Knowledge Graph Construction from Online Discussions: Modeling Context-Dependent Opinions and Their Evolution

Ren Shimoide[1], Rikuto Sasaki[1] and Kazuhiro Takeuchi[1]

[1]Osaka Electro-Communication University, Japan

Abstract

Knowledge extraction from discussions differs fundamentally from conventional knowledge graph construction based on individual texts. In online discussions, multiple topic triples are distributed across short comments whose interpretation depends on context, while subjects, premises, and referents are frequently omitted. We propose an ontology consisting of two-layers: a Comment Discussion Graph and a Topic Assumption Graph. The former represents participants, comments, reply relations, existing discussion labels, and Argumentation Scheme Relations, while the latter represents multiple topic triples and three core relation types: support, rebuttal, and qualification, together with other relations observed in the analyzed discussions. Rather than assigning stable opinions directly to participants, the ontology models opinions as context-dependent attitudes expressed in comments toward topic triples. This design enables the evolution of expressed attitudes to be represented explicitly. Using a publicly available dataset of social and political YouTube discussions (comprising comments and replies), we apply four coarse-grained ASR categories inspired by Walton's argumentation schemes. These categories are assigned to pairs consisting of a parent comment and its child reply. A locally executable LLM pipeline, featuring reproducible prompts, inference settings, inputs, and structured outputs, uses ASRs as signals for reconstructing and normalizing topic triples distributed across comments. We instantiate the ontology for multiple cases with different discussion quality labels, including cases classified as Constructive Disagreement. The resulting representations are used to examine structural characteristics associated with constructiveness, including the evolution and refinement of topic structures and expressed attitudes, rather than mere agreement or convergence toward a single opinion. We further demonstrate the practical utility of the analysis through competency questions and SPARQL queries that retrieve instances of topic evolution and transitions in expressed attitudes. The proposed framework provides an explainable knowledge representation for analyzing constructiveness in online discussions.

Keywords

Knowledge Graph Construction, Ontology Engineering, Argument Mining

1. Introduction

Knowledge graph construction from text often assumes that the input is a self-contained document or sentence from which entities and relations can be extracted. Online discussions violate this assumption. A comment may omit its subject, reuse an expression introduced several turns earlier, answer only one part of a previous comment, or introduce a condition that changes the interpretation of an earlier statement. Consequently, extracting isolated triples from each comment can produce fragments that are locally plausible but insufficient for reconstructing the discussion, reflecting the broader need to integrate locally extracted knowledge into a coherent graph structure [1].

This problem becomes particularly important when the goal is to represent disagreement and discussion quality. Agreement or disagreement alone does not show how a discussion develops. A disagreement can add evidence, narrow the scope of a claim, introduce a counterexample, distinguish conditions, or propose an alternative policy. Conversely, a sequence can remain destructive even when it contains nominally relevant replies if it does not develop the underlying topics. We therefore focus

Proceedings of the 5th International Workshop on Natural Language Processing for Knowledge Graph Construction (NLP4KGC 2026), co-located with the 25th International Semantic Web Conference (ISWC 2026), Bari, Italy, October 26, 2026

mi25a003@oecu.jp (R. Shimoide); takeuchi@osakac.ac.jp (K. Takeuchi)

0009-0006-4307-1767 (K. Takeuchi)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-19

CEUR-WS.org/Vol-4335/paper6.pdf

on the evolution of topic triples across reply chains rather than assigning a single stable opinion to each participant.

We use the Old but Gold (ObG) dataset introduced by Bassi et al. [2], which contains social and political YouTube discussion chains with comment and reply annotations. Figure 1 presents the proposed two-layer representation and its construction process using an ObG discussion chain.

