Eight Linked Views from Nine-Component Extraction: Decomposing Professional Ethics Narratives

Anonymous Author(s)

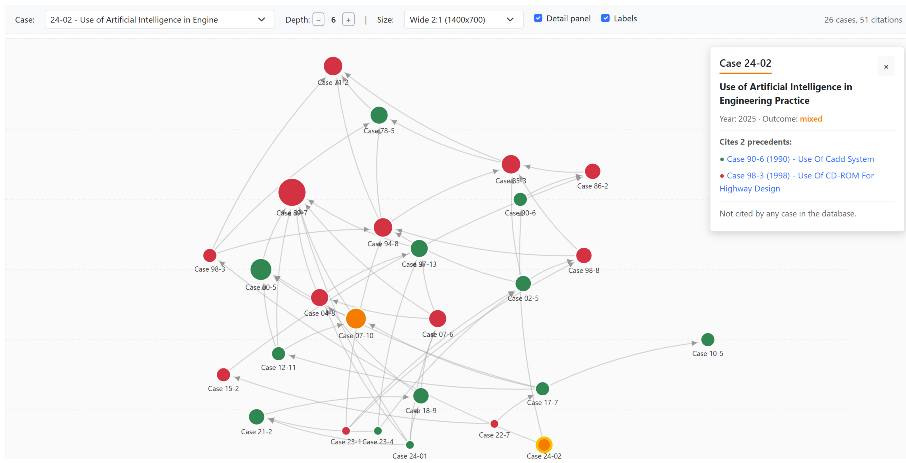

Figure 1: Precedent citation lineage for NSPE Case 24-02 (*Use of Artificial Intelligence in Engineering Practice*), showing a six-hop neighborhood of 26 cases and 51 citations. Nodes represent board opinions positioned by year. Directed edges point from citing case to cited precedent.

Abstract

This paper presents a system that decomposes engineering ethics narratives into nine typed components organized across three functional dimensions, then recomposes them through eight linked synthesis views. Three of the views support progressive disclosure, where collapsed headers carry count badges that readers can expand on demand. Shared entities with provenance annotations serve as typed links across views, so a reader can trace any extracted component through appearances in different views and back to the source text. An ontology-constrained extraction process populates the nine components, with a domain ontology that co-evolves as each processed case contributes new vocabulary. A Precedents view extends the architecture beyond individual cases by connecting opinions through board-authored citations and through shared ontology classes derived from constrained extraction, currently across a base of 119 processed opinions. A walkthrough of NSPE Board of Ethical Review Case 24-02 demonstrates how a large language model and ontology tooling decompose a single opinion into linked synthesis views.

HT '26, London, United Kingdom

ACM ISBN 978-1-4503-XXXX-X/2026/09
https://doi.org/XXXXXXX.XXXXXXX

CCS Concepts

• **Information systems → Ontologies**; *Web applications*; • **Human-centered computing → Hypertext / hypermedia**.

Keywords

professional ethics, narrative decomposition, multi-view synthesis, provenance tracking, ontology-constrained extraction, human-AI collaboration

ACM Reference Format:
Anonymous Author(s). 2026. Eight Linked Views from Nine-Component Extraction: Decomposing Professional Ethics Narratives. In *Proceedings of 37th ACM Conference on Hypertext and Social Media (HT '26)*. ACM, New York, NY, USA, 10 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

The National Society of Professional Engineers (NSPE) Board of Ethical Review has published over 650 advisory opinions since 1954 [14]. Each opinion applies provisions of the NSPE Code of Ethics to a specific fact pattern and follows a consistent internal structure of facts, ethical questions, discussion, code references, and conclusions. The accumulated opinions extend the practical scope of code provisions to situations the original text does not address. They form a reference archive for the board and the profession.

The internal structure supports sequential reading, but the components a practitioner needs to compare across cases, such as the professional roles at issue, the obligations triggered, the temporal sequence of actions and events, and the normative warrants the board applied, appear throughout without separation. A practitioner looking for prior opinions with analogous obligations or comparable role configurations cannot retrieve them by section alone, and full-text search matches word co-occurrence without distinguishing analytical role.

This paper describes a system that transforms these linear professional narratives into multi-path structures. An ontology-constrained extraction process [4] decomposes each opinion into nine typed components organized across three functional dimensions. The nine components draw on prior work in modular ethical architectures and role-generated professional obligations [6, 15]. The system then recomposes these components into eight linked analytical views, each framing the same case differently. Shared entities with provenance annotations link views to each other and back to the source text.

A view of precedential cases extends this linking across opinions through two layers, one drawn from board-authored citations extracted during synthesis and the other from constrained extraction against a shared ontology.

These architectural elements build on prior publications. The nine-component case representation and the extraction process that populates it have been described elsewhere [4], and a companion paper addresses component-aware case retrieval [3]. The present paper contributes the recomposition architecture: eight linked synthesis views with provenance chains and progressive disclosure, and a Precedents view that extends the linking across cases.

2 Background

2.1 Professional Ethics Cases as Structured Narratives

The regularity of NSPE Board of Ethical Review opinions makes them amenable to computational decomposition, but only if the decomposition schema reflects the analytical categories that practitioners apply when reading and comparing cases.

Ashley and McLaren [5] argued that professional ethics analysis operates through structured case comparison, where shared representation vocabularies enable identification of ethically relevant similarities and differences between cases. McLaren [12] demonstrated that ethical principles acquire their meaning extensionally, through accumulation of specific case applications rather than through abstract definition alone. In McLaren's account, each new board opinion refines or extends the scope of an existing principle by applying it to a novel factual situation. Where McLaren addressed how principles acquire meaning through cases, Oakley and Cocking [15] addressed where they originate. Professional roles generate specific obligations that constrain the ethical analysis. For example, an engineer's duty to verify design outputs follows from the engineer's role and commitment to public safety.

