Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Grounding Political Discourse in the Body: A Computational Approach to Analyzing U.S. Energy Narratives

Xiaoyu Li
National University of Defense Technology
Lei Gu
National University of Defense Technology
Haohan Meng
alan-mhh@qq.com

National University of Defense Technology

Article

Keywords: Embodied language, image schema, embodied simulation, semantic role labeling, narrative clustering

Posted Date: November 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10241267/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Image schemas, recurring pre-conceptual gestalts arising from sensorimotor experience, constitute the foundational level of embodied conceptual architecture, yet their systematic identification in naturalistic language data remains a persistent methodological challenge. Traditional introspective analyses in cognitive linguistics yield rich theoretical insights but face difficulties in scalability, inter-coder reliability, and systematic replication. This study addresses these challenges by developing a computationally assisted analytical framework that integrates semantic role labeling (AllenNLP) with large language model verification (DeepSeek R1) to detect image-schematic activations in discourse. Grounded in Fairclough's three-dimensional critical discourse analysis model, the pipeline extracts narrative triples ("agent–action–patient") from 10 U.S. energy policy documents (74,426 characters, 289 sentences), clusters them into conceptual domains via SentenceBERT embeddings and DBSCAN, and subjects the outputs to human–LLM collaborative validation. Against a manually annotated gold standard (50 triples from 29 sentences), the pipeline achieves an F1 score of 0.79, a 20-percentage-point improvement over the raw SRL baseline. Three image schemas (CONTAINER, PATH–GOAL, and FORCE) are identified as systematically activated through specific discourse strategies. By providing structured SRL output as input to the LLM rather than requesting annotation from scratch, the pipeline reduces hallucination while preserving the directional semantic relations needed for schematic analysis. This study offers a replicable, scalable methodology for investigating embodied language in naturally occurring discourse.

1 Introduction

Cognitive linguistics, grounded in embodied philosophy, holds that human language and thought are fundamentally shaped by bodily experience. Conceptual systems arise from recurring sensorimotor patterns encountered in everyday interaction with the physical and social world. At the foundation of this embodied conceptual architecture lie image schemas, recurring pre-conceptual gestalts such as CONTAINER, PATH–GOAL, and FORCE, that emerge from bodily engagement with the environment and structure more complex cognitive representations (Johnson, 1987). Conceptual metaphor theory further demonstrates how abstract domains are systematically understood through mappings from concrete, bodily-grounded source domains (Lakoff and Johnson, 1980). Together, these embodied mechanisms shape not only basic cognition but also higher-order reasoning and public communication. Sensorimotor experience provides the grounding for language comprehension and production, including the interpretation of complex social and political issues (Pelkey, 2023; Dove, 2023; Zinovyeva & Trapeznikov, 2025).

In political discourse, speakers routinely draw upon these embodied foundations to frame policy challenges in ways that resonate with audiences' sensorimotor understanding. Abstract policy concepts are mapped onto image schemas, transforming regulatory and geopolitical arguments into intuitively accessible narratives. Within U.S. energy policy discourse under the Trump administration's second term, the administration's "energy dominance" framing systematically recruits the CONTAINER schema to encode energy dependence as boundary violation, the PATH–GOAL schema to project a trajectory from

crisis to national renewal, and the FORCE schema in its “removal of restraint” variant to construe deregulation as liberation from constraints. These schematic activations render abstract policy positions viscerally persuasive, aligning policy arguments with bodily experiences of containment, purposeful movement, and the release of obstruction.

Traditional approaches to cognitive linguistic analysis of political discourse have relied on manual annotation conducted by trained analysts. While such approaches yield rich interpretive insights, they encounter persistent challenges in scalability, inter-coder reliability, and systematic replication when applied to large-scale policy corpora. Quantitative alternatives, such as topic modeling (e.g., LDA, BERTopic), offer broader coverage but fragment discourse into bags of words, while obscuring the directed semantic relations of “who does what to whom” that carry schematic content. With the rapid development of generative artificial intelligence, large language models now demonstrate their efficiency and interpretative power in political discourse analysis (Meizlish & Ziffo, 2026; Mochtak, 2025). The present study addresses these dual limitations by developing a LLM-assisted analytical pipeline grounded in Fairclough’s three-dimensional critical discourse analysis framework. Semantic role labeling is employed to extract narrative triples from policy sentences, preserving the directional structure of semantic relations. These triples are then encoded into dense vector representations and clustered to identify recurrent conceptual domains. A large language model (DeepSeek R1) is integrated into the workflow through structured prompt engineering for schema-informed annotation of discourse strategies. This pipeline enables systematic extraction of narrative structures while maintaining the interpretive depth necessary for identifying image-schematic activation across textual, discursive, and social levels of analysis.

This study makes three contributions to the field of embodied language research. First, it responds to calls for methodological innovation in studying language embodiment by presenting a replicable, computationally assisted pipeline that bridges natural language processing and cognitive linguistic theory. Unlike purely introspective approaches, the pipeline produces structured, verifiable outputs that can be evaluated against gold-standard annotations and replicated across corpora. Second, it advances understanding of how image schemas operate at multiple levels of discursive organization, providing systematic evidence that abstract political thought is grounded in sensorimotor experience. Third, it shows how a cross-disciplinary framework in which computational tools augment human interpretive analysis can address long-standing concerns about reliability and scalability in cognitive linguistic research (Meng et al., 2025; Fuchs, 2026). By constraining the LLM’s task to structured verification rather than open-ended annotation, the pipeline reduces the risk of model-generated confabulation while preserving the directional semantic relations needed for identifying embodied cognitive structures in discourse.

2 Literature Review

2.1 Energy Security Narratives as Social Practices

Discourse, in the field of Critical Discourse Analysis (CDA), is viewed as a form of social practice, shaping and being shaped by social reality (Fairclough and Fairclough, 2012). Within this tradition, narrative is understood as a primary mode through which discourse organizes meaning. As for energy politics, energy security narratives constitute a particularly consequential type of this discursive practice. Previous studies have adopted linguistic approaches, such as discourse-historical analysis (Ma and Zhao, 2025) and computational text analysis (Ba and Tan, 2025), to examine issues of energy transition, technological innovation, and climate change (Liu, 2024). This line of research has demonstrated how energy security narratives function simultaneously as products of existing power relations and as productive forces that construct cognitive frameworks for energy issues, shape public attitudes, and advance specific policy agendas.