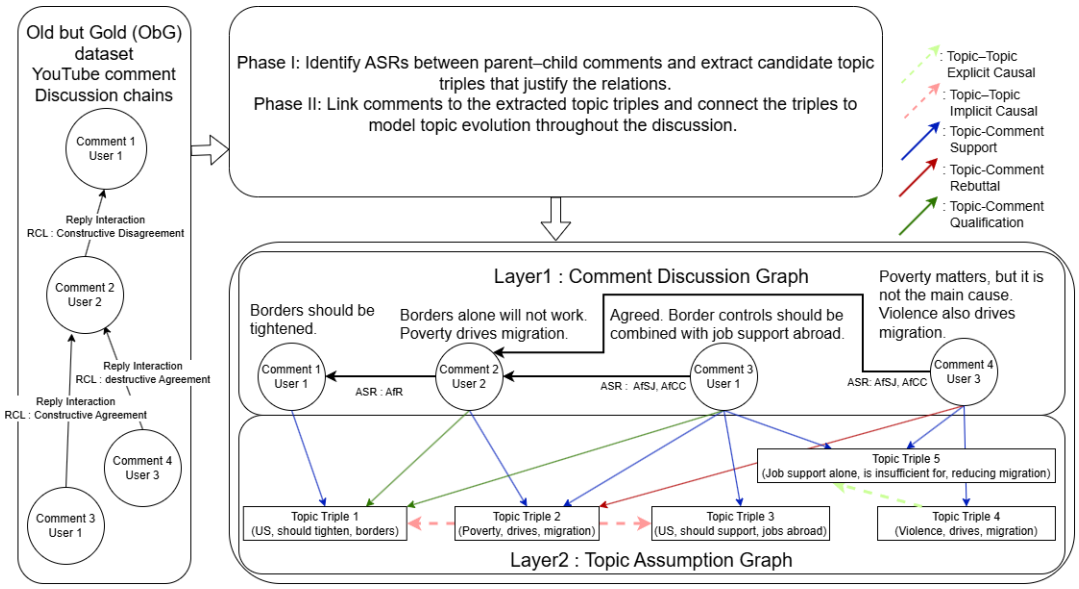


Figure 1: Overview of the proposed two-layer Discussion Knowledge Graph and its construction process. Black arrows indicate parent–child reply relations provided by the ObG dataset. In Phase I, ASRs are identified for parent–child comment pairs and retained in the Comment Discussion Graph. The identified ASRs are then used as signals to extract candidate topic triples. Explicit and implicit causal relations between the extracted topic triples are subsequently identified and are represented by green and red dashed arrows, respectively. In Phase II, comments are linked to the individual topic triples or relations between topic triples that they address. Blue, red, and green solid arrows indicate support, rebuttal, and qualification, respectively.

The two layers are connected through context-dependent attitudes expressed by comments. This design represents topic evolution and changes in expressed attitudes along a reply chain without assigning a fixed opinion to each participant. The main contribution of this paper is a two-layer knowledge representation that separates observable discussion structure from reconstructed topic structure and explicitly represents context-dependent attitudes expressed by comments toward topic triples and their relations. The proposed framework operationalizes this representation through the instantiation of actual discussion chains and formal analysis using a competency question and a SPARQL query.

2. Argumentative Relations and Discussion Topic Triples

Argument Mining is a research field that identifies argumentative components, such as claims and premises, and analyzes support and attack relations among them [3]. A prominent theoretical framework for such analysis is Walton's Argumentation Schemes, which categorize recurring patterns connecting premises and conclusions and support the analysis of inferential relations between utterances [4]. Existing approaches to argument analysis and stance detection typically define an issue or target in advance and assign labels such as favor, against, or neutral to individual utterances or participants [5]. ObG similarly uses the stance of each comment toward immigration and the agreement or disagreement between parent–child comment pairs [2]. However, in many informal discussions, including but not limited to social media, the particular aspect of an issue addressed by each comment is not necessarily explicit. Discourse relations may likewise require contextual reasoning beyond isolated utterances, including the interpretation of implicit connections and referential dependencies [6]. We therefore

use topic triples to represent the propositional content toward which comments express attitudes in their conversational context, and model how these topic triples and their relations evolve throughout a discussion.

To identify topics that are not explicit in individual comments, we use possible semantic relations between parent–child comments as contextual cues. For example, consider the parent comment A: "The city reduced the number of bus services," and its reply B: "Commuting became inconvenient." Comment B alone does not explicitly state what caused the inconvenience. We represent propositional content as a topic triple $t = (s, p, o)$, consisting of a subject, predicate, and object. From A and B, we construct the candidate topic triples $t_1 = (\text{the city}, \text{reduced}, \text{the number of bus services})$ and $t_2 = (\text{commuting}, \text{became}, \text{inconvenient})$, respectively. If a causal relation is hypothesized between the two comments, the hypothesis is represented as $h = \langle t_1, \text{causes}, t_2 \rangle$. We define four coarse-grained Argumentative Semantic Relations (ASRs), inspired by Walton's Argumentation Schemes and selected as contextual cues for reconstructing topic relations across related utterances in a discussion: Argument from Supportive Justification (AfSJ), Argument from Rebuttal (AfR), Argument from Causal Consequence (AfCC), and Argument from Analogy or Comparison (AfAC).

