Symbolic Mediation of Language-Based Decision Support in Tactical Contexts

Jaye Nias[1], Lashaun Baddol[1], Saurav K. Aryal[1], Jeremy Blackstone[1], Simone Smarr[1], Lucretia Williams[1], Gloria Washington[1]

[1]Human-Centered AI Institute, Howard University, Washington, DC, USA
jaye.nias@howard.edu, lashaun.baddol@bison.howard.edu, saurav.aryal@howard.edu, jeremy.m.blackstone@howard.edu, simone.smarr@howard.edu, lucretia.williams1@howard.edu, gloria.washington@howard.edu

Abstract

Language-based AI systems are increasingly explored as decision-support tools in tactical and operational contexts, where timely interpretation and action are critical. While recent advances enable language models to generate contextually grounded responses, their outputs are typically delivered as free-form text, placing interpretive burden on human decision-makers operating under time pressure. This paper addresses this representational gap by introducing a neuro-symbolic approach that mediates language model outputs through symbolic representations. We present a system architecture that transforms outputs from an already contextualized language model into standardized operational symbols using a symbolic inference layer grounded in rule-based mappings and confidence constraints. Rather than replacing language-based explanations, the system supports dual-mode presentation, preserving textual reasoning while rendering salient decision-relevant elements as symbols familiar to operational practice. This work contributes a representational framework for human-centered AI decision support in time-sensitive environments.

Introduction

Defense decision-making increasingly operates across a continuum from sensing and information fusion to judgment, coordination, and execution under uncertainty. AI systems are now routinely used to summarize data, assess situations, and support reasoning across this pipeline, particularly through language-based interfaces that synthesize complex operational inputs into textual analyses. While these systems offer expressive reasoning and explanatory power, their outputs are often delivered as unstructured language that does not align with the symbolic conventions, visual encodings, and representational constraints that guide time-sensitive tactical action. Operators frequently rely on compressed, standardized representations to support rapid interpretation, comparison, and shared understanding as situations evolve. This representational mismatch motivates hybrid approaches that bridge neural inference and symbolic structure. In this work, we introduce a neuro-symbolic architecture that treats representation as a core design problem, mediating language-based reasoning through a symbolic inference layer that selectively renders decision-relevant elements into operationally meaningful symbolic forms while preserving contextual richness and human judgment.

Motivation and Background

Operational decision-making environments place distinctive demands on how information is represented, interpreted, and acted upon. In tactical contexts, decision-makers must integrate incomplete and evolving information under time pressure, where delays or misinterpretations can have significant consequences. As a result, the effectiveness of computational decision-support systems depends not only on the accuracy of the information they provide, but on how well their outputs align with human cognitive processes and established operational practices.

Tactical Decision Making Under Stress

Long-standing defense research has shown that fluent information delivery alone does not ensure effective decision support. Early work in the Tactical Decision Making Under Stress (TADMUS) program emphasized that decision-support systems must account for cognitive workload, time pressure, and the representational demands placed on human operators (Morrison et al. 1996). This line of research highlights the importance of structuring system outputs in ways that support rapid interpretation and integration into human judgment, rather than treating computational outputs as self-sufficient answers (Klein 2017).

Representation and Human-Centered Decision Support

A central challenge highlighted in this literature is the representational gap between system outputs and operational decision-making needs (Endsley 2017). Text-centric interfaces place the burden of interpretation on the human operator, who must extract salient elements, assess relevance, and construct an operational picture before acting. Prior work in human-centered decision support demonstrates that when information is presented in forms misaligned with task demands, it can increase cognitive load or lead to inappropriate reliance, even when the underlying information is correct (Woods 2018; Parasuraman, Sheridan, and Wickens 2000). These findings underscore the importance of representation

Architecture Layer	System Components	Support Function
Neural Language Layer	Contextualized language model; domain sources	Contextual reasoning; explanation
Symbolic Mediation Layer	Parser; attribute structuring; rule-based mapping; confidence constraints	Logical structuring; uncertainty gating
Representational Interface Layer	Symbol library; dual-mode dashboard	Rapid interpretation; decision support

Table 1: Hybrid neuro-symbolic architecture and associated support roles.

as a design concern, not merely a presentation choice (Norman 2013).