As the world's largest producer and consumer of oil and gas, the United States faces energy security challenges inherent in its large-scale import-export energy market. Overcoming internal and external resistance to energy transition policies often requires leveraging energy security narratives as a discursive vanguard, with narrative discourse and legislative policy practice closely intertwined (Dunham and Schlosser, 2016), making this a persistent focus of scholarly attention. In recent years, U.S. media discourse on energy has frequently revolved around themes of political polarization, shifting policy priorities, and geopolitical competition, drawing attention to the challenges facing the U.S. in pursuing energy independence (Chen and Liu, 2025). During the Trump 1.0 era, the federal government's energy security discourse showed a clear preference for prioritizing fossil fuel development, calling for infrastructure construction and technological innovation as strategic means, and employment and security as strategic objectives, reflecting a strong ideology of energy realism and energy hegemonism (Zhao and Zhao, 2021). After his reelection, the securitization tendency in the energy narratives of Trump and his communication team has only intensified, becoming a "discursive weapon" for advancing his new energy policies.

2.2 Embodied-Cognitive Interpretation of Political Discourse

A growing body of research has examined political discourse through the lens of cognitive linguistics, revealing how embodied cognitive structures (e.g. image schemas and conceptual metaphors) systematically shape public understanding and ideological positioning. Image schemas, usually understood as recurring, pre-conceptual patterns that are derived from bodily experience (Johnson, 1987), operate at the foundational level of political conceptualization, providing the cognitive scaffolding through which abstract policy issues become intuitively accessible. Grammatical structures in political texts systematically activate image schemas to encode ideological positions (Hart, 2014), and even extends beyond linguistic modes into multimodal political discourse, where visual and spatial cues further reinforce schematic activation (Hart, 2024). A related line of research has examined the intrinsic embodied-cognitive clue of how political actors produce, distribute and validate conceptual metaphors. Political discourse is organized along three spatial axes, namely space, time, and modality (Chilton, 2004), through which metaphors behave as narrative "scenarios" rooted in embodied schemas (Musolff,

2016). The range of conceptual source domains available for political metaphor draws upon embodied experience, which exhibits both universal cognitive patterns and culturally specific variations. Bodily experience, human interaction with the physical environment, and cultural context jointly shape which metaphoric mappings become conventionalized in political discourse (Kövecses, 2015).

2.3 Narrative Semantic Mining Based on Semantic Role Labeling

In general, narrative semantic mining is mainly based on keywords (Lazard et al., 2015), topics (Nanyonga and Wild, 2025), or triples (Zhao et al., 2024). Among these, the triple, i.e., "agent–action–patient" (ARG0-V-ARG1), can serve as a narrative unit, revealing the directed process of a narrative through semantic interaction links. In recent years, semantic role labeling (SRL) technology, through which triples are identified (Chen et al., 2025), has continued to evolve to mine narrative construction patterns in tweets (Jing, 2021), blogs (Tangherlini et al., 2020), news (Shahsavari et al., 2020) speeches (Ash et al., 2024), and other types of texts. However, when processing domain-specific texts with complex syntactic structures, such as legal and security documents, the accuracy of these general-purpose SRL models declines. Consequently, some researchers have attempted to fine-tune domain models (e.g., LEGALBERT) on legal texts (Bakker, 2025) or improve the accuracy and stability of SRL output through cross-domain fusion, such as introducing large language models and curriculum learning frameworks, or combining SRL based on NLP tools with other techniques like bootstrapping (Chen et al., 2023) and change point detection (Zhao et al., 2024), thereby automating threat intelligence mining and tracking the dynamic evolution of collective narratives. Moreover, with the rapid advancement of large language models (LLMs), recent previous research has shown that zero-shot or few-shot LLMs can be leveraged for semantic role labeling (SRL) tasks, thereby reducing dependency on large-scale annotated corpora in low-resource languages (Senator et al., 2025; Smirnov & Larionov, 2025); however, their performance on legal texts remains suboptimal (Bakker, 2025).

Overall, previous studies have predominantly relied on qualitative, interpretive methods or simple corpus statistics. When applied to large-scale policy corpora, these approaches face challenges of scalability, inter-coder consistency, and systematic consistency. This limitation motivates the integration of computational methods with political discourse analysis, and thus manage to balance structured extraction from larger corpora with embodied deconstruction. This study attempts to advance: first, by building a human-AI collaborated workflow to conduct automated and systematic narrative semantic mining, so as to improve the efficiency of semantic mining of policy texts; second, by integrating narrative semantic mining with critical discourse analysis to combine quantitative extraction with qualitative interpretation, offering an interdisciplinary perspective for systematically uncovering clue of embodied experience behind the U.S. energy narrative.

3 Materials and Methods

3.1 Research Questions

Grounded in the embodied cognition framework outlined above, this study addresses three interrelated research questions:

(1) Which image schemas are systematically activated in the Trump administration's "energy dominance" policy discourse?

(2) Through what specific lexical-grammatical mechanisms are these schemas encoded in policy sentences?

(3) How do these schemas combine into higher-order conceptual integration strategies that jointly construct the cognitive architecture of the political narrative?

3.2 Model

3.2.1 AllenNLP

AllenNLP is an open-source NLP research library built on PyTorch, designed for high-level semantic tasks including semantic role labeling, textual entailment, and coreference resolution (Wallace et al., 2019). This study employs AllenNLP's pretrained BERT-based SRL model, trained on the OntoNotes 5.0 corpus with PropBank-style semantic role annotations. The model takes raw sentences as input and outputs predicate–argument structures in which each predicate is associated with labeled arguments such as ARG0 (agent), ARG1 (patient), and ARGM- (modifiers). Its robust out-of-the-box performance on general-domain texts and accessibility through a well-documented Python API make it suitable for integration into the computational pipeline developed here.

3.2.2 DeepSeek R1

DeepSeek-R1 possesses high-precision annotation proficiency for complex tasks and excels in automated corpus labeling, rendering it applicable to multifaceted annotation paradigms including metaphor and narrative annotation. Therefore, DeepSeek was selected as the experimental model in this study, invoked via APIs with default parameters uniformly adopted.

3.3 Methodology

3.3.1 Semantic Role Labelling

Semantic role labeling (SRL) identifies core arguments such as the agent (ARG0) and patient (ARG1) of a predicate (typically a verb) in a sentence, transforming unstructured text into triples of the form ARG0–V–ARG1. These triples serve as narrative units that preserve the directed semantic relation of "who did what to whom," providing a formal representation amenable to systematic schema detection. This study employs the pretrained BERT-based SRL model from the AllenNLP framework to automatically extract triples from policy sentences.