The resulting topic triples and their relations represent content omitted from individual comments and capture topic evolution throughout a discussion.

These definitions distinguish the observable reply structure from the topic structure reconstructed from the discussion. The former consists of participants, comments, and directed reply links, whereas the latter consists of context-dependent topic triples and relations between topic triples. The two structures are connected by the attitudes expressed by comments toward individual topic triples or relations between topic triples.

The ObG dataset contains social and political YouTube comment discussions with parent and child reply links, comment-level stance labels, and relation and quality annotations [2]. We use the original comment texts, reply structure, and existing annotations as input, while the ASRs, topic triples, relations between topic triples, and links from comments to topic triples or their relations are generated by our pipeline.

3. Proposed Method

LLMs have been applied to few-shot knowledge extraction, ontology learning, and the automation of ontology and knowledge graph construction [7, 8, 9]. Multi-stage LLM procedures have been proposed for causal relation extraction, where an initial judgment is subsequently examined by another model to improve reliability [10]. Candidate generation followed by semantic filtering has also been explored to improve the transparency and verifiability of LLM-assisted knowledge extraction [11].

To instantiate the proposed representation from discussion data, we implement a two-phase pipeline using locally executable LLMs. The two phases specify what graph structures are constructed, whereas a three-stage decision procedure specifies how each LLM-based judgment is generated, reviewed, and finalized. For each task, Stage 1 generates or selects an initial candidate under task-specific constraints. Stage 2 examines the initial candidate against the original input and the task definitions, and either retains or revises it. Stage 3 produces the final judgment by considering the original input together with the outputs of the preceding stages. Stages 1–3 use gemma4:26b, qwen3.6:27b, and gpt-oss:120b, respectively; reasoning is disabled for Stages 1–2 and enabled for Stage 3. Prompts, inference settings, inputs, intermediate judgments, and final structured outputs are recorded for reproducibility.

3.1. Phase I: ASR Identification and Topic Structure Reconstruction

Phase I proceeds in three steps. First, ASRs are identified for each parent–child comment pair. Second, the identified ASRs are used as signals, and each parent–child comment pair is provided as input to extract candidate topic triples. Third, relations between candidate topic triples are classified as explicit causal, implicit causal, or unrelated, and only the first two are retained. Figure 2 illustrates this process for `discussion_32291eaac8`.

