coherent narrative message. The movement has been framed in various ways throughout academic literature as well as in the media since the 1990s. Understanding what is being circulated online can support how decisions are made across various organizations (gov, NGOs, SM groups). It can also help us understand how the public consumes information on the issue which informs how practitioners share their messaging campaigns. The objective of this research is to provide a better sense of how the anti-femicide movement in Mexico is circulated on social media by understanding the narrative structure of post comments. The research explores the question, *How do we understand the narrative discourse of the online anti-femicide movement by using Abstract Meaning Representation (AMR) structure?*

This paper follows previous work on extraction and representation of public discourse narratives from text, by using Abstract Meaning Representation (AMR) and narratology (Banarescu et al., 2013; Onega and Landa, 1996; Pournaki and Willaert, 2024), but pushes further to identify a social movement narrative in multilingual text. AMRs are graph structures of a sentence's semantic representation, which essentially provide a root based ontology to understand a sentence's meaning (Banarescu et al., 2013). Narratology describes a narrative as a semiotic model of a series of events intentionally connected in a temporal or causal way (Onega and Landa, 1996). In their work, Pournaki and Willaert's (2024) use AMR to support continuing empirical analysis on research in political discourse and narratives. The authors use AMR to formulate a political narrative ontology of State of the European Union Addresses using political worldviews as the root and follow a process of first creating the graph representation of political meanings and then using narratology to filter the graphs for specific representations between actors and events (Pournaki and Willaert, 2024).

Pournaki and Willaert's (2024) found that the AMR Narrative structures enabled major actors, their roles, and evolving political priorities to be identified in the State of the European Union data through narrative shifts, emerging themes, and actor relationships. They describe the extraction and analysis of the formulated discourse structure as a type of ontology to represent stories of political 'actors', the events where these 'actors' figure, and trace the perspectivization of the events as core narrative signals (Pournaki and Willaert, 2024). The authors were able to identify latent goals of political actors and narratives that evolve based on geopolitical events (Pournaki and Willaert, 2024). Working within the space of computational social science, multilingual NLP, and Machine Translation, this present research explores the discourse coherence of social media comments to analyze the narrative structure of the anti-femicide movement in Mexico. This exploratory research follows the narrative discourse groundwork carried out by Pournaki and Willaert's (2024).

1.1 Anti-Femicide Movemnt Discourse Structure

As a case study for understanding discourse narratives on social media, the anti-femicide movement in Mexico is used in this research to identify how we as spectators interpret what is being said about the phenomenon. Femicide is aggravated homocide due to gender, a short description of the phenomenon can be found in Appendix A (Staudt et al., 2009).

The anti-femicide movement in Mexico, spreads across multiple online accounts and platforms, which initiates decentralized fragments of the movement's goals (D'Ignazio et al., 2020). Much of the movement's content is a variation of informational or a call to action, or a combination of these, and can be understood as a political, social, and emotional topic with complex implications (Rohm et al., 2023). In short, the content points towards the same cause, represented by various perspectives to be consumed by audiences to do what they will. Structuring anti-femicide social media content into graphs provides a way to group distinctive features of the narrative and identify relationships within and between each of these groups. Evaluating these relationships emphasizes how knowledge is presented to support meaning making for social movements.

2 Previous Literature

Discourse coherence and multilingual AMR for social movement data is a potentially novel approach. Previous work in narrative and discourse coherence for public datasets and multilingual AMR exists, but not in a combined fashion to understand a narrative of online social movement messaging. Below is a review of the literature in these fields.

2.1 Discourse Coherence with AMR

Conflicting and fragmented information is not ideal for social movement practitioners, especially when they use social media to spread their cause. Current communication practices of social movements almost always require an online presence to connect with members and potential supporters (Peterson-Salahuddin, 2022). To better interpret the multiple perspectives of online content about the social movement, discourse coherence is used here as a way to carryout the methods.

2.2 Multilingual AMR Processes

Because the structure of the AMR parser is inherently build to process English, translating sentences into English before parsing with the graph has produced better results (Regan et al., 2024; Wein and Schneider, 2024). For example, Wein and Schneider's (2024) attempt to identify divergences between AMR graphs parsed before and after text was translated to English. Compared to the graph that parsed the non-english text and then translated that graph to English they find that translating sentences into English before paring with the AMR graph produced higher Smatch scores (52%), a metric that calculates the degree of overlap between two semantic feature structures (Cai and Knight, 2013). Wein and Schneider's (2024) show lower similarity scores for Chinese text (.25) and Spanish text (.30), and the Smatch scores were about 10% lower than their English counterpart (41% - 42%).