3.3.2 Narrative Clustering

Narrative clustering groups semantically similar triples to identify recurrent patterns across the corpus. Triples are encoded into dense vector representations using SentenceBERT (all-mpnet-base-v2), and the DBSCAN algorithm (Density-Based Spatial Clustering of Applications with Noise) partitions the semantic space. This step operates at both the sentence level and the argument level, providing the empirical basis for the identification of recurrent conceptual domains. At the sentence level, all triples within each cluster are treated as semantically similar narrative units, and the most frequent triple functions as the cluster representative. At the argument level, entities are clustered, and the entity closest to the cluster centroid is selected as the representative entity for each group.

3.3.3 Prompt Engineering

In the human–LLM collaborative component of the textual pipeline, structured prompts are used to guide DeepSeek-R1 in validating and correcting SRL outputs. The prompt is designed according to six principles. (1) Role setting: the LLM is instructed to assume the role of a “policy text analysis expert” to anchor its interpretive frame. (2) Concept definition: drawing on frame semantics theory, explicit definitions of “semantic role labeling” and “triple” are provided to ground the model’s understanding. (3) Task decomposition: the verification task is broken down into three subtasks (correctness judgment, error correction, and role completion), each requiring specific output. (4) Few-shot examples: two representative cases, encompassing one correct and one incorrect triple with corresponding corrections, are supplied to calibrate the model’s expectations. (5) Chain-of-thought: the LLM is required to articulate its reasoning (identifying the main predicate, determining syntactic structure, and explaining argument assignment) before producing the final output. (6) Output format constraint: a JSON schema is mandated, specifying fields for original triples, correctness judgments, corrected triples, reasoning traces, and completion notes. This structured approach draws on chain-of-thought prompting methodology to ensure consistency across large-batch validation and to reduce the variability associated with open-ended LLM interactions. The full prompt template is provided in the Appendix in Supplementary Material.

3.4 Research Framework

This study adopts a mixed-methods approach. At the textual level, an integrated computational pipeline combining AllenNLP’s SRL model with DeepSeek-R1 performs automated triple extraction, clustering, and LLM-assisted validation. At the discursive and social levels, these computational outputs serve as the evidentiary basis for manual qualitative analysis grounded in Fairclough’s three-dimensional critical discourse analysis framework. The research framework is shown below (Fig. 1).

3.4.1 Data Collection and Preprocessing

This study uses the White House website as the sole retrieval platform, setting keywords such as “energy dominance” and “energy policy,” with the retrieval period from January 20, 2025 (Trump’s second-term inauguration) to February 24, 2026. Manual screening was applied to exclude texts that only mention energy without involving policy design or strategic narratives, giving priority to executive

orders, presidential proclamations, and White House statements, finally selecting 10 texts totaling 75,716 characters. Then, manual cleaning (removing file headers, separators, headers/footers, extra blank lines), programmatic processing (removing list markers, processing list structures, splitting natural sentences to meet SRL requirements), and manual verification, resulting in 289 sentences with a total of 74,426 characters.

3.4.2 Semantic Role Labelling and LLM-Assisted Verification

Using the Python 3.11 environment, the AllenNLP 2.10.1 library was configured via Anaconda Prompt. The pretrained BERT-based SRL model was applied to automatically label each sentence and extract ARG0–V–ARG1 triples. Although the SRL model captures the basic predicate–argument structure, passive voice constructions, nested clauses, and long-distance dependencies in policy texts frequently produce labeling errors, most notably the omission or misattribution of the agent role (ARG0) and the misidentification of noise predicates such as participial and adjectival forms. To address these issues, all SRL outputs were validated through the human–LLM collaborative workflow described in Section 3.3.3. DeepSeek-R1 processed the prompt template for each sentence, generating a structured JSON output that included correctness judgments, error-type classifications, and corrected triples where applicable. For passive or agent-omitted constructions, the LLM inferred the implicit ARG0 when recoverable from context and marked unrecoverable agents as [UNKNOWN]. A Python script performed post-processing denoising, and a human analyst conducted substantive content review to ensure semantic fidelity.

3.4.3 Narrative Entity Clustering

First, verb normalization was performed using the VerbAtlas framework. Since the SRL model outputs predicates in PropBank format, where the same event can be expressed by dozens of different verb forms, direct clustering would disperse them into different clusters. Therefore, the NLTK library was used to lemmatize verbs, and then the VerbAtlas framework was used to map them, normalizing similar predicates to ensure clustering validity. The original semantic units contained 268 verbs, of which 263 were mapped (98.1% accuracy), yielding 122 VerbAtlas semantic frames.

Second, each triple was converted into a sentence and Sentence-BERT embeddings were generated. Narrative units were simplified into triples to capture core semantics and reduce interference from irrelevant details. Thus, triples were recomposed into sentences, and the Sentence-BERT model all-mpnet-base-v2 was used to generate vector representations for each sentence.

Third, the DBSCAN clustering algorithm was applied to arguments. From each cluster, the entity with the shortest distance to the cluster center was selected as the representative entity, collectively presenting the most characteristic information of each narrative cluster with domain coverage and representativeness.

3.4.4 Discourse Strategy Analysis

The identification and interpretation of discourse strategies at the discursive and social levels are conducted through manual qualitative analysis. This methodological choice reflects the nature of the analytical task: whereas the textual level involves pattern recognition across a large number of units to which computational methods are well suited; while the discursive and social levels require contextualized interpretation, domain-specific knowledge, and theoretical reasoning that exceed the current capabilities of automated systems. The analysis proceeds through three stages, grounded in Fairclough's (1992) three-dimensional framework and informed by securitization theory (Buzan et al., 1998; Stivas, 2024).

In the first stage, two analysts with expertise in cognitive linguistics and political discourse independently examined the clustering outputs. Each analyst annotated recurring patterns of language use that function as discourse strategies, defined as systematic ways in which language constructs social identities, social relations, and systems of knowledge and belief (Fairclough, 1992). Annotations were required to be grounded in specific textual evidence from the clustering results and to specify, wherever applicable, the image-schematic patterns (CONTAINER, PATH–GOAL, FORCE) associated with each strategy.

