Supporting Professional Conference Organizers in Expert-Finding Tasks using Knowledge Graphs and Topic Modeling

Benedict Hartmann[1,*], Nico Lindemann[1]

[1]University in Hagen, Universitätsstr. 1, 58097 Hagen, Germany

Abstract

Professional Congress Organizers (PCOs) require transparent, evidence-based support to manage large collections of conference submissions, identify thematic structures, and design coherent sessions. This paper presents a lifecycle-aware prototype system that integrates Transformer-based topic modeling with knowledge graph–based semantic organization to support knowledge management and expert-finding–driven session planning in the Meetings, Incentives, Conferences, and Exhibitions (MICE) domain. The system combines configurable BERTopic-based thematic extraction with RDF-based knowledge graph persistence, complemented by organizer-facing browsing, visualization, and retrieval interfaces. The system is evaluated using a mixed-method approach: a Cognitive Walkthrough with domain experts assesses workflow transparency and traceability while revealing usability challenges at larger topic scales, and a quantitative evaluation on a real-world PCO dataset (7,604 abstracts) and the NeurIPS benchmark corpus (7,241 papers) measures topic coherence, diversity, and stability. Results show performance comparable to established baselines while qualitative evaluation exposes usability limitations of the prototype application to serve as input for future work.

Keywords

Professional Congress Organization, Topic Modeling, Knowledge Graphs, Information Extraction

1. Introduction

Meetings, conferences, and events remain central mechanisms for knowledge exchange and community building across domains such as academia and information technology [1, 2, 3]. Within the tourism sector, these activities are commonly attributed to the Meetings, Incentives, Conferences, Events (MICE) industry [4], which has been reported as a large global market (e.g., $645.7 billion in 2022) [5]. Research conferences form a substantial and particularly information-intensive subset of MICE [6]: they orchestrate large volumes of submissions, reviews, scheduling decisions, and post-event feedback under tight deadlines and quality constraints. *This scale and time pressure give rise to the requirement for systematic, tool-supported handling of conference data that enables organizers to manage thematic structure and expertise evidence consistently and efficiently across the conference lifecycle.* Over the last years, MICE has been shaped by several structural drivers. The COVID-19 pandemic accelerated virtual and hybrid event formats [7], while climate-related concerns intensified attention to sustainable and "green" meeting practices [8, 9]. At the same time, the Internet, the Web, and social media increased the availability and heterogeneity of digital traces produced before, during, and after events (e.g., submissions, proceedings, program metadata, recordings, slides, Q&A transcripts, and survey feedback). In parallel, advances in Artificial Intelligence (AI)—particularly Transformer-based architectures [10]—have improved the ability to learn semantic representations from complex data and to support natural language understanding, prediction, and decision support [11]. Collectively, these drivers amplify both the need and the opportunity for data- and evidence-driven support in conference organization, *thereby requiring systems that integrate heterogeneous and evolving conference data and translate it into reusable, decision-relevant*

Proceedings of TRI-IVIS 2026, the First International Workshop on Trustable, Reproducible, and Intelligent Information Visualization Systems, June 2026, Venice, Italy.

benedict.hartmann@fernuni-hagen.de (B. Hartmann); nico.lindemann@fernuni-hagen.de (N. Lindemann)
0000-0001-8932-5853 (B. Hartmann); 0009-0006-8232-6240 (N. Lindemann)