These characteristics distinguish professional ethics opinions from other narrative domains. Boards compose opinions to be read, compared, and applied by practitioners [5], and they cite prior opinions when evaluating new cases. The computational systems reviewed in the following subsection address different aspects of this decomposition and recomposition problem.

2.2 Computational Approaches to Document Decomposition and Multi-View Presentation

Prior systems have addressed decomposition, multi-view presentation, or cross-case linking separately, but not the combination that professional ethics narratives require.

Trigg and Weiser's [23] TEXTNET organized scholarly text segments through a taxonomy of typed links, distinguishing substance links that connect nodes within a work (such as support, citation, and background) from commentary links that attach criticism or response to those nodes. Typed link taxonomies fell out of common use as general-purpose hypertext systems adopted untyped links. The present system recovers the concept through a different mechanism. It assigns formal ontology types to the entities that anchor navigation rather than to the links themselves, so traversal between views is mediated by typed identifiers in the semantic web stack.

Conklin and Begeman's [7] gIBIS captured design deliberation as typed networks of Issues, Positions, and Arguments. A Position responds-to an Issue, and an Argument supports or objects-to a Position. gIBIS demonstrated that domain-specific type systems can structure argumentative content for nonlinear traversal among distributed team members, though all content required manual authoring.

Hobbs and Potts [9] extended this approach to narrative content through the Hyperscenario framework. Their Scenario Markup Language (SCML) defines a typed decomposition schema for scenarios with structural elements (settings, characters, events, actions) and relationships between them. SCML encodes structure independently from presentation, so the same underlying scenario can be rendered in multiple representational formats.

Szekely and Kejriwal [21] developed Domain-specific Insight Graphs (DIG) under DARPA's MEMEX program. DIG performs information extraction from unstructured web content to populate knowledge graphs. Users can build customized search interfaces for specific investigative domains without programming expertise.

Edge et al. [8] introduced GraphRAG, a system that uses LLMs to extract entities and relationships from document collections, construct a knowledge graph through community detection, and generate hierarchical summaries of entity clusters. The resulting hierarchy supports both query-focused summarization and interactive browsing across summary levels. Domain-tailored prompts guide extraction, but the entity types are inferred by the LLM for each corpus.

Mumford, Atkinson, and Bench-Capon [13] found that LLMs applied to European Court of Human Rights cases struggled with both outcome prediction and structured keyword classification of case components. The finding reinforces the need for domain-specific decomposition schemas when analytical components must be individually identified and compared across cases.

Johannesson and Lindeberg [10] address a related problem in regulatory text. Their ontology-informed method reads legal provisions through LLM-assisted analysis guided by a domain ontology, structures the output into annotated tables, and renders the results as interactive concept maps designed to make complex regulations accessible to non-experts. The three-stage method parallels the decomposition-to-view architecture described here, though it currently produces a single visualization from regulatory text where the present system generates multiple linked views from professional ethics opinions.

Coordinated multiple views (CMV) address a related design problem. Wang Baldonado, Woodruff, and Kuchinsky [24] established eight design guidelines for CMV interfaces, including the rule of diversity (use multiple views when data exhibits a diversity of attributes, models, levels of abstraction, or genres) and the rule of parsimony (minimize views to reduce context-switching costs). Roberts [17] surveyed the maturation of CMV from model development toward domain-specific application. Section 3 describes how the system adapts the CMV design principles for professional ethics narratives.

Zhou et al. [25] investigated how interactive explanation interfaces affect human verification of LLM chain-of-thought reasoning. In a between-subjects study with 125 participants, they compared static text explanations against interactive formats including structured text, pseudo-code, and node-link diagrams. The visual interface significantly improved error detection over the static text baseline.

These systems address different portions of the problem space. Typed link taxonomies (TEXTNET) and argumentation tools (gIBIS) impose analytical structure but require manual authoring. Hyperscenario decomposes and links elements within a single narrative, but the schema is general-purpose. DIG applies traditional information extraction to populate investigative search interfaces. GraphRAG uses LLMs with domain-tailored prompts to support both retrieval and interactive browsing, but entity types are inferred per corpus without formal ontological constraint. LLMs alone struggle with both outcome prediction and component-level classification in legal case analysis [13]. Ontology-informed legal visualization [10] combines domain ontology with LLM-assisted extraction but currently produces a single view. CMV research [17, 24] and interactive explanation interfaces [25] establish design principles and empirical support for multi-view presentation but operate on visual encodings or reasoning traces, not on professional narrative documents.

No existing system combines ontology-constrained LLM extraction from professional narrative documents with a domain-specific decomposition schema, multi-view analytical recomposition, and cross-case navigation through shared ontology classes and expert-authored citation links. The system described in this paper combines all four.

3 Narrative Architecture

Section 3 has three parts. Section 3.1 specifies the nine-component decomposition schema and the ontology-constrained extraction process that populates it. Section 3.2 describes the eight synthesis views that recompose the components into linked analytical perspectives. Section 3.3 covers the provenance metadata, visual threading, and entity navigation mechanisms that link views to one another and to the source text.

Table 1: Nine-component decomposition schema. Components are grouped by functional dimension.