3 Research Framework

3.1 Data Collection

Data collected from web-scraped hashtags and accounts include initial comments by the account holder, images, likes, hashtags, URLs, account information, and automated descriptions from Instagram. This work only uses the comments connected to the initial post by the account and a post-id. The comments are a few sentences or phrases in Spanish and English, and include emojis, other tagged accounts, and hashtags. Initially, 17,695 posts were collected and after removing duplicates, 12,357 posts were kept for this experiment.

3.2 Methods

This framework explores multilingual AMR, as shown in Fig. 1, first by implementing machine translation of the data from a mix of Spanish and English to English. Second, the translated data is processed and put through IBM's Transition-Based Neural parser to identify if the graph can provide a coherent structure (Zhou et al., 2021). Third, a narrative discourse structure, specifically pertaining to the anti-femicide movement in Mexico, will attempt to be interpreted through actors, events, and how they work with each other.

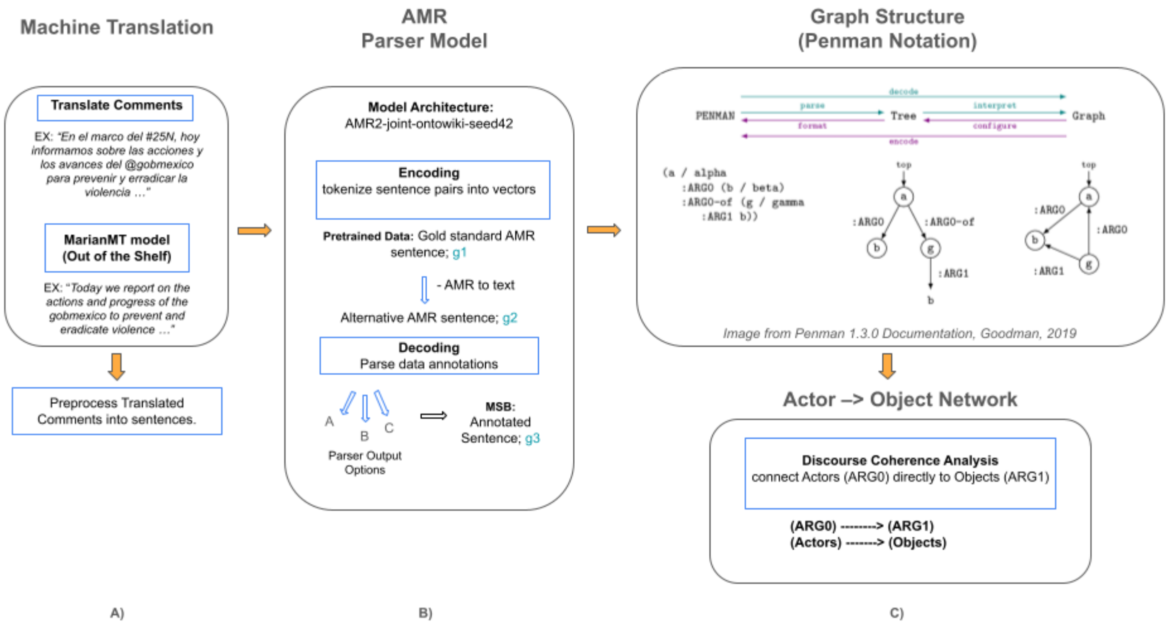

Figure 1: Analysis Framework: **A)** MarianMT model: A BartForConditionalGeneration model code with transformer encoder-decoders with 6 layers in each component. **B)** Encoding takes pretrained AMR sentences, g1, and identifies the text to relate to the most appropriate sentence resting to the data, g2. Decoding vectors to construct AMR graphs with the transition-based neural parser (This is where we parse the data to show word-node alignments).**C)** The AMR output is taken and parsed into a human and machine readable tree and graph structures using Penman Notation. More details can be found in Appendix C.

3.2.1 Translation

It was found that translating text from non-English to English before parsing the English based AMR model proved to show better semantic representation results (Wein and Schneider, 2024; Regan et al., 2024). The post comments were translated using a pre-trained model to show better results in downstream parsing. The HuggingFace translation tokenizer and model from MarianMT were used to translate the 12,378 Instagram post comments from a mix of English and Spanish to English (Junczys-Dowmunt et al., 2018). The particular Spanish to English model uses the same models as BART, and was pre-trained on the OPUS dataset, an open-source corpus that uses collected sentences from around the world, which makes it well equip to handle general-domain text data (Junczys-Dowmunt et al., 2018; Lewis, 2019).