knowledge artifacts. Conference organization involves multiple actors and interdependent tasks. Key actors include the Professional Congress Organizers (PCOs), authors, reviewers, chairs, panelists, and attendees. A PCO is responsible for acquiring and coordinating contributions, ensuring review quality, organizing sessions, and staffing roles that directly affect the perceived quality and fairness of the event. A recurring challenge in this context is identifying suitable experts for roles such as session participants, reviewers, chairs, or panelists. Experts can be characterized as individuals with outstanding knowledge and capabilities in particular topics, typically evidenced through publications, contributions, and community engagement. The process of identifying and selecting such individuals is commonly referred to as Expert Finding (EF) [12, 13, 14]. Importantly, EF decisions require evidence: organizers need to integrate information about people, topics, contributions, and relationships, and must often justify assignments (e.g., reviewer selection) under constraints such as conflicts of interest, interdisciplinarity, and time pressure. *This creates the requirement for EF support that not only retrieves candidates, but also makes underlying evidence paths transparent and inspectable for organizer decision-making.* From a systems perspective, EF in MICE is an information-intensive decision process that depends on the systematic transformation of raw, fragmented, and evolving data into actionable and inspectable knowledge. The Information Extraction Lifecycle (IELC) reference model by Tamla [15] provides a suitable conceptual framing for such transformations. The IELC distinguishes the stages of **Data Collection**, **Data Preparation**, **Information Extraction (IE)Information Extraction (IE)**, **Information and Knowledge Organization (IKO)**, and **Information Access and Retrieval (IAR)**. Data Collection gathers raw data from heterogeneous sources; Data Preparation cleans, aligns, and structures it for analysis; IE extracts structured information from unstructured or semi-structured inputs [16]; Information Organization structures extracted content to enable search [17, 18]; Knowledge Organization provides semantic representations (e.g., taxonomies, ontologies, semantic networks) that preserve relationships and meaning; and Information Access & Retrieval (IR) focuses on the representation, storage, and retrieval of information [19]. For MICE, these stages are not optional engineering details but determine whether expert-related evidence remains consistent, reusable, and auditable across event editions and changing tool landscapes. *Consequently, MICE EF systems must be designed in a lifecycle-aware manner that explicitly supports all IELC phases rather than isolated analytic steps.* Current conference organization commonly relies on Conference Management Systems (CMS) to manage workflows and content [20]. In parallel, Expert Finding Systems (EFS) support expert ranking and retrieval, typically based on textual evidence [13]. Two technical capabilities are particularly relevant for strengthening MICE EF along the IELC: Topic Modeling (TM) and Knowledge Graphs (KG). TM is an unsupervised technique to uncover latent thematic structure in large corpora, enabling scalable semantic analysis and clustering [21]. KGs represent entities and relations as nodes and edges, supporting semantic integration, traceable reasoning, and interoperability across heterogeneous sources [22]. In combination, topic structures can provide an operational "thematic backbone" for session planning, while knowledge graphs can preserve expert evidence as navigable relationships between authors, submissions, topics, and event context. *This implies the requirement for integrated TM and KG support that treats extracted topics and relationships as persistent, reusable artifacts for organizer-facing workflows.* Despite these advances, MICE PCOs continue to face rising operational costs and increasing data volume and heterogeneity, which challenge effective and efficient session planning and author management [8]. Concrete decisions such as reviewer assignment or panel composition increasingly require integrating heterogeneous evidence: public expert metadata (e.g., affiliations, publication records), event metadata (e.g., submissions, proceedings, program structure), and event execution data (e.g., survey results, Q&A transcripts). These evidence sources differ in format and quality, evolve over time, and frequently contain unresolved references. As a result, organizers must perform labor-intensive manual synthesis, and decision-making can become opaque, inconsistent, or hard to justify. *This establishes the requirement for systems that reduce manual effort by providing unified, explainable access to expert-related knowledge grounded in systematically organized data.*

Several prior research and innovation projects motivate this research as they have addressed individual phases of the IELC in adjacent domains, yet gaps remain when these contributions are transferred to MICE EF. In practice, observations in real-world PCO environments further highlight that conference data is fragmented across multiple platforms, subject to schema changes, and not readily available as a unified evidence base for expert-related decisions.