Dimension	Component	Definition
Contextual	Roles (R)	Professional positions that generate ethical obligations
Grounding	States (S)	Conditions describing circumstances at a point in time
	Resources (Rs)	Available knowledge, standards, and institutional documents
Normative	Principles (P)	Abstract ethical standards recognized by the profession
Structure	Obligations (O)	Concrete duties generated by roles and applicable principles
	Capabilities (Ca)	Actions a professional has the capacity to perform
	Constraints (Cs)	Limitations on what a professional may do
Temporal Dynamics	Actions (A)	Volitional conduct by participants
	Events (E)	Occurrences outside participant control

3.1 Nine-Component Decomposition Schema

Anderson, Anderson, and Armen [2] showed with MedEthEx that ethical decision-making becomes operationalizable when cases are decomposed before reasoning. MedEthEx scored cases along three prima facie duties from biomedical ethics and used inductive logic programming to learn decision principles distinguishing correct from incorrect actions. The present system extends the case-decomposition strategy to professional ethics opinions, producing typed components that populate linked analytical views for human-directed analysis.

The system extracts nine component types from professional ethics case text: Roles (R), Principles (P), Obligations (O), States (S), Resources (Rs), Actions (A), Events (E), Capabilities (Ca), and Constraints (Cs). These nine types synthesize two lines of work in computational ethics. Berreby, Bourgne, and Ganascia [6] proposed a modular architecture for representing ethical scenarios with distinct structural elements, and Oakley and Cocking [15] formalized how professional roles generate specific obligations that constrain ethical evaluation.

The nine component types appear across the analytical material in any given opinion. NSPE Case 24-02 (*Use of Artificial Intelligence in Engineering Practice*) names professional roles and situational conditions, cites ethical principles and the obligations they generate, references the competencies and constraints that bound participant choices, and describes a sequence of actions and precipitating events. Section 4 walks through the extraction in detail.

Table 1 groups the nine components into three functional dimensions. Contextual Grounding (R, S, Rs) establishes who is involved and under what circumstances. Normative Structure (P, O, Ca, Cs) captures the ethical requirements and capacities that apply. Temporal Dynamics (A, E) traces what happened and in what order. Section 4 illustrates each component through NSPE Case 24-02.

The extraction process operates through a LangGraph state graph that sequences three steps based on cross-component dependencies, one for each functional dimension shown in Table 1. Before each step, SPARQL queries to the domain ontology retrieve existing class definitions via Model Context Protocol (MCP) tool calls. These definitions constrain what the LLM produces. Additional MCP queries during extraction check whether proposed entities match existing ontology entries or represent novel discoveries requiring human review. The architecture is described in detail elsewhere [4]. This paper addresses the recomposition that transforms extracted components into linked analytical views.

3.2 Eight Synthesis Views as Linked Retellings

After decomposition, the nine component types populate eight synthesis views. Each view offers a different analytical perspective on the same extracted entities. Figure 2 shows the overall architecture, and Table 2 lists each view with its analytical question and content.

Wang Baldonado, Woodruff, and Kuchinsky's [24] Rule of Decomposition holds that complex data should be partitioned into multiple views to create manageable segments and to provide insight into interactions among different dimensions. Hobbs and Potts [9] argue the same for narrative content. Separating the structure of a scenario from its presentation enables multiple representational views of the same underlying content. Wang Baldonado et al. also identify parsimony as a design principle. The eight views represent the minimum set covering the analytical dimensions present in professional ethics opinions, with each answering a question that no other view addresses.

The first group presents or connects the full entity inventory. Entities displays all extracted components as a navigable graph organized by functional dimension. The Flow view links code provisions to ethical questions to conclusions, drawing edges to relevant entities across all nine types. In Provisions, Obligations and Principles are regrouped under NSPE Code Sections I, II, and III by institutional hierarchy. Precedents handles cross-case connections through two distinct link types described below.

The second group analyzes the ethical situation from complementary perspectives. Q&C presents the questions the board framed and the conclusions it reached, with analytical overlays. Each question includes emergence data identifying the triggering Events, Actions, and competing warrants that gave rise to it, and each conclusion includes resolution data identifying the determinative Principles and facts. The Decisions view organizes decision points where participants faced choices. Each decision point presents options, the board's assessment, and arguments for and against each course of action. The argumentative structure follows Toulmin's [22] model. In Timeline, Actions and Events appear in temporal sequence, drawn from Step 3 extraction. Decision points from synthesis nest beneath their corresponding timeline entries as collapsible subitems. Allen's interval algebra [1] provides the temporal relations that organize the sequence, such as before, during, overlaps, and meets.

The third group is the Narrative view, which integrates eight of the nine component types into a participant-centered retelling. Characters come from Roles, ethical tensions from Principles, Obligations, and Constraints, opening conditions from States and Resources, and narrative events and decision moments from Actions and Events. Capabilities are the exception. They serve as preconditions in the Decisions view rather than as narrative content. The Narrative organizes content around the participants and their ethical stakes, while the Timeline organizes by temporal sequence.

The inventory and mapping views reorganize extracted entities into different navigational structures. The interpretive and integrative views present content that does not appear in the original case text, produced through synthesis over the extracted components. Successive synthesis steps produce the emergence and resolution overlays in Q&C, the argumentative structure in Decisions, the causal flow in Timeline, and the character-driven account in Narrative. Synthesis is implemented through LLM calls orchestrated by LangGraph, with the ontology providing the same constraint that guides extraction. Each synthesis step receives the typed entities from prior steps as input. Human review applies to both extraction and synthesis.

The Precedents view extends the architecture beyond individual cases. Two link types connect opinions across the case base. Explicit links arise from board-authored citations. During synthesis, the system extracts references to prior opinions from the Discussion section. Each citation identifies a prior opinion the board judged relevant, the passage in which the citation appeared, the Principle for which it was invoked, and a citation type of supporting, distinguishing, analogizing, or overruling. The taxonomy follows the legal-CBR tradition of Ashley and McLaren [5]. The authority of the explicit layer derives from direct expert judgment about intertextual relevance. Implicit links arise from constrained extraction against a shared ontology. When two cases share an Obligation class such as "verify AI-generated output," the shared vocabulary creates a navigable connection. The Precedents view presents both layers with distinct visual treatment, so the provenance of each link type is immediately apparent. The retrieval methodology underlying the implicit layer is described in concurrent work [3].