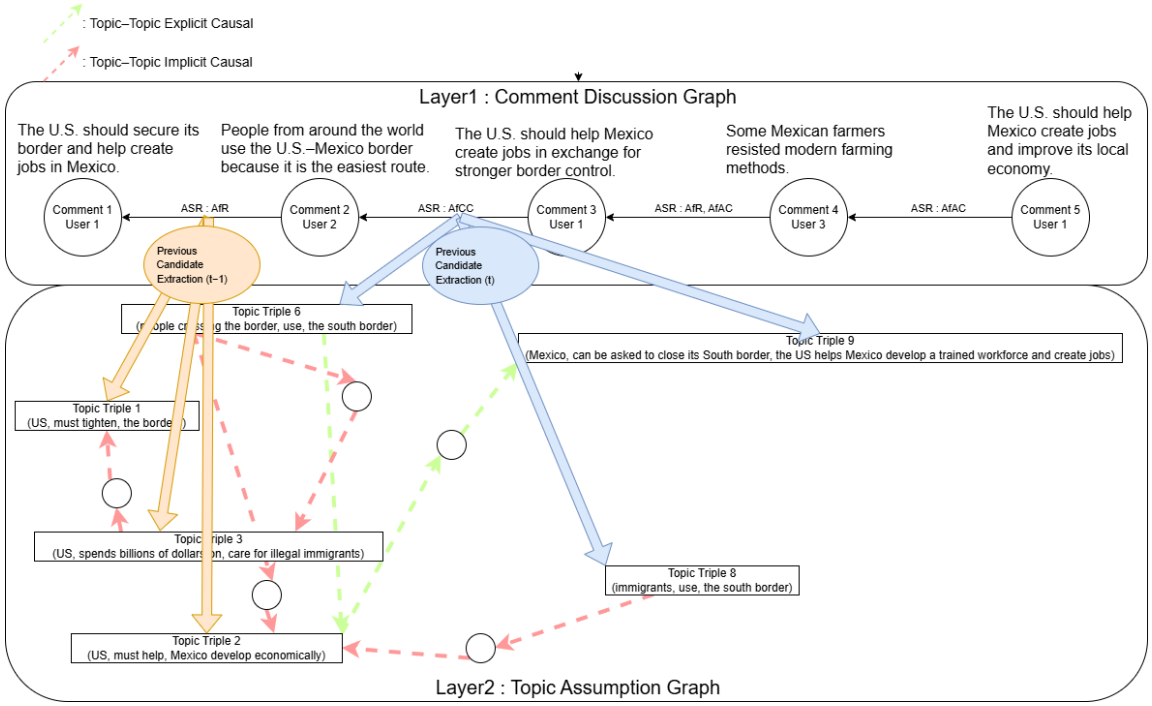


Figure 2: Parent–child comment pairs are processed sequentially along the reply chain, and ASRs are identified for each pair. The identified ASRs are used as contextual signals to extract and normalize candidate topic triples. Relations between the extracted topic triples are then classified as explicit causal, implicit causal, or unrelated. Green and red dashed arrows indicate the retained explicit and implicit causal relations, respectively.

3.2. Phase II: Comment to Topic linking

Phase II links comments to individual topic triples and relations between topic triples reconstructed in Phase I. Figure 3 illustrates this process for `discussion_32291eaac8`.

Phase II proceeds in two steps. First, the discussion is traversed from the root comment toward its descendants, and the relation of each comment to individual topic triples and relations between topic triples is identified. Second, for each identified target, the expressed attitude is classified as support, rebuttal, or qualification. Treating relations between topic triples as addressable targets allows the method to distinguish rejection or qualification of a connection from rejection of either topic triple independently. The resulting cross-layer links connect the Comment Discussion Graph with the Topic Assumption Graph.

4. Case Studies on Discussion Chains

The two chains presented here demonstrate the proposed representation and its queryability in detail; additional discussion instances generated by the same pipeline are released as machine-readable JSON in our repository[1], providing inspectable evidence that the same pipeline has been used to instantiate the proposed representation beyond these two cases. The first chain, presented in Section 3, illustrates the construction of the topic structure and its cross-layer links through Phases I and II. The second chain is instantiated as a completed Discussion Knowledge Graph, over which a SPARQL query is executed to retrieve changes in expressed attitudes toward the same topic triple along a reply path. Figure 4 presents the graph constructed for `discussion_0f94a5abc0`. The discussion concerns border enforcement and the respective responsibilities of Texas and the federal government. The graph separates topic triples distributed across the comments and represents explicit and implicit causal relations between them, together with the support, rebuttal, and qualification expressed by individual comments.