3.2.2 Parsing the AMR

The translated text was then preprocessed by splitting the text into sentences and grouping these by post-ids. The data is then put through the Transition-Based Neural parser (Wang et al., 2015), using the specific model AMR2-joint-ontowiki-seed42, which uses an ensemble distillation of Smatch-based models and Bayesian interpretation which obtained state-of-the-art results for cross-lingual AMR parsing and domain adaptation (Lee et al., 2021). This was the same model used by Pournaki and Willaert's (2024) which will support replication efforts. The model, AMR2-joint-ontowiki-seed42, first encodes tokenized sentences into vectors then decodes those vectors into an AMR graph showing word-node alignment (Lee et al., 2021).

3.2.3 Building the Narrative Structure

The narrative structure is then created by parsing the sentences into the PENMAN graph notation where edges are directed towards the argument connections in an igraph (Goodman, 2020) .

4 Analysis

4.1 AMR Parser Model

The parsed AMR data resulted in a structured dataset after using the model, AMR2-joint-ontowiki-seed42 from Lee et al.'s (2021). Here, GPU units were needed to parse the large dataset of now 24, 173 sentence large dataset. For this parsing, GPU processing was used, first as g3.medium, GPU A100, with 8 CPU cores, 30 GB of RAM and 60 GB of root disk. But this task proved to take over 10+ hours of Burn rate, promoting the parsing to batch the data and move to a g3.XL which took 3 hours.

The model was run twice, once without alignment, and again with alignment, i.e. the alignment parameter set to False and True. This was done because the initial 'False' alignment proved to identify no predicate cases, meaning there would be nothing to align an actor and event in the network. The second had a similar outcome and failed to find alignment identifiers. In this case, predicates found in the sentence are considered alignment identifiers and connect actors to events. Particularly, while processing the data into the customized PENMAN notation (Goodman, 2020) found in Pournaki and Willaert's (2024)'s narrative table builder no predicates were found.

4.2 Discourse Coherence

The pivot to analyze Actor-Object network relations proved to find clusters of expected discourse pairs, but also found outliers that support that narrative. Specifically, action-based discourse communicated across Instagram can be understood as calls to action, and the network shows that actors found in the parser outnumber the objects. This seems expected as social movement discourse online is backed by calls to action, demands, and lists of grievances (Bennett and Segerberg, 2012). The objects in the network can be understood to support the actors in the network, potentially by supporting actions expressed in the comments. A graph of this can be found in Appendix D.

Specifically, the centrality of our entity nodes, regardless of actor or object, show how often a node appears in a sentence and its referential properties. Degree centrality measures the local centrality measures of the immediate neighborhood links within the network (Zhang and Luo, 2017).

$$DC(v_i) = \frac{1}{N-1} \sum_{j=1}^{N} \alpha_{i,j} \tag{1}$$

where N is the total number of nodes, $\alpha_{i,j} = 1$ denotes a link between v_i and v_j, as long as j doesn't equal i. In this case, the top in-degree centrality measures how often an entity is referred to in the network (Zhang and Luo, 2017). Here, the top entities seem to describe emotion, community, and ideas. Table 1 shows that "we" (56), "pain" (50), "abuse" (50), and "feminism" (48) are the subject of the sentences and potentially push forward a narrative within the context of these top words. Appendix E has Table 2 which shows the Out-degree entities, or how often an entity appears as the referring node are: "we" (81), "pain" (77), and "positive" (73) have the highest out-degree in the network.

Entity & (Entity Label)	In-Degree Centrality
we (w)	56
pain-01 (p)	50
abuse-01 (a)	50
feminism (f)	48
so (s)	48
violence (v)	47
this (t)	47
abuse-02 (a2)	44
we (w2)	41
i (i)	40

Table 1: Top 10 Most Influential Entities

5 Discussion and Findings

This replicated Pournaki and Willaert's (2024) study of political narrative discourse using social media anti-femicide movement data can not be one for one replicate because of the structure of the social media comment sentences. Social media posts compared to political speeches have more jargon and dialectical elemnts and need a different AMR parser. Here, the AMR parser did not identify strong enough predicates to build a relational network. This was because the data consisted of phrases, hashtags, and thoughts, not full sentences.

The discourse coherence graph analyzed here is based on actor-object relations to help identify which actors are related to which topics, and review objects central to the discussion. A degree centrality analysis of the network shows that top parsed entities to describe emotion, community, and ideas which is in-line with typical online social-movement discourse (Rohm et al., 2023).

5.1 Future Work

Using parsers that support the type of text that is used in social media online should be explored to replicate the actor-predicate-object structure of the original research. The anti-femicide data has images, incorporating them can provide an extra layer of coherence.