Three views implement progressive disclosure following Shneiderman's "overview first, zoom and filter, then details-on-demand" [20]. The Provisions view groups provisions by NSPE Code section. Collapsed headers show provision counts and entity totals. A first expansion reveals the full provision text and a summary of applicable entity types. A second expansion into an entity type group reveals individual entities with ontology definitions. The pattern follows the layered structure familiar from legal information systems, where headnote summaries classified by topic provide scannable overviews without requiring the reader to examine each opinion in full.

The Decisions view implements three levels of progressive disclosure. Each decision point header shows a focus identifier, description preview, option count, and board choice indicator. The first expansion presents the options, the board's choice, and a summary of the argumentative structure. The second discloses the full Toulmin argument with claim, warrant, backing, data, rebuttal, and qualifier. This three-level hierarchy parallels the typed deliberation structure in Conklin and Begeman's gIBIS [7], where Issues contain Positions that are supported or opposed by Arguments.

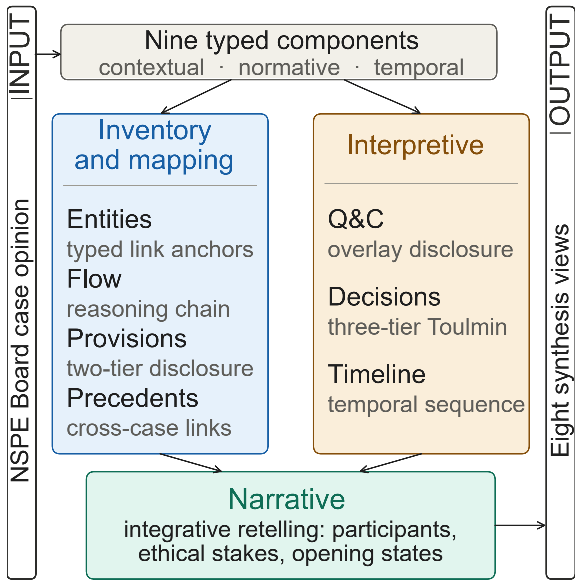

Figure 2: Architectural overview of the eight synthesis views. Nine typed components from the three-step ontology-constrained extraction process populate views organized into three groups. Citation labels beneath each group identify the design and reasoning traditions the views draw on.

Table 2: Eight synthesis views organized by group. Each view answers a distinct analytical question using the same extracted entities. Inventory and mapping views reorganize extracted entities; interpretive and integrative views present additional analytical content produced through synthesis.

Group	View	Analytical Question	Content
Inventory &	Entities	What components were extracted?	Graph of all nine types with filters, search, definitions, ontology links
Mapping	Flow	How does the board's reasoning chain proceed?	Graph linking code provisions to questions to conclusions with entity connections
	Provisions	Which code sections apply?	Code section groups with progressive disclosure to entity type groupings
	Precedents	What other cases relate?	Explicit board citations and implicit similarity-based matches through shared ontology classes
Interpretive	Q&C	What questions arose and how were they resolved?	Questions with emergence overlays linked to conclusions with resolution overlays
	Decisions	What choices existed and what arguments support each?	Causal-normative links, decision points with options, and pro/con arguments
	Timeline	What happened in what order?	Actions and Events in temporal sequence from Step 3 extraction, with synthesis decision points nested hierarchically and causal flow
Integrative	Narrative	Who was involved and what was at stake?	Characters with roles, ethical tensions, opening states

Q&C presents questions and conclusions as compact headers with count badges. Expanding a question opens emergence overlays naming the triggering events, actions, and competing warrants. Expanding a conclusion opens resolution overlays naming the determinative Principles and facts. Collapsed headers in all three views provide enough information for readers to assess relevance before expanding. Pirolli and Card [16] call this information scent. The principle connects progressive disclosure to reader agency across the architecture.

The coordination mechanism linking these views is shared entity identity. An Obligation entity that appears in the Entities view is the same entity referenced in the Provisions, Q&C, and Decisions views. Persistent entity identifiers maintain the link across views, so a reader can follow any component through appearances in different analytical contexts.

3.3 Provenance, Visual Threading, and Entity Navigation

Every extracted entity has attached provenance metadata recording the source text span in the original case narrative, the extraction step that produced it, a confidence classification of ontology match or novel proposal, and a human review status of approved, modified, or rejected. This metadata creates a traceable chain from the original document through decomposition to synthesis. The chain functions analogously to a web annotation layer [18], where annotations reference specific text passages and carry structured metadata. A reader examining an Obligation in the Provisions view can trace it to the sentence in the original case narrative where the board identified that Obligation, the ontology definition that constrained the classification, and the expert review decision that approved or modified it. Three step types organize the extraction history. Readers can filter by contextual, normative, or temporal entries and expand any record to inspect the full provenance.

Every entity mention across all eight views renders as an interactive inline label. The label displays the entity name in the canonical component color. Hovering opens a popover with the component type badge, extraction step badge, ontology definition, and entity URI. A blue border on a label marks an ontology-sourced entity whose class exists in the shared vocabulary and may appear in other cases. Plain labels mark case-local entities specific to a single case.

Clicking an entity label opens the full ontology definition and class hierarchy for that entity. The click also triggers a cross-view search that locates the same entity across the other synthesis views and presents navigation options. A reader examining an Obligation in the Provisions view can follow a link to the same Obligation in the Q&C view, the Decisions view, or the Precedents view, which lists other cases sharing the same Obligation class. This cross-view navigation mechanism implements what Wang Baldonado et al. [24] call self-evidence, where relationships among views become apparent through direct interaction.