Data Collection, Preparation, and IE. In the early IELC stages, PCO workflows depend on multiple data sources (e.g., separate databases for authors and submissions) and must frequently cope with system migrations and legacy data unavailability [23]. In addition, PCOs must integrate multiple APIs with custom schemas that require deep domain knowledge [24], while unresolved entity references across sources complicate alignment. These conditions directly impact downstream extraction: topic identification from submissions is essential for session planning, yet reading and synthesizing large numbers of submissions is time-consuming, motivating automated IE support. The **Realizing an Applied Gaming Eco-system (RAGE)** project (2015–2019) developed the Content and Knowledge Management Ecosystem Portal (KM-EP) as a collaborative content management system supporting document management and data retrieval from remote Internet sources (e.g., GitHub) [15]. However, RAGE did not support collection and management of MICE-specific data and metadata from conference and research databases, nor did it address the conference-specific sources required for EF. Similarly, **SNERC** and **Framework-Independent Toolkit for Named Entity Recognition (NER) (FIT4NER)** applied NER to support domain experts in IE through automated entity identification in applied gaming and medical contexts. FIT4NER enabled framework-independent NER development and comparative training across Stanford CoreNLP, spaCy, and Hugging Face Transformers, including format conversion utilities and reproducible runs. However, these projects were not tailored to conference EF: they did not address extraction of MICE-specific entities and relations (e.g., authorship and co-authorship), did not operationalize vocabulary development for thematic abstractions (e.g., stable topic labels), and did not consider multimodal conference artifacts (e.g., slides, recordings) that increasingly shape event knowledge. **Motivating Statement 1 (MS1):** Integrated, lifecycle-aware support is required to ingest heterogeneous and evolving MICE data and to perform conference-specific information extraction—particularly topic modeling and entity relationship extraction—so that thematic structures can be reused for session planning.

Information & Knowledge Organization. Beyond extraction, PCOs require representations that preserve relationships between entities such as authors, submissions, sessions, and editions. This includes entity resolution across sources (e.g., authors appearing under variants of names; sessions instantiated across times) and maintaining relationship integrity despite changing system architectures during event transitions. These requirements extend beyond conventional content storage: they call for robust semantic structures that enable traceable evidence paths from experts to contributions and topics. The KGen4MICE initiative aimed to support MICE PCOs via a smart service platform for conference organization and execution tasks [25], exploring aggregation of expert data (e.g., author names, transcripts, abstracts) to seed semantic structures for event management and EF. However, robust MICE-specific ontologies and knowledge-graph representations remained underdeveloped, limiting the ability to preserve and query expert evidence as stable relationships. **Motivating Statement 2 (MS2):** MICE expert finding requires semantic representations that preserve relationships between authors, submissions, and topics across changing system architectures, motivating the use of robust MICE-specific ontologies and knowledge graphs.

Information Access & Retrieval. Finally, even well-structured knowledge is only valuable if it can be accessed by organizers through usable and trustworthy interfaces. In practice, PCOs require exploratory access patterns (e.g., faceted filtering by topic, track, affiliation, or time), overviews to triage large submission sets, and interaction paths that connect high-level thematic views to actionable details about submissions and expert profiles. Retrieval mechanisms that do not provide transparent provenance (i.e., which evidence supports a recommendation) reduce trust and hinder justification of decisions. Prior work in other domains has explored shared platform architectures and modular AI toolchains. For instance, **AI4H3** and **Cloud-based Information Extraction (CIE)** studied cloud-based architectures integrating Transformer-based Machine Learning (ML) techniques and NER for medical text analysis and model training [26, 27, 28]. CIE operationalized cloud-based NER development on AWS and Azure, supporting dataset onboarding into managed storage and exposing trained models via APIs and UI. The **KMEP4NCoEGPS** aggregated scholarly articles with PDF metadata analysis and TM for thematic groupings, producing static visualizations. While these efforts contribute architectural and methodological building blocks for Information Access & Retrieval, they do not directly support MICE PCOs in accessing and inspecting expert-related knowledge for EF (e.g., topic relationships between submissions and authors, explainable evidence paths, and organizer-facing exploration). **Motivating Statement 3 (MS3):** Organizer-facing mechanisms are required that provide transparent, usable, and provenance-aware access to expert-related knowledge, enabling PCOs to explore topics, inspect relationships, and justify expert-related decisions. Taken together, existing work demonstrates the feasibility of collaborative knowledge management platforms and AI techniques across domains such as gaming, healthcare, and environmental analytics. However, their coverage of the IELC remains incomplete when transferred to MICE EF. Prior systems typically address individual lifecycle phases in isolation: some focus on data ingestion and storage without supporting MICE-specific sources and schemas; others advance information extraction techniques without modeling conference-relevant entities, relationships, or stable vocabularies; knowledge-organization approaches often remain conceptual or lack robust ontological and graph-based implementations; and access and retrieval mechanisms frequently do not provide schema-aware, provenance-grounded, and organizer-facing interaction patterns suitable for expert-related decision-making. As a consequence, PCOs still lack integrated support for tracing, inspecting, and justifying expert-finding decisions across the full conference lifecycle. Addressing these limitations requires a comprehensive, end-to-end treatment of the IELC tailored to MICE expert finding. Such a system must support the ingestion and preparation of heterogeneous conference data; perform conference-specific information extraction, including automated topic modeling and entity relationship extraction; organize extracted evidence in flexible, relationship- preserving semantic structures such as ontologies and knowledge graphs; and provide organizer-facing access and retrieval mechanisms that enable exploration, inspection, and justification of expert recommendations. Only by integrating these lifecycle stages can expert-finding support move beyond isolated analytic tools toward reusable, auditable, and decision-relevant knowledge infrastructures for conference organization.