[1] https://github.com/RenShimoide/discussion-knowledge-graph-data

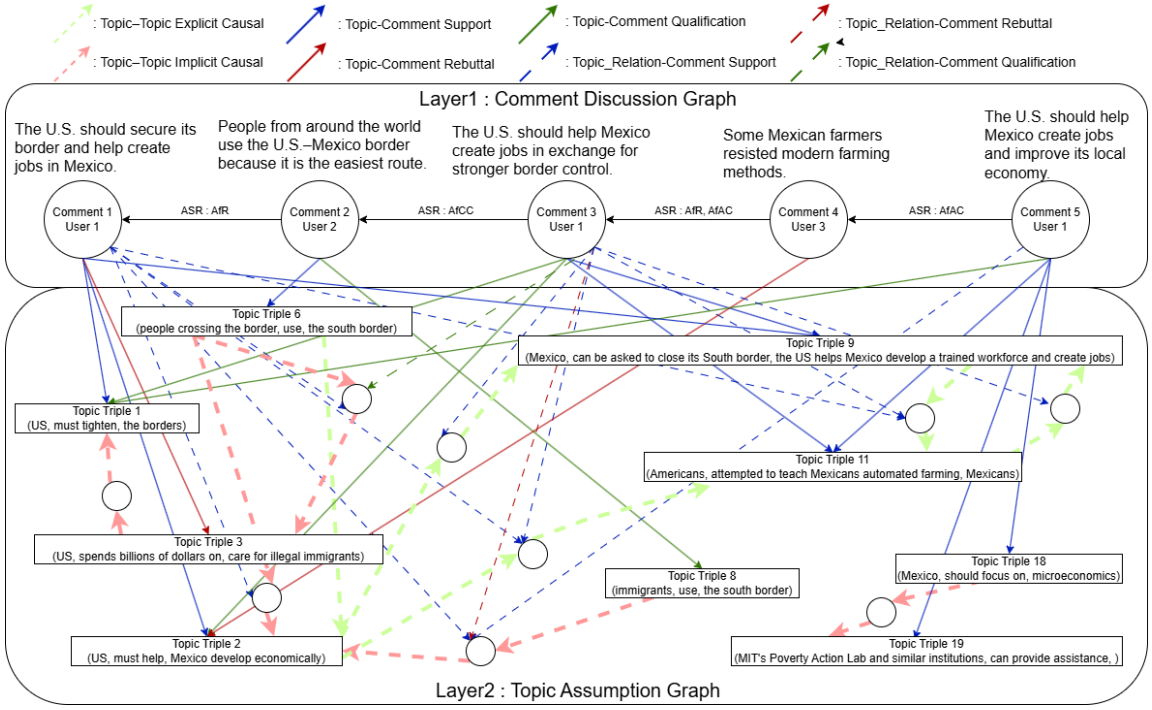


Figure 3: Phase II for `discussion_32291eaac8`. Comments are linked to the individual topic triples or relations between topic triples that they address. Solid arrows indicate links to individual topic triples, whereas dashed arrows indicate links to relations between topic triples. Blue, red, and green arrows represent support, rebuttal, and qualification, respectively.