Each component type uses a canonical color assigned during extraction and preserved through all synthesis views. A reader tracking "orange = Principles" recognizes that visual thread from extraction through review and into synthesis without consulting documentation. The consistent color scheme creates implicit typed links across views, an instance of the Rule of Consistency identified by Wang Baldonado et al. [24].

The ontology evolves through use. Novel classes approved through human review enter the shared ontology and become available for all subsequent cases. The analytical vocabulary grows as the case base expands, and new connections between cases emerge as shared classes accumulate.

Section 4 walks through these architectural features using Case 24-02.

4 Walkthrough Demonstration

This section illustrates the architecture of Section 3 through NSPE Board of Ethical Review Case 24-02, *Use of Artificial Intelligence in Engineering Practice* [14]. The case concerns an engineer who used AI language processing software to draft a technical report and AI-assisted drafting tools to generate preliminary engineering

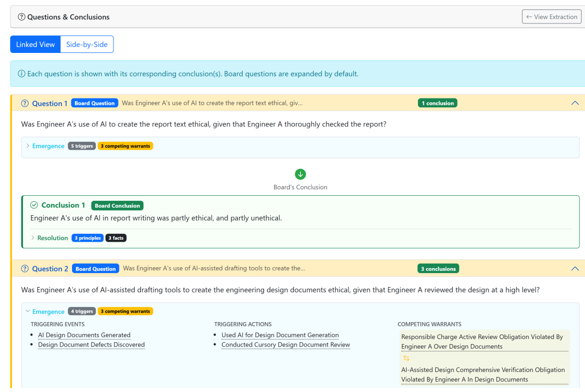

Figure 3: Q&C view in linked mode for Case 24-02. Question 1 (top) is paired with its board conclusion through a green arrow labeled "Board's Conclusion." Count badges on the Emergence and Resolution rows summarize the triggers, competing warrants, determinative principles, and supporting facts attached to the question and conclusion. Below, Question 2's Emergence row is expanded to show the full triggering events, triggering actions, and competing warrants.

design documents. The engineer conducted a thorough review of the report but only a cursory review of the design documents. The board evaluated three questions about the ethical permissibility of these actions under the NSPE Code of Ethics.

Case 24-02 implicates all nine ethical components across the three functional dimensions. Four professional roles appear, situational conditions shift across the case, and a temporal sequence of early actions constrains later ethical evaluations. The board reached mixed conclusions, finding some uses of AI ethical and others unethical. The distinction rests on the engineer's level of independent professional judgment in each context.

4.1 Synthesis Views in Practice

The system processed Case 24-02 through three sequential steps [4]. Extraction produced 213 entities across all nine components, and synthesis produced the analytical content for the eight views described in Section 3.2. The walkthrough below covers five views: Q&C, Timeline, Decisions, Narrative, and Precedents. These five span the inventory, interpretive, and integrative groups and illustrate progressive disclosure, cross-view entity coordination, and cross-case navigation.

Questions and Conclusions. Figure 3 shows the Q&C view in linked mode, where each question appears with its corresponding conclusions. The board posed three questions about the ethical permissibility of the engineer's actions. Synthesis decomposed these into 21 analytical questions, each paired with triggering events, triggering actions, and competing warrants. It also produced 28 conclusions linked to the questions with the basis for each. Question 1, for example, asks whether the engineer's use of AI to create report text was ethical, given that the engineer thoroughly checked the report. Five triggers and three competing warrants are summarized as count badges on the Emergence row, which is collapsed in the figure but expandable on demand. Below Question 1, Conclusion 1 finds the engineer's use of AI in report writing partly ethical and partly unethical. Three determinative Principles and three supporting facts are summarized as count badges on the Resolution row, expandable on demand to reveal the full entries.

Question 2 turns to AI use in design and asks whether the engineer's use of AI-assisted drafting tools to create engineering design documents was ethical, given that the engineer reviewed the design at a high level. This question links to three conclusions. Board Conclusion 2 finds the use of AI-assisted drafting tools not unethical per se. Analytical Conclusion 5 qualifies the board's permissive finding with a competence threshold tied to Code provisions I.2 and II.2.a.

Board Conclusions are direct determinations by the board, while Analytical Conclusions are inferences drawn from the board's analysis during synthesis. Type badges adjacent to each conclusion header make the distinction visible. Emergence overlays surface the triggers and competing warrants for each question. Resolution overlays surface the determinative Principles and facts for each conclusion. Readers can expand any overlay to inspect the linked entries.

In Figure 3, each question is paired with the conclusions that answer it. A reader can navigate downward from question to conclusions, or upward from a conclusion to the question it resolves.

Timeline. The Timeline view (Figure 4) presents the case as a temporal sequence drawn from Step 3 extraction. Entries are classified as Actions or Events and arranged with sequence numbers and linked entity badges. For Case 24-02, the sequence spans thirteen entries covering the engineer's conduct and the situational events that preceded and followed each decision. Entry 6 ("Used AI for Design Document Generation") is expanded in the figure to reveal three nested decision points, each with available options, the board's choice, and a link to the full argumentative treatment in the Decisions view. A collapsible Causal Flow section below the sequence traces `enables` links between Actions and the Events they produced.