In the second stage, the two analysts compared their independent annotations. Areas of agreement were consolidated; discrepancies were resolved through structured discussion and re-examination of the source texts. This process yielded a consensus typology of three discourse strategies (threat construction, adversarial polarization, and restoration narrative), each defined in terms of its discursive function and its cognitive-linguistic substrate.

In the third stage, the consensus annotations were mapped onto the three-dimensional framework. The textual level was accounted for by the image-schematic profiles of the discourse strategies. The discursive level was accounted for by the ways in which these strategies combined to produce a coherent political narrative, integrating the four conceptual domains (Administrative Entities, Policy Actions, Energy Resources, and National Security) into a unified thematic structure. The social level was accounted for by situating these discursive operations within the genre conventions of executive orders, the political-economic conditions of the post-2025 conjuncture, and the embodied-cognitive foundations that enable their persuasive force.

4 Results

For evaluating the results, we post-processed the generated outputs, converting the JSON output into a list of labels per word to ensure alignment with the required format. We then selected 10% of the sentences (29 sentences) from the full text as a test set. Two independent coders manually annotated these sentences with (ARG0-V-ARG1) triples following a predefined annotation scheme: red for agents (ARG0), yellow for actions (V), and cyan for patients (ARG1). Sentences marked in gray were excluded from annotation, as they contained only nominal or descriptive statements without transitive event structures. Inter-annotator agreement was measured using Fleiss' kappa. A domain expert adjudicated

sentences where annotators disagreed, retaining one unique annotation per sentence as the gold standard. The same 29-sentence test split was used to evaluate both the raw AllenNLP output and the LLM-assisted output, thereby providing a consistent basis for measuring the effectiveness of our approach. We then used precision, recall, and F1 scores to evaluate the performance of the models. For argument matching, we adopted a lenient criterion: a predicted triple was considered a true positive if the verb matched and the core content nouns of both ARG0 and ARG1 overlapped with the gold standard, without requiring exact boundary alignment.

The annotations on the sentences yielded an inter-annotator agreement of $\kappa \approx 0.84$, exceeding the 0.75 threshold, thereby confirming the reliability of the manual annotation. After expert adjudication, the resolved annotations were used as the gold standard.

4.1 Results of Semantic Role Labeling

To evaluate the effectiveness of the prompt-driven verification pipeline, we present a representative example of LLM's output from the verification process. The following example illustrates how the LLM validates correct SRL extractions (Fig. 2) .

As this example illustrates, the AllenNLP SRL model incorrectly identifies the adjective "pronounced," which does not encode a substantive action, as a verb. Moreover, the resulting triple is incomplete, as it lacks ARG1. The LLM effectively corrects this error, filtering out the noise predicate and retaining only triples in which both ARG0 and ARG1 are present and the verb encodes a genuine transitive event.

To systematically assess the pipeline, 29 sentences were manually annotated as a gold standard comprising 50 (ARG0-V-ARG1) triples. Both the raw AllenNLP output and the LLM-assisted output were evaluated against this gold standard using triple-level precision, recall, and F1 scores. The results are presented in Table 1.

Table 1
Evaluation results of SRL triple extraction

Model	TP	FP	FN	Precision	Recall	F1
Raw AllenNLP	38	41	12	0.481	0.760	0.589
AllenNLP + DeepSeek	42	14	8	0.750	0.840	0.7925

Note: The evaluation counts all predicted triples, including those for sentences outside the gold-standard annotation scope. The gold standard comprises 50 (ARG0-V-ARG1) triples from 29 manually annotated sentences. A predicted triple was considered a true positive if the verb matched and the core content nouns of both ARG0 and ARG1 overlapped with the gold standard (only correct role type assignment is required without precise boundary matching).

As Table 1 shows, the LLM-assisted pipeline yields a substantial improvement over the raw AllenNLP baseline. The F1 score rises from 0.59 to 0.79, an increase of 20 percentage points. This improvement is driven primarily by a sharp reduction in false positives (from 41 to 14), which in turn raises precision

from 0.48 to 0.75. Recall also improves, from 0.76 to 0.84, as the LLM recovers several agent roles that the raw SRL model missed and corrects misattributed agents.

The 41 false positives in the raw AllenNLP output can be traced to two main sources. First, AllenNLP frequently treats participial and adjectival forms as main verbs, producing noise triples from non-core predicates such as living, fixed, forgotten, growing, pronounced, including, authorized, and appropriated. These forms serve a modifying or descriptive function rather than encoding transitive events, yet AllenNLP assigns them argument structures. Second, AllenNLP tends to truncate or misassign the agent role (ARG0): in passive constructions, it either leaves ARG0 empty or extracts only a fragment (e.g., “our Nation’s” instead of “our Nation’s diminished capacity”), while in active constructions it occasionally confuses the instrument with the agent (e.g., labeling “emergency fuel waivers” as ARG0 of allow instead of “the Administrator”). The LLM correction step addresses both issues: it filters out noise verbs by evaluating whether a predicate encodes a semantically complete transitive event, and it repairs truncated or misassigned agent arguments through contextual inference.

4.2 Micro-Level Analysis: Argument Clustering

The clustering algorithm can extract entities and perform entity clustering, but due to the large number and diversity of entities, they still need to be grouped and named according to their meanings to clearly present narrative focal points (Jia et al., 2026). After clustering entity vectors, 34 valid clusters were obtained, such as “American energy,” “our Nation’s coal industry,” etc. After manual screening and grouping, the cluster groups and entities are presented as Table 2.

Table 2
Argument clustering results

Cluster Group	Entities
Administrative Entities	Agencies (48); Council (5); Secretary (4); the U.S (3); States (5); possible authorities (4); authority granted by law (3); the Nation (3); Trump (4); OMB Director (6); Director of the Fish and Wildlife Service (5); agency does not have agency actions (3)
Policy Actions	appropriate action (3); policies or regulations (3); enforcement discretion (4); such actions (3); assessment (3); State laws (7); policies concerning the “beginning of construction” (3); agency actions identified as unduly burdensome (3); repeal or modifications of green energy tax credits (3); guidance, regulations, programs, and policies (3)
Energy Resources	our grid (37); coal (31); year - round sale of E15 gasoline (4); National Defense Stockpile (3); our electric grid upon which defense-industrial production depend (3); Nation’s energy supply that may be subject to emergency treatment (4); adverse environmental consequences (3); energy development, siting, or production activities on Federal land (3); cost of energy (3)
National Security	national security (7); a threat to national security (3); foreign adversaries (3)

After clustering entity vectors, 34 valid clusters were obtained and manually grouped into four cluster groups according to semantic affinity.