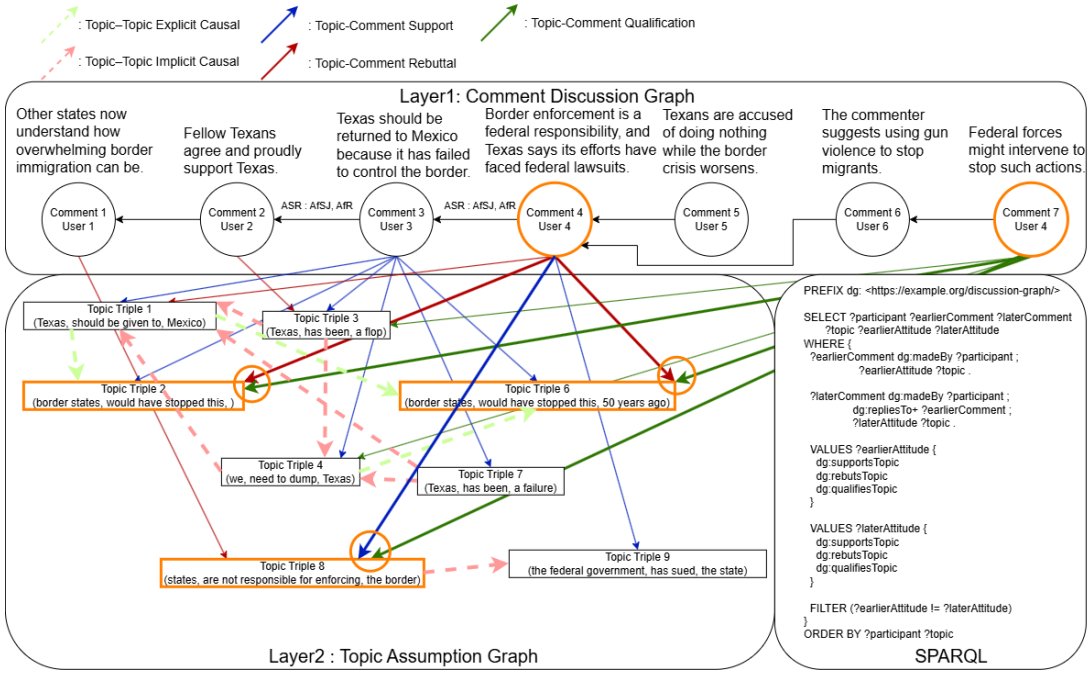


Figure 4: Illustrative Discussion Knowledge Graph for `discussion_0f94a5abc0` with the results of a SPARQL query for changes in expressed attitudes. Orange outlines highlight the earlier and later comments and the shared topic triples included in the query results. Thickened links indicate the attitude relations participating in the returned bindings, and each orange circle encloses a pair of links with different predicates, marking a detected change in expressed attitude.

To demonstrate queryability, we consider the competency question: *Which participants express different attitudes toward the same topic triple at different points along a reply path?* The SPARQL query retrieves pairs of comments by the same participant that occur along a reply path, are linked to the same topic triple, and use different attitude predicates.

The example shows that the proposed representation captures changes in comment-specific attitudes without assigning a single stable opinion to each participant. It also supports queries concerning rebutted causal relations, qualification of topic triples, and refinement of the topic structure across a discussion.

5. Conclusion

This study presents an integrated framework for analyzing online discussions. It combines a two-layer knowledge representation model, the instantiation of actual discussion chains as Discussion Knowledge Graphs, and the formulation of analytical tasks as SPARQL queries. The case studies make explicit several aspects that cannot be directly captured by the ObG discussion quality labels or comment-level agreement and disagreement alone, including the evolution and refinement of the topic structure, rebuttal and qualification directed at relations between topic triples, and context-dependent expressed attitudes. The SPARQL analysis further shows how changes in attitudes toward the same topic triple can be retrieved along a reply path. These results demonstrate the feasibility and queryability of the proposed representation; quantitative evaluation against human annotations and baseline methods remains future work.

Declaration on Generative AI

Generative AI tools were used to translate text originally drafted by the authors in their native language into English, and for linguistic polishing and proofreading. The authors reviewed the final manuscript and take full responsibility for its content.

References

[1] R. Yang et al., Graphusion: A RAG Framework for Scientific Knowledge Graph Construction with a Global Perspective, Companion Proc. of the ACM Web Conference, 2025, pp. 2579–2588.

[2] D. Bassi et al., Old but Gold: LLM-Based Features and Shallow Learning Methods for Fine-Grained Controversy Analysis in YouTube Comments, Proc. of the 12th Argument Mining Workshop, 2025, pp. 46–57.

[3] J. Lawrence, C. Reed, Argument Mining: A Survey, Computational Linguistics, 2020, pp. 765–818.

[4] D. Walton, C. Reed, F. Macagno, Argumentation Schemes, Cambridge University Press, 2008.

[5] D. Küçük, F. Can, Stance Detection: A Survey, ACM Computing Surveys, 2020, pp. 1–37.

[6] Y. Miao, M.-Y. Kan, Discursive Circuits: How Do Language Models Understand Discourse Relations?, EMNLP, 2025, pp. 32558–32577.

[7] A. M. P. Braşoveanu, L. J. B. Nixon, A. Weichselbraun, A. Scharl, Framing Few-Shot Knowledge Graph Completion with Large Language Models, NLP4KGC, 2023, 13 pages.

[8] R. M. Bakker, D. L. Di Scala, M. H. T. de Boer, Ontology Learning from Text: An Analysis on LLM Performance, NLP4KGC, 2024, pp. 70–87.

[9] V. K. Kommineni, B. König-Ries, S. Samuel, Towards the Automation of Knowledge Graph Construction Using Large Language Models, NLP4KGC, 2024, pp. 19–34.

[10] T. Ugai, Causal Relationship Extraction Using AI Multi-Agents and Large Language Models for Causal Knowledge Graph Generation, Special Interest Group Technical Report of JSAI, 2025, 10 pages, (in Japanese).

[11] R. Sasaki et al., SynSem-Align: Ontology-Driven KG Extraction via Syntactic Candidate Mining and Paraphrase-Based Equivalence Filtering, ISWC 2025 Companion, 2025, 7 pages.