Decisions and Arguments. The Decisions view (Figure 5) presents the same case from an argumentative perspective. Six causal-normative links map the engineer's Actions to the Obligations they fulfill or violate. "Used AI for Design Document Generation" fulfills zero Obligations and violates nine. "Conducted Thorough Report Review" fulfills three and violates two. Below the causal-normative links, the view shows 18 decision points where participants faced choices among alternatives. A decision point header includes a description preview, the number of options available, the board's choice, and a confidence percentage. Expanding a decision point reveals the three options with the board's choice marked, and a further expansion discloses the full Toulmin argument described in Section 3.2. Collapsed headers at both tiers let readers scan rapidly before committing to any analytical thread. Figure 5 also shows progressive disclosure on the Causal-Normative Links side, where "Input Confidential Data into Public AI" is expanded to show the specific Obligations that action fulfills and violates.

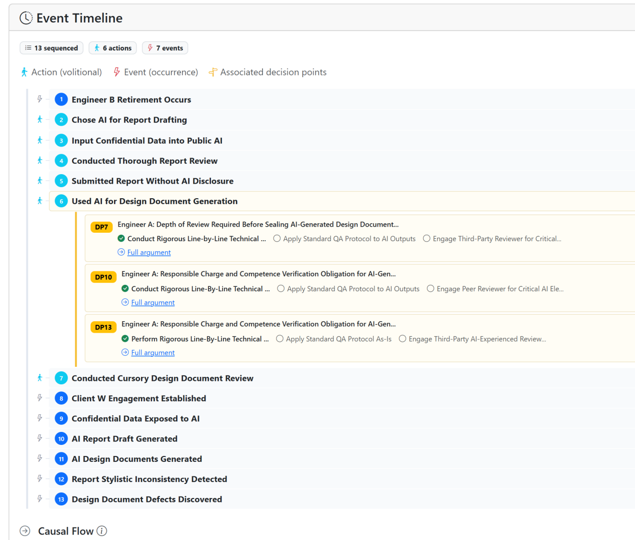

Figure 4: Timeline view for Case 24-02. Thirteen entries drawn from Step 3 temporal extraction appear in sequence, with Actions (volitional conduct) and Events (occurrences outside agent control) distinguished by icon. Entry 6 ("Used AI for Design Document Generation") is expanded to reveal three nested decision points (DP7, DP10, DP13), each showing available options with the board's choice marked by a green check and a "Full argument" link to the Decisions view. The Causal Flow section header at the bottom traces `enables` links between actions and events on demand.

Narrative. The Narrative view (Figure 6) organizes the case around participants and their ethical stakes. For Case 24-02, eight characters drawn from extracted Roles appear with profiles and protagonist or stakeholder badges. Each character represents a facet of an underlying professional position. Engineer A, for example, appears in four Roles that range from environmental consultant to engineer in responsible charge. Sixteen ethical tensions drawn from Obligations and Constraints appear in the Ethical Tensions section below the Characters view as paired entity badges identifying competing duties. An opening context paragraph drawn from States establishes the circumstances at the start of the case in second-person address. The Timeline sequences Actions and Events chronologically. The Narrative groups the same entities by participant and ethical stake.

Precedents. Case 24-02 has two board-cited precedents (Figure 7), both extracted from the Discussion section of the opinion. Case 98-3, *Use of CD-ROM for Highway Design*, carries a "distinguishing" badge marking a relevant difference. Case 90-6, *Use of CADD System*, carries an "analogizing" badge marking a relevant similarity.

The lower section displays ten similar cases identified through shared ontology classes and embedding similarity [3]. Case titles indicate thematic clusters such as competence in design services, signing and sealing plans not prepared by the engineer, and professional competence in structural design. One case, *Use of CADD System*, appears in both layers and carries a "cited" badge in the

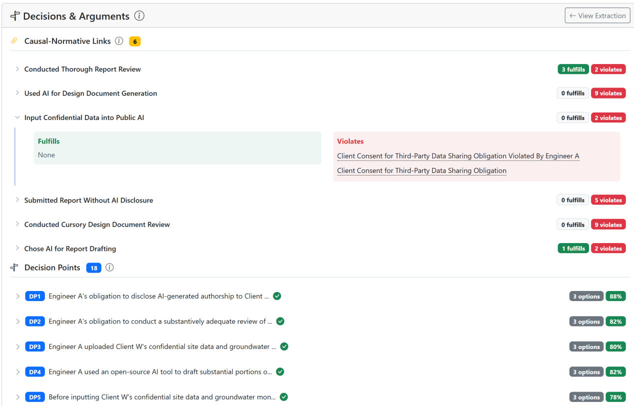

Figure 5: Decisions view for Case 24-02 showing two-tier progressive disclosure in the Causal-Normative Links section. Six causal-normative links (top) map the engineer's actions to the obligations they fulfill or violate, with fulfills and violates counts shown on each collapsed header. "Input Confidential Data into Public AI" is expanded to reveal a side-by-side Fulfills and Violates panel naming the specific obligations on each side. Below, the Decision Points section with an "18" badge lists decision point headers (DP1 through DP5 visible), each carrying an option count and a board-choice confidence percentage.

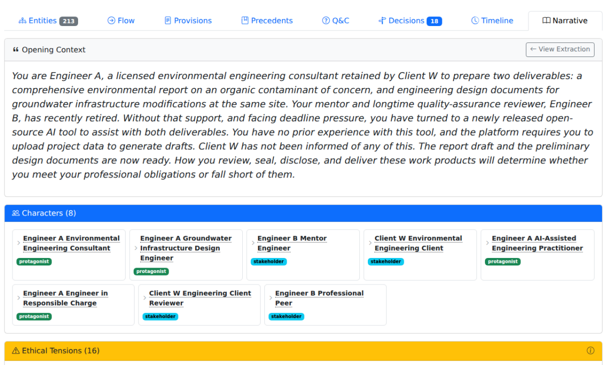

Figure 6: Narrative view for Case 24-02. The Opening Context paragraph (top) establishes the starting circumstances of the case in second-person address. Below, the Characters section displays all eight character profiles drawn from extracted Roles, with protagonist badges on Engineer A's four facets and stakeholder badges on the other four. The Ethical Tensions header at the bottom indicates sixteen paired tensions (cropped from this view).

implicit list. The overlap confirms that the computed similarity aligns with the board's own judgment for that case.

