

ProEthica: A Professional Role-Based Ethical Analysis Tool Using LLM-Orchestrated, Ontology-Supported Case-Based Reasoning

Christopher B. Rauch, Rosina O. Weber

Drexel University
3675 Market Street
Philadelphia, PA 19104 USA
cr625@drexel.edu, rw37@drexel.edu

Abstract

Professional ethics committees currently lack structured tools to identify relevant ethical concepts from complex narratives and compare them against prior decisions. ProEthica analyzes professional ethical scenarios against established codes and precedent cases. The system uses large language models (LLMs), leveraging their natural language processing capabilities to extract nine types of components (Roles, Principles, Obligations, States, Resources, Actions, Events, Capabilities, and Constraints) from case text and scenario descriptions. Domain-specific ontologies provide precise definitions that constrain LLM output to match formal concept specifications, ensuring consistency across extraction and validation. Case-based reasoning identifies precedent cases and analogous situations relevant to the scenario under analysis. The current implementation processes engineering ethics cases with complete provenance tracking and ontology-driven validation. The framework supports extension to other professional domains with established codes and precedent systems.

Code — <https://github.com/cr625/proethica>

Video Demo — <https://www.proethica.org/demo>

ProEthica is an analysis tool that combines case-based reasoning and ontology-supported validation, orchestrated by LLMs via LangGraph, to help professional ethics committees analyze ethical scenarios against established standards and precedents and to determine whether actions by professionals are ethical. The system extends our prior research on precedent-based role ethics (Rauch et al. 2025) and builds on McLaren’s approach (McLaren 2003), which defines ethical principles through specific case examples from engineering board decisions. The framework provides a methodology applicable across multiple professional domains.

ProEthica employs a nine-component formal framework to capture the essential components of professional ethical analysis and to describe structured representations of codes of ethics and precedent cases for each domain. This framework, defined as $D = (R, P, O, S, Rs, A, E, Ca, Cs)$, synthesizes nine concepts from the computational ethics literature into a unified structure for the first time. The components organize according to logical dependencies

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: ProEthica core components organized by function: Processing (LLM and LangGraph orchestration), Knowledge (ontology definitions and case repository), and Infrastructure (Model Context Protocol (MCP) and hybrid database storage). Bidirectional arrows indicate interactive querying between Processing and Knowledge, while the dashed arrow shows direct data persistence to Infrastructure.

in profession-based ethical decision analysis and build on Berreby, Bourgne, and Ganascia (2017)’s modular architecture and Oakley and Cocking (2001)’s role-generated obligations.

Roles (R) define professional positions that generate abstract Principles (P), which specify concrete Obligations (O) that form the deontic foundation. States (S) and Resources (Rs) establish the situational context and available knowledge within which these obligations apply. Events (E) precipitate Actions (A) that professionals take. Capabilities (Ca) and Constraints (Cs) provide guardrails that ensure actions align with obligations and professional standards.

Example: Processing NSPE Case 24-2

We illustrate the interactive workflow using National Society of Professional Engineers (NSPE) Board of Ethical Review (BER) Case 24-2 (*Use of Artificial Intelligence in Engineering Practice*), with the full example available at <https://proethica.org/cases/7>.

As a Board-adjudicated case, Case 24-2 includes both the fact pattern and the Board’s ethical analysis, questions, and conclusions. Such cases enable system validation and precedent database construction used for analyzing future scenarios where outcomes are not yet known. After uploading the case text, the user reviews the formatted narrative and initi-