Administrative Entities encompass federal agencies, the President, the OMB Director, state governments, and other institutional actors. These entities function primarily as syntactic subjects (ARG0) in the extracted triples and carry limited independent semantic weight. Their prevalence reflects the genre of executive orders, in which every action must be attributed to a statutorily authorized actor.

Policy Actions include both broad policy concepts and specific regulatory instruments: enforcement discretion, repeal of green energy tax credits, and guidance documents targeting previous administration policies. These entities represent the instrumental means through which the “energy dominance” agenda is operationalized. Their semantic content is procedural rather than image-schematic: they specify what is done, not how it is cognitively construed.

Energy Resources constitute the most semantically dense and metaphorically rich group. It spans traditional energy sectors (coal, electricity, gasoline), defense-related energy assets (the National Defense Stockpile, electricity required by the defense industrial base), and territory-linked resources (energy activities on Federal land). The original policy texts characterize these resources through explicit conceptual metaphors such as “liquid gold” and “wealth beneath our feet,” which map energy onto the source domains of VALUABLE POSSESSION and BURIED TREASURE. The co-occurrence of “National Defense Stockpile” and “defense-industrial production” within this group further activates a STRATEGIC ASSET framing, construing energy resources as military matériel that requires protection. These metaphoric mappings are not sentence-level phenomena; they operate across multiple sentences and documents, forming the conceptual substrate on which the discourse strategies in Section 4.3 draw.

National Security is the smallest group (13 entities) but the most consequential for the discourse’s cognitive architecture. It consists of “national security,” “a threat to national security,” and “foreign adversaries.” These entities function as the target domain toward which the other three groups are discursively oriented. Energy resources are framed as objects of national security concern, policy actions as instruments of national security protection, and administrative entities as agents of national security enforcement. This target-domain anchoring is what transforms a sectoral energy policy into a securitized national agenda.

The four cluster groups collectively map the conceptual landscape of the “energy dominance” discourse. Administrative Entities and Policy Actions provide the institutional and instrumental infrastructure. Energy Resources supply the metaphorically elaborated source domain. National Security provides the target domain that anchors the entire discursive configuration. These domains, however, are not themselves image schemas or discourse strategies. Their integration into coherent cognitive-narrative structures is the function of the three strategies analyzed in Section 4.3: agentive assertion, which draws on the FORCE schema to construe institutional action as compulsion; oppositional framing, which draws on the CONTAINER schema to demarcate in-group and out-group boundaries; and prophetic narration, which draws on the PATH–GOAL schema to project a trajectory from present crisis to future restoration.

4.3 Macro-Level Analysis: Discourse and Social Practice

Narratives are the vehicle through which political actors perform the speech acts of securitization (Stivas, 2024), and discourse strategies are the concrete operations through which morality and authority are constructed (Arafat and Abdullah, 2025). Drawing on securitization theory and grounded in the four semantic cluster groups identified above, the manual qualitative analysis described in Section 3.4.4 identified three discourse strategies: threat construction, adversarial polarization, and restoration narrative.

4.3.1 Discourse Practice

Threat construction frames energy dependence as an existential security threat. The White House defines insufficient energy development as “an imminent and growing threat to American prosperity and national security,” warning that “without immediate remedial measures, this situation will sharply worsen in the near future.” This strategy recruits the CONTAINER schema to encode foreign energy reliance as boundary violation (dependence on external sources is construed as penetration of the national container) and the FORCE schema in its “removal of restraint” variant to construe deregulation as liberation from externally imposed constraints. The dissolution of the Interagency Working Group on the Social Cost of Greenhouse Gases, accompanied by the determination that the “social cost of carbon” calculation suffers from “logical flaws, politicization, and lack of legislative basis,” exemplifies this strategy at the institutional level: the elimination of scientific frameworks that justify regulatory constraint is itself an act of schematic force: removing an obstacle that blocks the path to energy dominance.

Adversarial polarization constructs binary oppositions along two axes. The first axis opposes the current administration to previous administrations: green transition policies are stigmatized as a “Green New Scam” waging “a war on American energy,” while New York State’s climate policy is characterized as a “climate change extortion law.” Entrepreneurs, workers, and local voters in traditional energy-producing states are integrated into a community of the oppressed, granting traditional energy policies moral legitimacy that transcends economic calculation. The second axis opposes “America First” to “foreign adversaries”: the green transition is condemned as “shipping energy jobs from Texas to Tehran, from the Midwest to Moscow, from Baton Rouge to Beijing,” portraying China, Iran, and Russia as actors that persistently erode U.S. energy dominance. This dual-axis strategy jointly activates the CONTAINER schema to demarcate in-group and out-group boundaries, both domestic (the previous administration and its allies) and foreign (competitor nations), while channeling the FORCE schema’s compulsion subtype into an us-versus-them dynamic that elevates domestic policy disagreements into conflicts with external forces.

Restoration narrative projects a trajectory from present crisis to national renewal. The negative pole emphasizes catastrophic deterioration without intervention; the positive pole declares that “expanding reliable and affordable energy production will lower inflation, grow the economy, create high-paying jobs,

and restore peace," and that "American energy dominance is delivering tangible benefits to hardworking families across the nation." This strategy recruits the PATH–GOAL schema to organize policy actions into a temporal sequence (threat diagnosis, remedial intervention, and prosperity restoration) with energy dominance positioned as both the means and the endpoint. Against the backdrop of the U.S.–Iran conflict disrupting global energy markets, this complete securitization chain of threat–response–benefit, led by fear mobilization and counterbalanced by hope, fosters public acceptance of extraordinary measures such as declaring a national energy emergency.

4.3.2 Social Practice

The persuasive force of those discourse strategies is enabled by the interplay of genre conventions, political-economic conditions, and embodied-cognitive foundations that constitute the social practice dimension of Fairclough's model.

At the level of genre, the executive order affords particular forms of schematic activation. Its characteristic imperative mood ("shall identify and exercise"), delegation of authority to named officials, and legal boundary-drawing clauses ("Nothing in this order shall be construed to...") create a textual environment in which the FORCE schema is grammatically encoded: the President's directive carries illocutionary force that compels institutional action. The genre's requirement that every action be attributed to a statutorily authorized actor accounts for the predominance of Administrative Entities in the argument clustering results. The genre's extensive use of disclaimer and jurisdictional clauses, in turn, explains why the majority of sentence clusters lack image-schematic profiles: large portions of these texts are not narratively elaborated but legally secured.