Accordingly, this paper addresses the following research questions: **RQ1:** How can MICE PCOs be supported in knowledge management and session planning through a lifecycle-aware system that integrates heterogeneous conference data with automated thematic extraction? **RQ2:** How can Transformer-based topic modeling and knowledge graphs be combined to provide reusable, explainable, and organizer-facing access to expert-related knowledge in MICE settings?

To address these research questions, this work follows the systems development methodology proposed by Nunamaker et al. [29], which structures research through iterative phases of observation, theory building, systems development, and experimentation. Within this framework, the study derives a lifecycle-aware system design grounded in the state of the art, implements a prototype that integrates topic modeling and knowledge-graph–based organization, and evaluates the resulting system with respect to both technical quality and organizer-facing utility. The evaluation assesses the quality, stability, and interpretability of extracted thematic structures as well as the usability, transparency, and inspectability of the knowledge made available to PCOs for session planning and expert-finding tasks. The remainder of this paper is structured as follows. Section 2 reviews the state of the art in science and technology along the phases of the IELC and identifies remaining challenges for MICE expert finding. Section 3 presents the modeling and design of the proposed system, including use cases, information models, and architectural decisions. Section 4 describes the implementation of the prototype and its realization of the proposed design. Section 5 reports on the qualitative and quantitative evaluation of the system. Finally, Section 6 concludes the paper and outlines directions for future work.

2. State of the Art in Science and Technology

This section reviews the state of the art along the phases of the IELC in the MICE context and synthesizes remaining challenges that motivate the system design and evaluation.

Data Collection, Data Preparation, Information Extraction. Data collection in MICE environments is frequently affected by access and continuity problems, for example when legacy systems become unavailable after platform migrations or organizational transitions [23]. Data preparation is further inhibited by integration challenges with external systems and heterogeneous endpoints [30]. In practice, PCOs typically operate multiple tools (e.g., submission and review platforms, programme schedulers, registration systems, publication repositories, and survey tools) with non-aligned schemas, inconsistent identifiers, and varying data quality. Such fragmentation introduces recurring challenges in entity resolution (e.g., author name variants, affiliation changes, duplicate profiles) and provenance tracking (i.e., understanding which source and transformation produced a given record). These issues are consistent with broader enterprise integration problems where multiple APIs and custom schemata impede systematic data reuse and automation [24]. Consequently, architectures are required that support controlled and repeatable ingestion, normalization, and alignment across sources, while preserving provenance metadata to ensure downstream robustness and auditability.