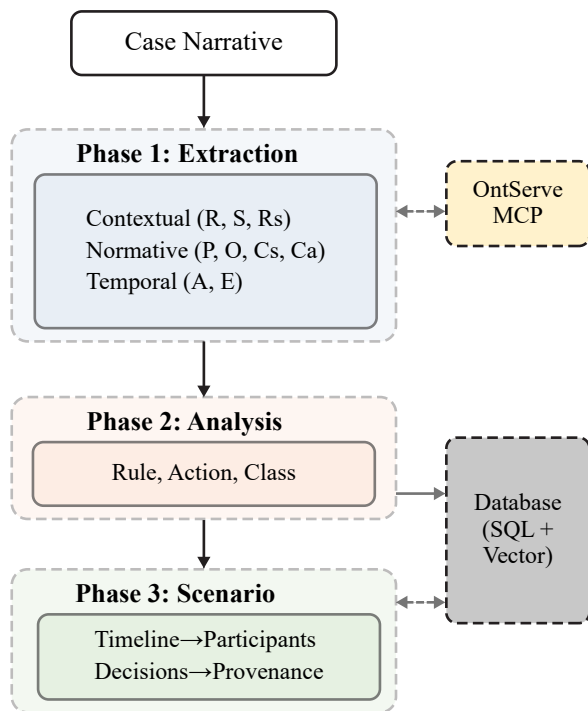


Figure 2: ProEthica three-phase architecture for ethical case analysis. Phase 1 performs multi-pass concept extraction using ontology-validated definitions via MCP queries to OntServe. The Contextual Framework pass extracts roles (R), states (S), and resources (Rs); the Normative Requirements pass identifies principles (P), obligations (O), constraints (Cs), and capabilities (Ca); and the Temporal Dynamics pass captures actions (A) and events (E). Phase 2 analyzes institutional rules, maps actions to normative structures, and classifies ethical transformations. Phase 3 generates interactive scenarios with participant mapping, decision points, and provenance tracking.

ates extraction.

Phase 1: Multi-Pass Extraction

Table 1 shows the three-pass extraction process for Case 24-2. Each pass builds on prior results. The system identifies both ontology matches and new discoveries requiring user review.

Phase 2: Case Synthesis and Analysis

After extraction completes (134 total entities), the user initiates six-part case analysis. The system identifies 2 principle tensions (*efficiency versus competence*) and 2 obligation conflicts (*timely delivery versus validation*). It maps 6 actions taken and 5 alternatives not pursued (*hire specialist, decline project*), then classifies the case as *transfer transformation* with pattern `ai_competence_boundary_violation`. The user reviews analysis results, adjusts tension characterizations, and

Pass 1: Contextual Framework (49 entities)

- Roles (4): e.g., *Engineer, Client*
- States (16): e.g., *Engineer lacks AI competence*
- Resources (29): e.g., *NSPE Code II.1.a*

User: Reviews entities, edits 2 definitions, approves 1 new class

Pass 2: Normative Requirements (74 entities)

- Principles (18): e.g., *Hold paramount public safety*
- Obligations (18): e.g., *Verify AI-generated designs*
- Constraints (18): e.g., *Cannot certify beyond competence*
- Capabilities (20): e.g., *Can hire specialists*

User: Modifies 3 definitions, approves 2 new classes

Pass 3: Temporal Dynamics (10 entities + 18 relations)

- Events (3): e.g., *Client requests AI design*
- Actions (7): e.g., *Uses AI without verification*
- Temporal Relations (12): e.g., *request before decision*
- Causal Chains (6): e.g., *lacks competence → violates obligation*

User: Adjusts 1 causal chain, commits entities and relations

Table 1: Three-Pass Extraction Results for NSPE Case 24-2

approves classification for database storage.

Phase 3: Interactive Scenario Generation

The user initiates scenario generation, which executes 6 stages. These stages include timeline construction, participant mapping (4 participants with LLM-enhanced profiles), relationship networks linking participants to ethical tensions, decision point identification, causal chain visualization, and normative framework display with links to extracted Principles, Obligations, and NSPE Code provisions. The output is presented in an interactive web interface with entity provenance, timeline visualization, decision tree representation, code of ethics provision links, and comparison to 23 similar precedents.

Conclusion

ProEthica demonstrates a generalizable methodology for implementing professional ethics analysis across multiple domains. Previous systems like SIROCCO (McLaren 2003) and MedEthEx (Anderson, Anderson, and Armen 2006) were applied in specific domains. ProEthica provides an extensible framework through the formal tuple and three-layer ontological architecture enabling structured extraction from professional codes and cases. This architecture separates domain-specific content from the underlying extraction and reasoning mechanisms. The current implementation demonstrates the framework using engineering ethics.

Future development includes extension to medical and legal ethics domains and creation of an interface for scenario generation based on user-supplied vignettes.

References

- Anderson, M.; Anderson, S. L.; and Armen, C. 2006. MedEthEx: A Prototype Medical Ethics Advisor. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, 1759. AAAI Press.
- Berreby, F.; Bourgne, G.; and Ganascia, J.-G. 2017. A declarative modular framework for representing and applying ethical principles. In *16th Conference on Autonomous Agents and MultiAgent Systems*, 96–104.
- McLaren, B. M. 2003. Extensionally defining principles and cases in ethics: An AI model. *Artificial Intelligence*, 150(1-2): 145–181.
- Oakley, J.; and Cocking, D. 2001. *Virtue Ethics and Professional Roles*. Cambridge University Press.
- Rauch, C. B.; Molineaux, M.; Mainali, M.; Sen, A.; Floyd, M. W.; and Weber, R. O. 2025. Role-Based Ethics for Decision-Maker Alignment. In *2025 IEEE Conference on Artificial Intelligence (CAI)*, 1209–1212.