At the level of political economy, the "energy dominance" discourse emerges at a conjuncture marked by several intersecting conditions: the U.S. position as the world's largest oil and gas producer, the domestic backlash against green transition policies in traditional energy-producing states, the geopolitical disruption of global energy markets following the U.S.–Iran conflict, and the political imperative to deliver tangible economic benefits to core electoral constituencies. These conditions render the CONTAINER schema's boundary-violation logic intuitively resonant (the nation's energy resources are "ours," and their exploitation by external actors constitutes a violation), and the PATH–GOAL schema's crisis-to-renewal trajectory politically expedient, offering a narrative resolution to the contradictions between environmental commitments and economic interests.

At the level of embodied cognition, the persuasive force of these strategies draws on the pre-conceptual, sensorimotor character of image schemas. The CONTAINER schema arises from the bodily experience of boundedness—the body as a container, the experience of inside versus outside—and its recruitment in threat construction and adversarial polarization renders abstract geopolitical dependencies viscerally felt as intrusions. The PATH–GOAL schema arises from the bodily experience of purposeful movement from a starting point to a destination, and its recruitment in the restoration narrative transforms policy sequences into intuitively compelling journeys. The FORCE schema, operating subsidiarily, draws on the bodily experience of compulsion and the release of obstruction. These schemas constitute the cognitive

conditions of possibility for the discourse's persuasive effect, grounding abstract policy arguments in sensorimotor experiences shared by all human beings.

5 Discussion

The results demonstrate that the prompt-driven verification pipeline substantially improves the quality of SRL-extracted triples from policy texts. The key insight is that providing the LLM with structured SRL output as input, rather than asking it to annotate from scratch, constrains the task to validation and targeted repair, which reduces hallucination and format inconsistency, two common failure modes of generative LLMs when applied to specialized annotation tasks. Moreover, this pipeline achieves competitive performance without any domain-specific fine-tuning, using only a set of carefully engineered prompts and a modest number of few-shot examples, making it a viable and lightweight alternative for narrative mining in specialized policy domains where annotated corpora are scarce and fine-tuning is impractical.

From the perspective of embodied cognition research, this pipeline offers a methodologically rigorous approach to addressing the reliability challenges that have long plagued cognitive linguistic studies of political discourse. Traditional introspective analyses, while theoretically grounded, rely on the researcher's interpretive judgment, which may vary across analysts and limit systematic replication. By integrating automated SRL extraction with LLM-assisted verification against a gold-standard annotation set, this study provides a systematic quantitative framework for identifying image-schematic activations in large-scale corpora. The substantial improvement in precision (0.48 → 0.75) and the resulting F1 score of 0.79 offer empirical credibility for claims about the embodied grounding of political language.

Furthermore, the pipeline's contribution to embodied language research extends beyond performance metrics. The structured SRL output functions as an intermediate representation that preserves the directed semantic relations essential for schema identification. Rather than relying on the LLM to infer schemas from raw text—"a task that risks projecting anthropocentric patterns onto linguistic data"—the pipeline extracts predicate-argument structures first, then uses the LLM to validate and correct these structures. This design choice ensures that the final annotation reflects genuine linguistic patterns, while preserving the interpretability necessary for cognitive linguistic analysis.

The findings of this study speak to the broader theoretical debate between embodied and amodal approaches to language comprehension that motivates this Research Topic. Amodal theories hold that linguistic meaning is processed through abstract, propositional representations independent of sensorimotor systems; embodied theories argue that language comprehension recruit modality-specific, bodily-grounded simulations (for reviews, see Dove, 2023; Pelkey, 2023). The systematic activation of CONTAINER, PATH–GOAL, and FORCE schemas across the policy corpus provides corpus-level evidence consistent with the embodied position: abstract policy concepts (energy dependence, regulatory constraint, national renewal) are not merely described through metaphorical language but are cognitively constituted through sensorimotor schemas that render them intuitively accessible to audiences. This

interpretation fits with Fuchs's (2026) argument that concepts remain anchored in corporeal and intersubjective experience, with image schemas such as CONTAINER forming the embodied basis of categorization even at high levels of abstraction. On the methodological side, the pipeline addresses a common criticism of embodied approaches: that the evidence base relies too heavily on introspective analysis and small-N experimental designs that are difficult to replicate. By showing that image-schematic patterns can be systematically identified in naturalistic discourse with quantifiable performance metrics (F1 = 0.79), this study strengthens the empirical foundations of embodied language research. Together with recent work that brings large language models into cognitive linguistic analysis (Meng et al., 2025), this study adds to the cross-disciplinary research agenda that this Research Topic aims to advance. Future work that combines this pipeline with neuroimaging or behavioral paradigms could further test the cognitive reality of the schemas identified here.

Several limitations qualify these findings. First, the corpus size is limited. When supervised machine learning is applied to large-scale political texts, a substantial number of manually annotated training data points is typically required. Constrained by the selection of research objects, the corpus size is relatively small, which to some extent limits model training and the generalizability of results. Second, the adaptability of the annotation tool needs improvement. The AllenNLP BERT-based SRL model used in this study was trained on general news text (OntoNotes 5.0). When processing legal and administrative texts with highly nested structures, passive voice, and complex clauses, its accuracy in extracting narrative triples declines, and the entity clustering process still requires manual intervention to verify outputs, thus limiting the degree of automation and normalization in semantic mining. Third, the persistent difficulty with agent (ARG0) identification remains. Error analysis reveals that the LLM occasionally assigns the syntactically nearest subject rather than the discourse-level agent in complex sentences, and struggles with the deeper syntactic reasoning required to recover implicit agents in passive and nested constructions. Finally, the discourse strategy analysis, while conducted with methodological rigor, remains an interpretive exercise and is subject to the epistemological constraints of qualitative CDA research.

6 Conclusion

This study demonstrates that image schemas are systematically activated in U.S. energy dominance discourse through threat construction, adversarial polarization, and restoration narrative strategies. These schemas do not merely decorate political language; they constitute the sensorimotor architecture through which abstract policy concepts are rendered intuitively graspable, offering corpus-level evidence that political thought is grounded in embodied experience.

Methodologically, the prompt-driven SRL–LLM pipeline addresses a longstanding challenge in embodied language research: the reliance on introspective analysis. By achieving an F1 score of 0.79 against a manually annotated gold standard without domain-specific fine-tuning, the pipeline provides a computational framework for detecting semantic structure and embodied schemas in naturalistic

discourse at scale, contributing to the methodological innovation and replicability agendas that define the trend of cognitive linguistics.