Within the IELC, IE provides the transition from raw textual artifacts to structured signals that can inform organizer decisions. TM is a particularly relevant extraction capability because it enables scalable discovery of thematic structure in large corpora [21]. TM methods can be categorized into algebraic, fuzzy, probabilistic, and neural approaches [31]. Algebraic models represent documents and terms as matrices and apply factorization; fuzzy models model graded memberships between documents and topics; probabilistic models such as LDA remain widely used due to their simplicity and interpretability [32]; and neural approaches, increasingly based on Transformer representations, have gained prominence in recent years. In the conference context, topic modeling has been used for session planning and program construction. The POST project clusters conference abstracts using LDA and suggests session topics based on inferred topic distributions [33]. Related studies report similar feasibility for unsupervised grouping and program assistance [34, 35, 36]. More recent systems have begun to incorporate neural semantic representations: for example, NeSyChair employs neural methods such as BERT to identify similarities across paper texts and metadata for schedule construction [37]. Nevertheless, many widely used conference management platforms (e.g., Indico) focus primarily on workflow management and do not integrate NLP-driven automation into organizer-facing decision processes [38]. Methodologically, evaluation of topic modeling for such decision-support settings should combine qualitative inspection (e.g., manual topic plausibility checks and label interpretability) [39] with quantitative indicators capturing quality trade-offs such as topic coherence, diversity, and stability [40]. Overall, the literature indicates increasing application of neural topic modeling in CMS-like settings, yet robust lifecycle integration—including systematic ingestion, persistence of extraction artifacts, and reuse across event editions—remains limited. **Remaining Challenge 1 (RC1):** A lifecycle-aware system is required that systematically ingests and prepares MICE data and applies state-of-the-art topic modeling to produce reusable thematic structures for session planning, together with an evaluation of extraction quality, stability, and organizer-facing utility.

Information & Knowledge Organization via Knowledge Graphs. IKO operationalize how extracted signals are stored, linked, and made queryable. In practical CMS deployments, operational constraints such as licensing, hosting, and maintenance costs can influence architectural choices [30]. From a knowledge organization perspective, KG provide a principled means to represent entities and relations as directed graphs, enabling semantic integration and traceable reasoning across heterogeneous sources. Built on standards such as RDF [41] and OWL [42], knowledge graphs are typically guided by ontologies that define entity types, relationships, and constraints. Their construction can be manual, automated through NLP-based extraction, or hybrid, trading off quality and scalability. In CMS-like settings, knowledge graphs can enrich content organization by supporting semantic tagging and classification, powering semantic search and recommendation, and enabling provenance-aware inspection. Integration between TM and KG has been studied, for example via approaches such as TMKGE [43], which leverage semantic relations and graph embeddings to improve topic coherence and interpretability. However, persistent challenges remain in operationalizing these approaches under real-world MICE conditions: heterogeneous schemas, identity drift (e.g., changing affiliations and author identifiers), relationship integrity across tool migrations, and the need for representations that are usable by organizers rather than only by technical experts. Recent directions point toward deeper integration with Large Language Models (LLMs)-based components for semantic reasoning and natural language interaction; yet such integration increases the importance of grounding, provenance, and auditable evidence paths when used for expert-related decisions. **Remaining Challenge 2 (RC2):** A MICE-specific semantic representation is required that structures extracted information about experts, submissions, and topics in a knowledge graph to preserve relationships and evidence across changing system architectures and to enable traceable access and retrieval.

Information Access & Retrieval for organizer workflows. The final IELC phase concerns IAR supporting access and retrieval of knowledge. Usability and user experience in CMS often suffer due to complex interfaces, particularly for users unfamiliar with advanced analytics functionality [30]. For PCO-facing tasks such as session planning and expert finding, access challenges manifest as limited support for exploratory search (e.g., faceted filtering by topic, track, affiliation, time, or role), insufficient overviews and dashboards to triage large submission sets, and weak interaction patterns for moving from high-level thematic views (topics, clusters, expert lists) to actionable details (individual submissions, author profiles, and their evidence). Moreover, retrieval results that do not expose provenance and evidence paths (i.e., which abstracts, publications, or relationships support a recommendation) can reduce trust and hinder the ability to justify decisions to program chairs or stakeholders. These issues are amplified by performance and scalability constraints, such as slow response times during peak usage or high-latency queries over large indices and graphs [30]. Finally, privacy and security requirements are critical because conference data can include sensitive information (e.g., reviewer identities, conflict of interest relations, unpublished submissions, and internal evaluations), motivating role-based access control, data minimization, and auditable access patterns across retrieval and browsing functions. **Remaining Challenge 3 (RC3):** Organizer-facing systems are required that provide usable, scalable, and provenance-aware interfaces for accessing and retrieving knowledge about submissions and authors, including transparent evidence paths that support PCO decision-making in MICE expert finding.