Implications for AI-Supported Decision Systems

Recent advances in language-based AI systems have expanded interest in using natural language interfaces for decision support in defense settings. While these systems offer new flexibility for information synthesis, they also risk reintroducing representational burdens if outputs are not aligned with operational practices. Neuro-symbolic approaches suggest one pathway for integrating the flexibility of language-based reasoning with structured representations that support interpretability and human-AI teaming (Gunning 2017). This perspective treats symbolic structure as a complementary layer that shapes how AI-generated information is interpreted and used in decision-making contexts.

System Architecture and Symbolic Mediation

The system described in this paper operationalizes a neuro-symbolic approach to structuring language-based decision support for tactical contexts. Rather than introducing a new language model, the design focuses on how outputs from an already contextualized language model can be mediated through symbolic representations that better align with operational decision-making practices. This section describes the system architecture, symbolic inference process, and representational schema in a unified manner to support clear conceptual and visual mapping.

Hybrid Neuro-Symbolic Architecture

The system integrates neural language-based reasoning with symbolic mediation through a modular architecture organized into three complementary layers: a neural language layer, a symbolic mediation layer, and a representational interface layer. Each layer plays a distinct role in shaping how language-based AI outputs are transformed into decision-relevant representations for tactical contexts.

- **Neural Language Layer:** User queries are processed by a contextualized language model grounded in domain-relevant documentation. This layer produces natural language responses that support contextual reasoning, explanation, and sensemaking, capturing nuance, qualifiers, and conditional structure relevant to operational decision-making.
- **Symbolic Mediation Layer:** Outputs from the neural language layer are passed to a symbolic inference layer that selectively interprets decision-relevant elements of the text. Through parsing, lightweight structuring, and rule-based mapping with confidence constraints, this layer reshapes language-based reasoning into structured symbolic representations without fully formalizing or replacing the original text.
- **Representational Interface Layer:** The system presents both language-based and symbolic outputs through a dual-mode dashboard interface. Textual reasoning is preserved to maintain context and traceability, while inferred symbols are rendered to support rapid interpretation, orientation, and comparison. This layer enables users to flexibly engage with different representations depending on task demands and time pressure.

Together, these layers reflect the complementary roles of language and symbols in decision support: language supports rich reasoning and explanation, symbolic mediation imposes representational discipline and constraint, and the interface layer enables rapid, human-centered interpretation without automating decisions.

Information Flow and Dual-Mode Presentation

System outputs are generated and presented through the following pipeline:

1. **User Query:** A natural language query is submitted to the contextualized language model.
2. **Language-Based Reasoning Output:** The model produces a natural language response containing actionable details and contextual qualifiers.
3. **Symbolic Mediation:** Decision-relevant attributes are identified and structured for symbolic interpretation without attempting to exhaustively formalize the language output.
4. **Dual-Mode Dashboard Interface:** Outputs are presented through a dual-mode interface that maintains both representational forms:
   - **Textual Reasoning Panel:** Preserves the full language-based response to support contextual understanding and traceability.
   - **Symbolic Representation Panel:** Displays inferred symbols that convey salient operational information at a glance.

By maintaining both representations simultaneously, the interface supports flexible sensemaking as task demands and time constraints evolve.

Symbolic Inference Pipeline

Symbolic inference functions as a representational mediation layer rather than a comprehensive formalization of language model output. The objective is to selectively surface decision-relevant elements that can be reliably expressed through standardized operational symbols. The inference pipeline consists of the following stages:

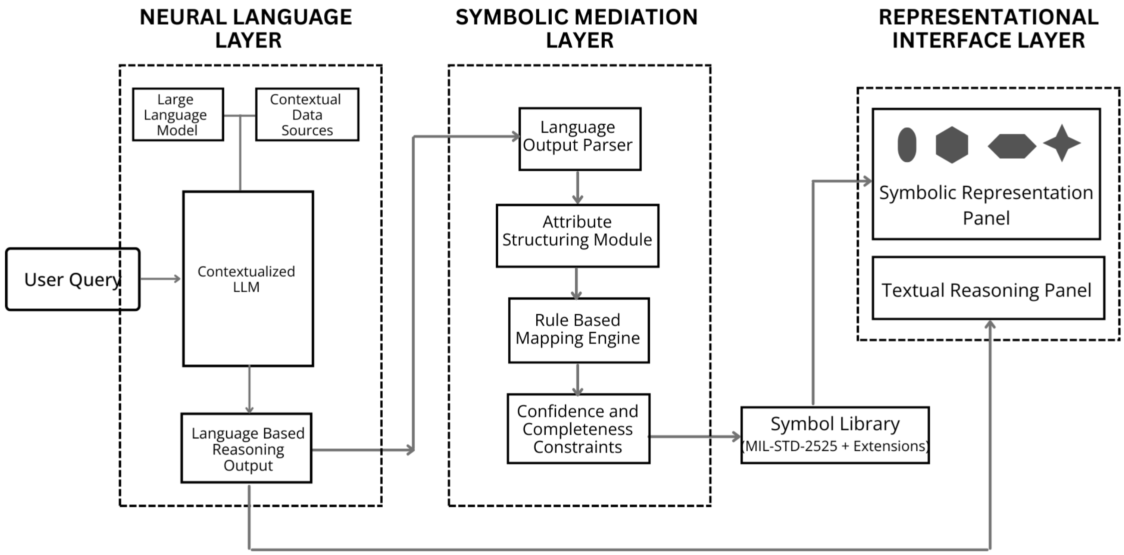

Figure 1: Hybrid neuro-symbolic architecture illustrating the neural language layer, symbolic mediation layer, and representational interface layer.

Language Output Parsing

- **Language Output Parser:** Parses language-based reasoning outputs to identify entities, operational attributes, and contextual qualifiers relevant to tactical decision-making, including references to objects of interest, inferred status or affiliation, spatial or temporal cues, and conditional language.

Attribute Structuring

- **Attribute Structuring Module:** Transforms extracted elements into structured representations that serve as inputs for symbolic mapping. This structuring is intentionally lightweight to support adaptability across scenarios while avoiding brittle over-specification.

Rule-Based Mapping and Constraints

- **Rule-Based Mapping Engine:** Applies rule-based logic that encodes domain knowledge about how combinations of object type, affiliation, and status correspond to operational meanings.
- **Confidence and Completeness Constraints:** Determines whether symbolic representations should be generated. When conditions for reliable inference are not met, symbolic output is withheld to avoid introducing misleading cues.

Symbol Schema and Representation

Symbol generation draws on a shared representational foundation while remaining extensible:

- **Symbol Library:** Grounded in standardized military symbology defined by MIL-STD-2525, supporting familiarity and consistency with established operational practices.
- **Custom Symbol Extensions:** Enable representation of concepts not fully captured by existing standards, allowing the symbolic schema to evolve alongside operational needs.

By treating symbolic inference and schema design as an extensible representational framework rather than a fixed taxonomy, the system maintains alignment with established practices while accommodating emerging representational requirements.

Dual-Mode Dashboard Interface

The system presents outputs through a dual-mode dashboard interface in Figure 2 designed to support interpretation under varying cognitive and temporal demands.

The textual reasoning panel preserves the full language model response, providing contextual detail, justification, and traceability. This view supports deliberation, verification, and reflection, particularly in situations where users need to understand how conclusions were reached. In parallel, the symbolic representation panel displays only those symbols that can be reliably inferred by the symbolic inference layer. These symbols are intended to convey salient operational information at a glance, supporting rapid orientation and comparison without requiring users to parse extended text.

The current interface in Figure 2 serves as a pilot instantiation of this design, illustrating how symbolic mediation can be operationalized without constraining future interface or workflow integrations. By presenting both representations