Future work should extend this approach across languages and policy domains, and integrate behavioral or neuroimaging paradigms to test the cognitive reality of computationally identified schemas. This study adds to a growing body of cross-disciplinary research that brings computational methods to bear on foundational questions in embodiment of language.

Declarations

Data Availability

The policy documents analyzed in this study are publicly available from the White House website (https://www.whitehouse.gov). The annotated dataset and the computational pipeline code are available from the corresponding author upon reasonable request.

Competing interests

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Ethical approval

This article does not contain any studies with human participants performed by any of the authors.

Informed consent

This article does not contain any studies with human participants performed by any of the authors.

Generative AI Statement

The authors declare that a generative AI model (DeepSeek-R1) was employed exclusively as a data analysis tool for semantic role labeling verification and narrative annotation, as described in Section 2.3.3. No generative AI was used in the writing of this manuscript.

Author Contribution

XL: Conceptualization, Data curation, Investigation, Visualization, Writing – original draft, Writing – review & editing. LG: Data curation, Formal analysis, Methodology, Resources, Writing – original draft. HM: Conceptualization, Supervision, Validation, Writing – review & editing.All authors reviewed the manuscript.

References

1. Arafat BA, Abdullah ATHB (2025) From device to deed: A framework for mapping rhetorical moves to illocutionary effects in political speech. J Pragmat Discourse Anal 4(5):1–20. https://doi.org/10.32996/jpda.2025.4.5.1
2. Ash E, Gauthier G, Widmer P (2024) RELATIO: Text semantics capture political and economic narratives. Political Anal 32(2):115–136
3. Bakker RM, Schoevers AJ, van Drie RAN, Schraagen MP, de Boer MHT (2025) Semantic role extraction in law texts: A comparative analysis of language models for legal information extraction. *Artificial Intelligence and Law*. Advance online publication. https://doi.org/10.1007/s10506-025-09437-x
4. Ba Y, Tan H (2025) Moral foundations of energy transitions in the United States: A cross-sector discourse analysis. Energy Res Social Sci 125:104090
5. Chen C, Liu R (2025) How media discourse constructs the role of new energy policies in addressing energy challenges: A comparative analysis of China and the U.S. Energy Policy 206:114807. https://doi.org/10.1016/j.enpol.2025.114807
6. Chen H, Zhang M, Li J, Zhang M, Øvrelid L, Hajič J et al (2025) Semantic role labeling: A systematical survey. arXiv. https://doi.org/10.48550/arXiv.2502.08660
7. Cheng S, Li Z, Wei T (2023) A threat intelligence entity relation extraction method integrating bootstrapping and semantic role labeling. J Comput Appl 43(5):1445–1453
8. Chilton P (2004) Analysing political discourse: Theory and practice. Routledge
9. Dove GO (2023) Rethinking the role of language in embodied cognition. Philosophical Trans Royal Soc B 378(1870):20210375. https://doi.org/10.1098/rstb.2021.0375
10. Dunham IM, Schlosser K (2016) Energy security discourses and environmental protection measures in U.S. federal energy legislation: An introductory exploration. Extractive Industries Soc 3(1):86–94. https://doi.org/10.1016/j.exis.2015.11.015
11. Fairclough I, Fairclough N (2012) Political discourse analysis: A method for advanced students. Routledge
12. Fuchs T (2026) The experiential basis of concepts: Integrating embodied and enactive accounts. Front Psychol 16:1710102. https://doi.org/10.3389/fpsyg.2025.1710102
13. Hart C (2014) *Discourse, grammar and ideology: Functional and cognitive perspectives*. Bloomsbury
14. Hart C (2024) Language, image, gesture: The cognitive semiotics of politics. Cambridge University Press
15. Jia K, Fu M, Zhou L (2026) Consensus and divergence in U.S. and British strategic communities' science and technology policies toward China: Narrative mining based on RELATIO text semantic capture (2021–2025). J Intell 45(1):39–47
16. Jing E, Ahn Y-Y (2021) Characterizing partisan political narrative frameworks about COVID-19 on Twitter. EPJ Data Sci 10(1):53

17. Johnson M (1987) The body in the mind: The bodily basis of meaning, imagination, and reason. University of Chicago Press. https://doi.org/10.7208/chicago/9780226177847.001.0001
18. Kövecses Z (2015) Where metaphors come from: Reconsidering context in metaphor. Oxford University Press
19. Lazard AJ, Scheinfeld E, Bernhardt JM, Wilcox GB, Suran M (2015) Detecting themes of public concern: A text mining analysis of the Centers for Disease Control and Prevention's Ebola live Twitter chat. Am J Infect Control 43(10):1109–1111. https://doi.org/10.1016/j.ajic.2015.05.025
20. Liu L (2024) Exploring the emerging trends of energy discourse: A bibliometric analysis. Energy Strategy Reviews 52:101338. https://doi.org/10.1016/j.esr.2024.101338
21. Ma J, Zhao X (2025) Zero-sum game and co-opetition: A contrastive study of the energy discourses on China by US and EU institutions based on discourse-historical approach. J Ludong Univ (Philosophy Social Sci Edition) 42(1):35–47
22. Meizlish T, Ziffo C (2026) Evaluating an LLM's performance in annotating discourse strategies. Corpus Pragmat, *10*(1)
23. Meng H, Li X, Sun J (2025) Large language models prompt engineering as a method for embodied cognitive linguistic representation: A case study of political metaphors in Trump's discourse. Front Psychol 16:1591408. https://doi.org/10.3389/fpsyg.2025.1591408
24. Mochtak M (2025) Chasing the authoritarian spectre: Detecting authoritarian discourse with large language models. Eur J Polit Res 64(3):1304–1325
25. Musolff A (2016) Political metaphor analysis: Discourse and scenarios. Bloomsbury
26. Nanyonga A, Wild G (2025), January Analyzing aviation safety narratives with LDA, NMF and PLSA: A case study using Socrata datasets. In *2025 2nd International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE)*. IEEE. https://doi.org/10.1109/RMKMATE64874.2025.11042356
27. Pelkey J (2023) Embodiment and language. Wires Cogn Sci 14(5):e1649. https://doi.org/10.1002/wcs.1649
28. Senator F, Lakhfif A, Zenbout I, Boutouta H, Mediani C (2025) Leveraging ChatGPT for enhancing Arabic NLP: Application for semantic role labeling and cross-lingual annotation projection. IEEE Access 13:3707–3725. https://doi.org/10.1109/ACCESS.2025.3525493
29. Shahsavari S, Holur P, Tangherlini TR, Roychowdhury V (2020) Conspiracy in the time of corona: Automatic detection of emerging COVID-19 conspiracy theories in social media and the news. J Comput Social Sci 3(2):279–317. https://doi.org/10.1007/s42001-020-00086-5
30. Smirnov IV, Larionov D (2025) LLM for semantic role labeling of emotion predicates in Russian. Supercomputing Front Innovations 12(3):31–46
31. Stivas D (2024) Variations in the intensity of the securitization narratives at the EU level: Securitizing the European refugee crisis. Front Political Sci 6:1460531. https://doi.org/10.3389/fpos.2024.1460531