Tracing an Entity Across Views. The Obligation "verify AI-generated output" was extracted in Step 2 as an instance of an existing ontology class linked to NSPE Code Section II.2.a. This Obligation appears across four views. In Provisions, it sits under Section II

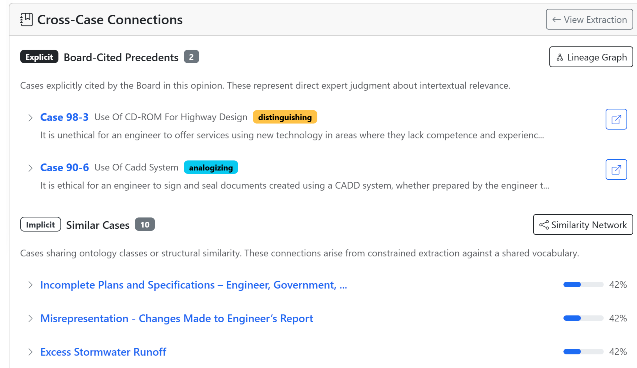

Figure 7: Cross-Case Connections for Case 24-02. Board-Cited Precedents (Explicit, top) carry citation-type badges (distinguishing, analogizing). Similar Cases (Implicit, below) are identified through shared ontology classes. *Use of CADD System* appears in both layers and carries a "cited" badge in the implicit list.

(Rules of Practice) alongside other competence-related Obligations. In Q&C, it is a determinative factor in Conclusion 1. In Decisions, it is one of the nine Obligations violated by "Used AI for Design Document Generation." In Precedents, it links Case 24-02 to cases where the board addressed verification duties in technology-dependent practice, including Case 98-3.

4.2 Cross-Case Navigation

The reader can click "Case 98-3" in the cited precedents list to open the full analysis of that case, including a nine-component decomposition and a Precedents view with further citations. Each step follows a board-authored link, building a citation chain analogous to following footnotes in a legal opinion. The similar cases list offers a second path. An expanded entry shows the shared ontology classes that produced the similarity score. If Case 24-02 and a similar case both contain the Obligation class "verify AI-generated output," a click on that shared class navigates to the entity in the other case, where the same Obligation appeared in a different factual context.

5 Discussion and Future Work

5.1 Contributions in Hypertext Terms

The architecture presented in this paper instantiates a hypertext operation on professional ethics narratives. Linear board opinions, written as continuous prose for sequential reading, are decomposed into nine typed components and recomposed through eight linked synthesis views. The result is a multi-path structure where readers choose which analytical perspective to pursue and at what depth.

The eight views differ from coordinated multiple views in information visualization. CMV systems present different visual encodings of a shared dataset, coordinated through linking and brushing [24]. The synthesis views offer analytical retellings of the same case organized along different principles. Provisions follows institutional code hierarchy, Timeline follows temporal sequence and causal flow, and Decisions follows deliberative structure. Each view answers a question no other view addresses, and the coordination mechanism is shared entity identity across all eight views.

The Precedents view contributes a cross-case linking architecture with two layers of distinct epistemic weight. One layer is authoritative, drawn from board-authored citations. The other is computed from shared ontology classes. The separate visual treatment of the two layers allows readers to distinguish between connections that the board itself considered relevant and connections that emerge from shared analytical vocabulary. This two-layer structure extends the typed linking tradition of Trigg and Weiser's [23] TEXTNET and Conklin and Begeman's [7] gIBIS to cross-document intertextual navigation in a professional domain.

Progressive disclosure across Provisions, Decisions, and Q&C views places depth under reader control. Collapsed headers provide information scent for assessing relevance before expansion. Readers decide how far to follow any thread. The ontology-constrained extraction process and synthesis produce the analytical structures, human experts review them, and readers navigate them. The architecture does not generate ethical judgments about the conduct described in the cases.

5.2 Limitations

The current implementation processes engineering ethics cases from a single professional domain (NSPE Board of Ethical Review). Extension to a domain such as legal ethics requires a domain-specific ontology and a corpus of published opinions with consistent internal structure. The nine-component decomposition schema is designed to generalize across professional domains, but it has not been tested outside engineering ethics.

The case base contains 119 cases with section-level processing. Cross-case linking through ontology classes depends on the breadth of the shared vocabulary, and a small case base produces sparse implicit links in the Precedents view. As the case base grows and the ontology accumulates entity classes through extraction and human review, the density and utility of cross-case connections should increase.

Section 4 demonstrates the architectural claims in Section 3 through a single case.

5.3 Future Directions

The nine-component representation and eight-view architecture provide infrastructure for interactive scenario generation. The Event Calculus [11, 19] fluent tracking already present in the Timeline view supports counterfactual exploration. Alternative choices at a decision point would generate scenarios that trace consequences through obligation activation and constraint violation. A comparison of these scenarios against the board's conclusions would let readers evaluate the alternative courses of action.

Extension to legal ethics opinions from bar association committees, which follow internal structures comparable to NSPE BER opinions, would test the generality of the nine-component decomposition schema. A concurrent study validates the component-aware retrieval methodology underlying the Precedents view [3]. Formal evaluation with professional ethics practitioners would complement the retrieval validation with evidence about how readers use cross-case navigation and progressive disclosure during case analysis.

During the preparation of this work, the authors used Claude (Anthropic) to assist with manuscript editing. The authors reviewed and edited all content and take full responsibility for the publication's content.