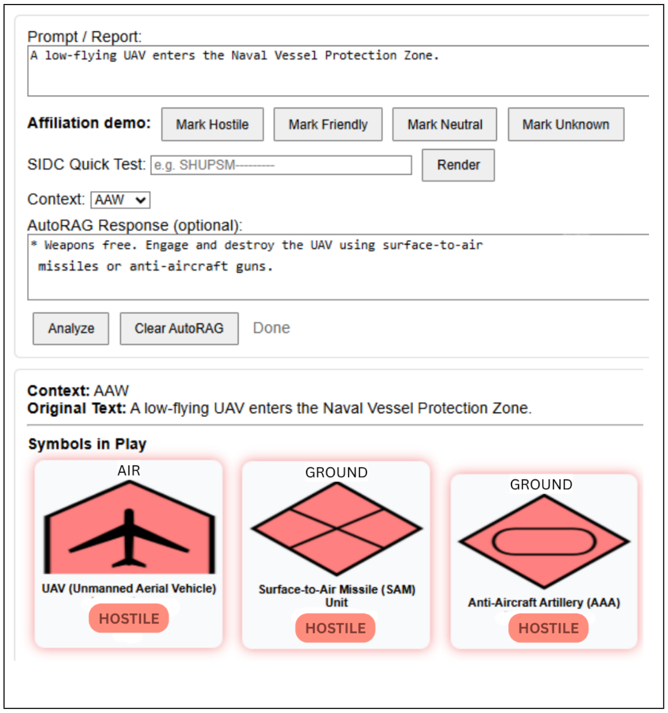

Figure 2: Dual-mode dashboard interface illustrating the parallel presentation of language-based reasoning and inferred symbolic representations.

concurrently, the dashboard supports adaptive sensemaking. Users can rely on symbolic cues when time pressure is high, while retaining access to language-based explanations when greater context or justification is needed. This design reflects the system's role as a decision support aid that shapes interpretation rather than automating decision-making.

Discussion

This work offers several observations about the role of representation in language-based decision support for tactical contexts. First, introducing a symbolic mediation layer provides a means of making salient elements of language model outputs immediately legible without requiring operators to parse extended textual responses. By selectively rendering decision-relevant attributes as standardized symbols, the system supports rapid orientation while preserving access to contextual detail through accompanying language-based reasoning.

Second, symbolic mediation functions as a diagnostic mechanism rather than a purely translational one. In cases where symbols cannot be confidently inferred, the absence of symbolic output becomes informative, signaling ambiguity, missing information, or implicit assumptions in the underlying language-based reasoning. More broadly, these observations highlight the importance of treating representation as a first-class design concern in AI-enabled decision support. Language models are often evaluated primarily on fluency and plausibility, yet operational use requires attention to how information is structured, constrained, and rendered actionable. Symbolic mediation offers a complementary representational strategy that bridges expressive reasoning and operational practice, supporting shared understanding and adaptive sensemaking across the sensing-to-execution pipeline.

Limitations and Future Directions

This work is limited in scope by its focus on system architecture and representational design rather than empirical evaluation with operational users. The symbolic inference mechanisms rely on rule-based mappings informed by representational analysis and operational standards, and have not yet been validated in live or simulated decision-making settings. As such, the system does not claim to improve decision accuracy or performance, but instead illustrates a design approach for structuring language-based reasoning in support of human judgment. Future work will examine how language model outputs align with symbolic representations across decision-relevant concepts, informing refinement of mapping rules, confidence constraints, and symbolic schemas. Additional directions include integrating the system into training and rehearsal workflows and investigating how symbolic mediation supports trust calibration, explanation, and shared situational awareness in human–AI teams. By treating representation as an active design layer rather than a passive output format, this work points toward a broader research agenda for aligning language-based AI systems with the cognitive and operational demands of time-sensitive decision-making environments.

Acknowledgments

This material is based upon work supported by the Office of Naval Research under Grant No. N00014-22-1-2714. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research, the Department of Defense or the Department of War.

References

Endsley, M. R. 2017. Toward a theory of situation awareness in dynamic systems. In *Situational awareness*, 9–42. Routledge.

Gunning, D. 2017. Explainable artificial intelligence (xai). *Defense advanced research projects agency (DARPA), nd Web*, 2(2): 1.

Klein, G. A. 2017. *Sources of power: How people make decisions*. MIT press.

Morrison, J. G.; Kelly, R. T.; Moore, R. A.; and Hutchins, S. G. 1996. Tactical decision making under stress (TADMUS) decision support system. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*.

Norman, D. 2013. *The design of everyday things: Revised and expanded edition*. Basic books.

Parasuraman, R.; Sheridan, T. B.; and Wickens, C. D. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, 30(3): 286–297.

Woods, D. D. 2018. Toward a theoretical base for representation design in the computer medium: Ecological perception and aiding human cognition. In *Global perspectives on the ecology of human-machine systems*, 157–188. CRC Press.