32. Tangherlini TR, Shahsavari S, Shahbazi B, Ebrahimzadeh E, Roychowdhury V (2020) An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, Pizzagate and storytelling on the web. PLoS ONE 15(6):e0233879. https://doi.org/10.1371/journal.pone.0233879
33. Wallace E, Tuyls J, Wang J, Subramanian S, Gardner M, Singh S (2019) AllenNLP Interpret: A framework for explaining predictions of NLP models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing: System Demonstrations* (pp. 7–12). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-3002
34. Zhao W, Guo S, Lerman K, Ahn Y-Y (2024) Discovering collective narrative shifts in online discussions. In *Proceedings of the Eighteenth International AAAI Conference on Web and Social Media* (Vol. 18, pp. 1472–1483). AAAI Press. https://doi.org/10.1609/icwsm.v18i1.31427
35. Zhao X, Zhao L (2021) A study of Trump administration's energy strategy discourse: A corpus-based critical discourse analysis. J China Univ Petroleum (Social Sci Edition) 37(1):28–35. https://doi.org/10.13216/j.cnki.upcjess.2021.01.0004
36. Zinovyeva ES, Trapeznikov VP (2026) International political bias in large language models: A critical discourse analysis of narratives in ChatGPT, LLaMA, Gemini, and DeepSeek. Polis Political Stud 2026(1):157–177. https://doi.org/10.17976/jpps/2026.01.11

Figures

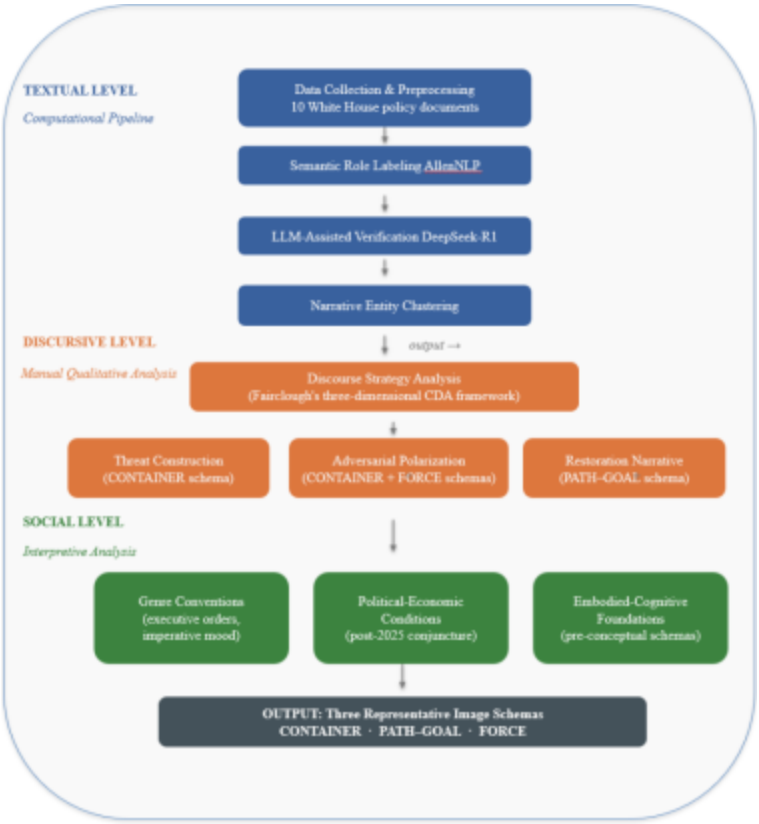


Figure 1

Research Framework

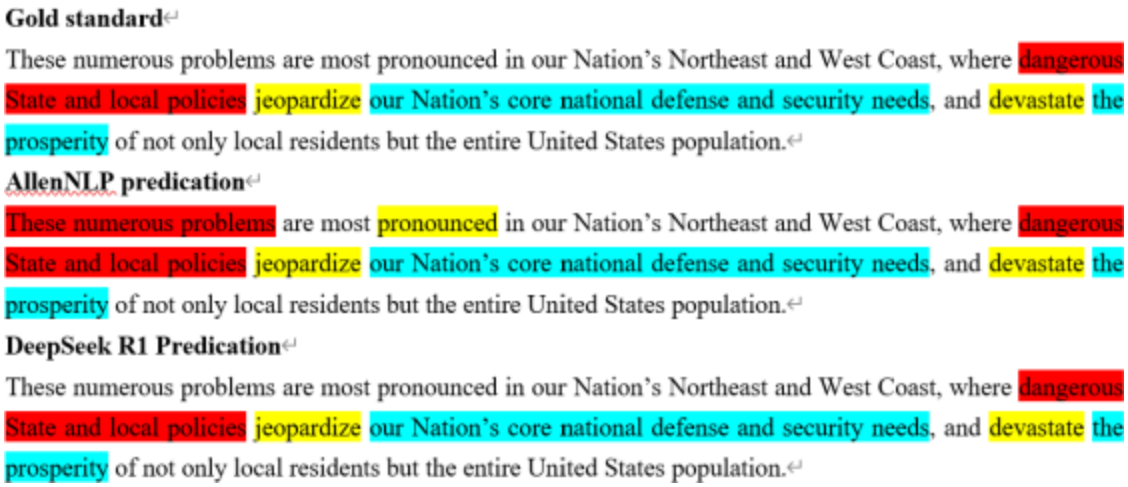


Figure 2

Performance comparison of different approaches for semantic role labeling

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryMaterial.zip
- Appendix.docx