6 Limitations

6.1 Processing Power

While attempting to reproduce this work access to large compute and storage are necessary for carrying out the models. Furthermore, researchers can only use Linux or mac based products which provides a steep curve to learning and replicability. The author(s) of this work have the fortune of attaining access to HPC cloud computing, but others might not.

6.2 Data

The data here was not analyzed in a temporal fashion similar to the reproduced paper so understanding clear actor-event relationships is not exactly apparent. Although this is an interesting finding for this work, it might have different implications for the graph structure and its latent connections.

The original text in this dataset include emojis, other account holders, and specialized hashtags which were not taken into account during this research. This could be an integral part of the semantic infrastructure, especially in the world of social media.

Similarly, the data also include .jpg images that come with the post. This would be an interesting addition to understanding the AMR parser, but again, more computing power would be necessary for this.

Alternatively, some of the comments are clearly bots attempting to promote or sell content that is not associated with the topic. This could potentially bias the graph structure. Identifying these bots or noise in the data could be a helpful solution to better these findings.

7 Ethical Considerations

This study offers theoretical contributions to research on multilingual AMR and social movement networks online. It holds practical implications for computational social scientists using human computation to understand social phenomena, as well as useful tools for members of the social movement. The practical implications of this work come from the analysis of a narrative network and resources for its reproduction to provide insights for social movement groups in need of online contextual information. This work will be included in the overall creation of an Open Science project that holds a code-base for the analysis. A de-identified dataset of collected Instagram posts and their metadata will also be included for future research. Through these efforts this research will follow the FAIR and CARE principles to manage the data as an Open Source project.

References

Heather Robin Agnew. 2015. Reframing "femicide": Making room for the balloon effect of drug war violence in studying female homicides in mexico and central america. *Territory, politics, governance*, 3(4):428–445.

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 discourse*, pages 178–186.

W Lance Bennett and Alexandra Segerberg. 2012. The logic of connective action: Digital media and the personalization of contentious politics. *Information, communication & society*, 15(5):739–768.

Shu Cai and Kevin Knight. 2013. Smatch: an evaluation metric for semantic feature structures. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 748–752, Sofia, Bulgaria. Association for Computational Linguistics.

Catherine D'Ignazio, Helena Suárez Val, Silvana Fumega, Harini Suresh, and Isadora Cruxên. 2020. Feminicide & machine learning: detecting gender-based violence to strengthen civil sector activism.

Michael Wayne Goodman. 2020. Penman: An open-source library and tool for AMR graphs. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 312–319, Online. Association for Computational Linguistics.

Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. Marian: Fast neural machine translation in C++. In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.

Young-Suk Lee, Ramon Fernandez Astudillo, Thanh Lam Hoang, Tahira Naseem, Radu Florian, and Salim Roukos. 2021. Maximum bayes smatch ensemble distillation for amr parsing. *arXiv preprint arXiv:2112.07790*.

Mike Lewis. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.

Susana Onega and José Angel García Landa. 1996. *Narratology: an introduction*. Routledge.

Chelsea Peterson-Salahuddin. 2022. Posting back: Exploring platformed black feminist communities on twitter and instagram. *Social Media+ Society*, 8(1):20563051211069051.

Armin Pournaki and Tom Willaert. 2024. A graph-based approach to extracting narrative signals from public discourse. *arXiv preprint arXiv:2411.00702*.

Michael Regan, Shira Wein, George Baker, and Emilio Monti. 2024. Massive multilingual abstract meaning representation: A dataset and baselines for hallucination detection. *arXiv preprint arXiv:2405.19285*.

Ricardo Henry Dias Rohm, Gabriel de Souza Valuano, and Sofia de Souza Ferreira Xavier. 2023. The role of digital activism within social movements to foster democratization in the global south: the case of ni una menos. *Espacio abierto*, 32(1):29–50.

Kathleen A Staudt, Tony Payan, and Z Anthony Kruszewski. 2009. *Human rights along the U.S.-Mexico border: gendered violence and insecurity*.

Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015. A transition-based algorithm for amr parsing. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 366–375.

Shira Wein and Nathan Schneider. 2024. Assessing the cross-linguistic utility of abstract meaning representation. *Computational Linguistics*, pages 1–55.

Melissa W Wright. 2011. Necropolitics, narcopolitics, and femicide: Gendered violence on the mexico-us border. *Signs: journal of women in culture and society*, 36(3):707–731.

Junlong Zhang and Yu Luo. 2017. Degree centrality, betweenness centrality, and closeness centrality in social network. In *2017 2nd international conference on modelling, simulation and applied mathematics (MSAM2017)*, pages 300–303. Atlantis press.

Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo, and Radu Florian. 2021. Amr parsing with action-pointer transformer. *arXiv preprint arXiv:2104.14674*.

A Appendix: Little context for the anti-femicide movement

The anti-femicide movement in Mexico grew from years of femicide that began at the border of Mexico and Texas in cd. Juarez Chihuahua and moved across the country (Wright, 2011). The first case was documented around 1993 and described as aggravated homicide due to gender, specifically females, sexual violence, dehumanization through the image of discarded female remains, or other visual representation of ultraviolence against a woman leading to murder or death (Staudt et al., 2009; Agnew, 2015).

In 2024 The Met Gala reimagined Geroges Bizet's Carmen in a contemporary setting along the U.S.-Mexico border. The crux of the Opera's ending has the titular character, Carmen, murdered by her lover and the crime is described as "matters of a jealous heart" (Cracknell, 2024). This new interpretation did not describe Carmen's tragic murder as a lover's quarrel but was intended by the production team to be understood as a femicide (Cracknell, 2024). The notion that femicide is a modern-day occurrence that happens in a contemporary setting is indicative of the bravado in this scenario.

B Appendix: Accounts and Hashtag Names in Data Collection

The hashtags include #mexicofeminicida, #méxicofeminicida, #niunamenos, #niunamashmo, #vivasnosqueremos, #noestamostodas, #feminicidiosenmexico, #elmachismomata, #yotecreo, #8M, #nonoscallamosmas, #violenciadegenero, #violenciamachista, #femicidioemergencianacional, #25N, #UnDíaSinNosotras. The accounts include: @apartados8km, @noestamostodas, @siwapazyjusticia, @womansonfire,@brujamixteca, @feminicidiocdmx_, @demachosahombres, @abogadafemina, @redpsicofem.jrz, @colectivoyositecreo, @antimonument_vivasnosqueremos.

C Appendix: Extended Analysis Framework

A) MarianMT model: A BartForConditionalGeneration model code with transformer encoder-decoders with 6 layers in each component.Translations should be similar, but not identical to output in the test set linked to in each model card. A.2) Group comments by id and split by sentence. **B)** AMR2-joint-ontowiki-seed42, Encoding takes pretrained AMR sentences, g1, and identifies the text to relate to the most appropriate sentence resting to the data, g2. Decoding vectors to construct AMR graphs with the transition-based neural parser (This is where we parse the data to show word-node alignments).The model diststills the sentences on Instagram comments to identify the 'silver' standard sentence using 'off the shelf" parsers and Maximum Bayes Smatch

Ensemble (MBSE). This process produces a single annotation, g3, by choosing from various state-of-the-art off-the-shelf parser outputs (A, B, C), for each input sentence. A, B, and C come from existing AMRs or their modified versions. MBSE is only applied to unlabeled English sentences to produce. **C**) The AMR output is taken and parsed into a human and machine readable tree and graph structures using Penman Notation. A string is parsed into a tree, and a tree is interpolated into a graph. Going from a graph to a tree is called configuration, and from a tree to a string is formatting. The interpolated graph is then used to identify the actor and objects they're directly connected to. This creates an actor network structure of the discourse.

D Appendix: Actor-Object Network

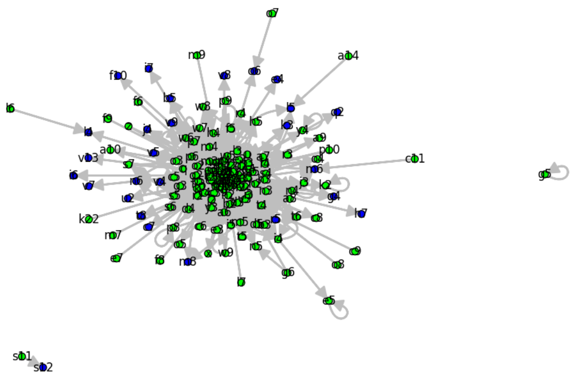

Figure 2: **Actors in Green, Objects in Blue.** Actor-Object Network where Actors are central in the network. The actors in green outnumber the objects in the network and refer to a type of action-based discourse communicated across Instagram. This seems expected as social movement discourse online is backed by calls to action, demands, and lists of grievances (Bennett and Segerberg, 2012). The objects in the network can be understood to support the actors in the network, potentially by supporting actions expressed in the comments.

E Appendix: Out-Degree Centrality

Entity & (Entity Label)	Out-Degree Centrality
we (W)	81
pain-01 (p)	77
positive (p2)	73
we (w2)	72
you (y)	70

Table 2: Top 10 Most Influential Entities