Summary In summary, the state of the art reveals lifecycle-spanning gaps. Data collection and preparation remain fragmented across heterogeneous tool ecosystems and changing schemas, limiting consistent reuse of conference data. Topic modeling is increasingly applied in CMS-like settings, yet often remains an isolated analytic step rather than a persistent extraction artifact integrated into organizer workflows. Knowledge-graph approaches offer relationship-preserving representations, but are rarely operationalized with MICE-specific ontologies and identity management that remain stable under system change. Finally, access and retrieval mechanisms frequently lack exploratory interaction, scalability, and provenance transparency required for expert-related decisions. These remaining challenges motivate an end-to-end, IELC-aligned approach that integrates topic modeling and knowledge graphs to support PCOs in knowledge management and session planning as enabling functions for MICE expert finding.

3. Modeling and Design

This section presents the modeling and design of the proposed system. In line with the requirements derived in the remaining challenges identified of section 2, the design follows a *User-Centered Systems Design (UCSD)* approach and is explicitly aligned with the phases of the IELC. The modeling focuses on the needs and tasks of MICE PCOs as the primary users and makes the dependencies between data ingestion, IE, knowledge organization, and organizer-facing access explicit.

3.1. Actors and Use Case Model

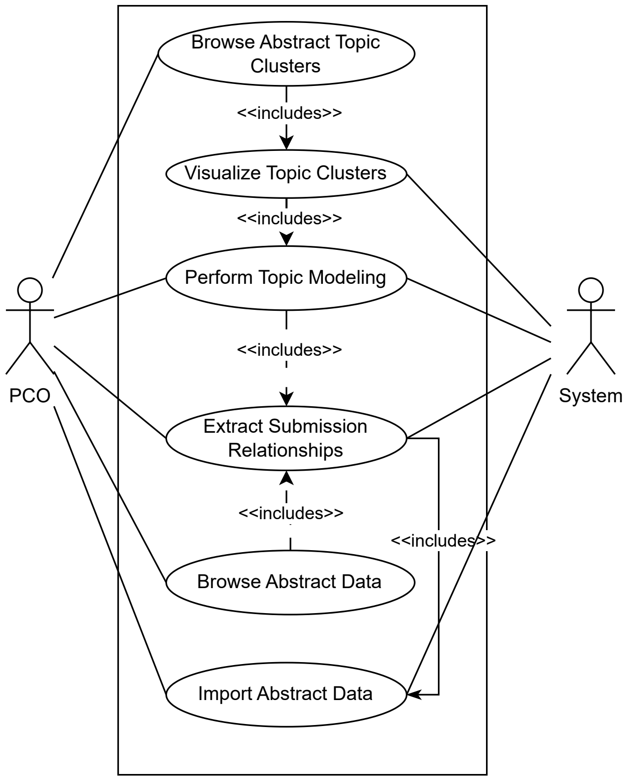

Figure 1: UML Use Case Diagram showing the application use cases.

The primary actor interacting with the system is the PCO, who is responsible for importing data, initiating extraction processes, and exploring the resulting knowledge structures for session planning decisions. The system itself is modeled as a secondary actor that executes automated processing steps such as topic modeling and relationship extraction. Figure 1 presents the UML Use Case diagram of the application. The diagram captures the organizer-facing and system-driven interactions and explicitly encodes their dependencies. The use cases include the import of abstract data, which is required before any extraction or exploration can take place. Once abstracts are ingested, the system can perform topic modeling to derive thematic clusters from the submission corpus and extract submission relationships such as authorship and co-authorship. These extracted structures enable browsing of abstract data as well as visual exploration of topic clusters, which supports session planning and thematic overview tasks. The `<<includes>>` relationships shown in Figure 1 make lifecycle dependencies explicit: browsing and visualization depend on prior topic modeling, which in turn depends on successful data import. By modeling these dependencies directly in the use case diagram, the system design operationalizes RC1 and ensures that organizer-facing functionality is grounded in prepared and extracted data rather than ad-hoc analysis results.