References

[1] James F. Allen. 1983. Maintaining Knowledge About Temporal Intervals. *Commun. ACM* 26, 11 (1983), 832–843. doi:10.1145/182.358434
[2] Michael Anderson, Susan Leigh Anderson, and Chris Armen. 2006. MedEthEx: A Prototype Medical Ethics Advisor. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*. AAAI Press, 1759–1765.
[3] [Anonymous]. 2026. Component-Aware Case Retrieval for Professional Ethics with Ontology-Constrained LLM Extraction. In *International Conference on Case-Based Reasoning*. Springer, Cham, Switzerland. Under review.
[4] [Anonymous]. 2026. ProEthica: A Professional Role Based Ethical Analysis Tool Using LLM-Orchestrated, Ontology Supported Case Based Reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 40. AAAI Press, Washington, DC, USA, 41673–41675.
[5] Kevin D. Ashley and Bruce M. McLaren. 1995. A CBR Knowledge Representation for Practical Ethics. In *Advances in Case-Based Reasoning*. Springer, 180–197.
[6] Fiona Berreby, Gauvain Bourgne, and Jean-Gabriel Ganascia. 2017. A Declarative Modular Framework for Representing and Applying Ethical Principles. In *16th Conference on Autonomous Agents and MultiAgent Systems*. IFAAMAS, São Paulo, Brazil, 96–104.
[7] Jeff Conklin and Michael L. Begeman. 1988. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. In *Proceedings of the 1988 ACM Conference on Computer-Supported Cooperative Work (CSCW '88)*. ACM, 140–152. doi:10.1145/62266.62278
[8] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2025. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. *arXiv preprint arXiv:2404.16130* (2025). doi:10.48550/arXiv.2404.16130
[9] Reginald L. Hobbs and Colin Potts. 2000. Hyperscenarios: A Framework for Active Narrative. In *Proceedings of the 38th Annual ACM Southeast Conference (ACMSE '00)*. ACM, 201–210. doi:10.1145/1127716.1127760
[10] Paul Johannesson and Jöran Lindeberg. 2026. Ontology-Informed Design of Legal Visualisations. In *Conceptual Modeling*, Dominik Bork, Roman Lukyanenko, Shazia Sadiq, Ladjel Bellatreche, and Oscar Pastor (Eds.). Springer Nature Switzerland, Cham, 341–357. doi:10.1007/978-3-032-08623-5_18
[11] Robert Kowalski and Marek Sergot. 1986. A Logic-Based Calculus of Events. *New Generation Computing* 4 (1986), 67–94.
[12] Bruce M. McLaren. 2003. Extensionally Defining Principles and Cases in Ethics: An AI Model. *Artificial Intelligence* 150, 1–2 (2003), 145–181. doi:10.1016/S0004-3702(03)00135-8
[13] Jack Mumford, Katie Atkinson, and Trevor Bench-Capon. 2024. Unravelling the ECHR: Components of Legal Case Analysis. In *Legal Knowledge and Information Systems (JURIX 2024) (Frontiers in Artificial Intelligence and Applications, Vol. 395)*. IOS Press, 107–118. doi:10.3233/FAIA241238
[14] National Society of Professional Engineers. 2025. Board of Ethical Review Case Archive. https://www.nspe.org/career-growth/ethics/board-ethical-review-cases. Accessed March 2026.
[15] Justin Oakley and Dean Cocking. 2001. *Virtue Ethics and Professional Roles*. Cambridge University Press, Cambridge, UK.
[16] Peter Pirolli and Stuart K. Card. 1999. Information Foraging. *Psychological Review* 106, 4 (1999), 643–675. doi:10.1037/0033-295X.106.4.643
[17] Jonathan C. Roberts. 2007. State of the Art: Coordinated & Multiple Views in Exploratory Visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007)*. IEEE, 61–71. doi:10.1109/CMV.2007.20
[18] Robert Sanderson, Paolo Ciccarese, and Benjamin Young. 2017. *Web Annotation Data Model*. W3C Recommendation. W3C. https://www.w3.org/TR/annotation-model/
[19] Murray Shanahan. 1999. The Event Calculus Explained. In *Artificial Intelligence Today: Recent Trends and Developments*, Michael J. Wooldridge and Manuela M. Veloso (Eds.). Lecture Notes in Computer Science, Vol. 1600. Springer, 409–430. doi:10.1007/3-540-48317-9_17
[20] Ben Shneiderman. 1996. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proceedings of the IEEE Symposium on Visual Languages*. IEEE, 336–343. doi:10.1109/VL.1996.545307
[21] Pedro Szekely and Mayank Kejriwal. 2018. Domain-specific Insight Graphs (DIG). In *Companion Proceedings of The Web Conference 2018 (WWW '18)*. ACM, 433–434. doi:10.1145/3184558.3186203

[22] Stephen E. Toulmin. 2003. *The Uses of Argument* (updated ed.). Cambridge University Press, Cambridge, UK.
[23] Randall H. Trigg and Mark Weiser. 1986. TEXTNET: A Network-Based Approach to Text Handling. *ACM Transactions on Information Systems* 4, 1 (1986), 1–23. doi:10.1145/5401.5402
[24] Michelle Q. Wang Baldonado, Allison Woodruff, and Allan Kuchinsky. 2000. Guidelines for Using Multiple Views in Information Visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '00)*. ACM, 110–119. doi:10.1145/345513.345271
[25] Runtao Zhou, Giang Nguyen, Nikita Kharya, Anh Totti Nguyen, and Chirag Agarwal. 2026. Improving Human Verification of LLM Reasoning through Interactive Explanation Interfaces. In *Proceedings of the 31st International Conference on Intelligent User Interfaces (IUI '26)*. ACM, 456–473. doi:10.1145/3742413.3789134