3.2. Information Model

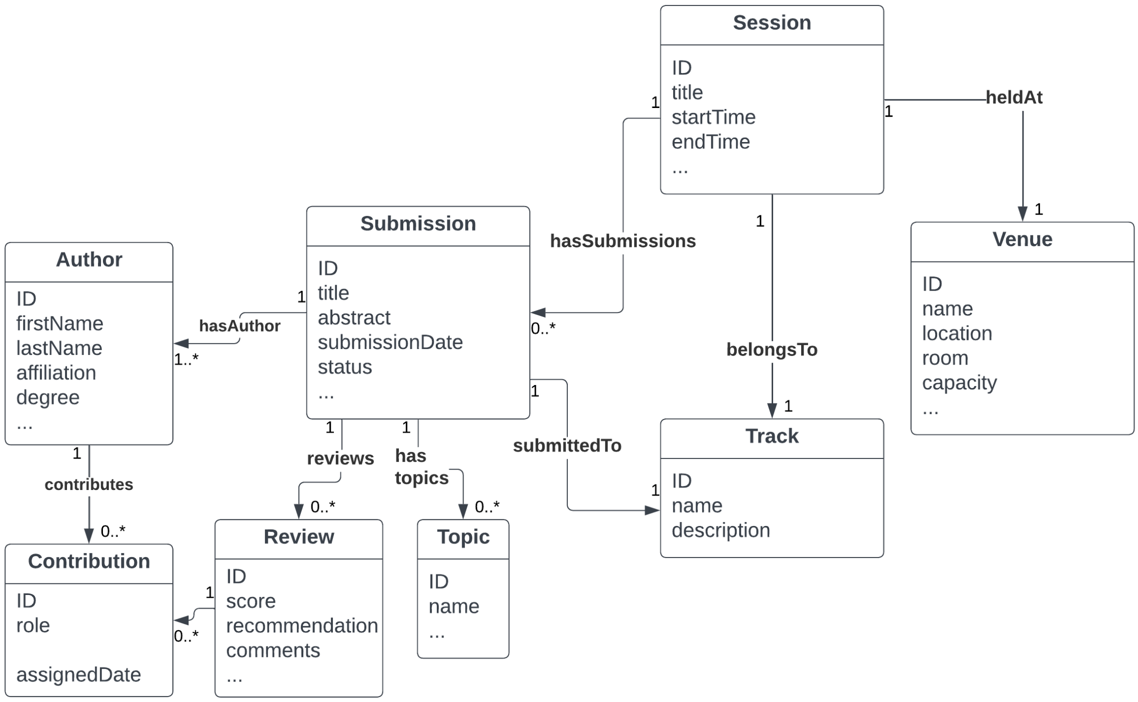

Figure 2: UML Information Model showing the relationship between experts, documents, and topics.

To support expert finding as an evidence-based and relational task, the system relies on an explicit information model that captures the entities and relationships relevant to conference organization. Figure 2 shows the UML information model underlying the application. The model is centered on three core entities: **Author**, **Submission**, and **Topic**. Authors represent experts with associated metadata such as names, affiliations, and degrees. Submissions represent conference abstracts and papers, including their textual content and submission metadata. Topics represent thematic clusters derived from topic modeling. These entities are connected through explicit semantic relationships: authors are linked to submissions via authorship and co-authorship relations, submissions are associated with topics through classification relations, and authors are transitively connected to topics through their authored submissions. Additional entities such as **Review**, **Contribution**, **Session**, **Track**, and **Venue** capture organizational and contextual aspects of conferences, including review outcomes, assigned roles, scheduling, and physical or virtual locations. Together, these entities enable expertise to be represented as structured, relational evidence rather than isolated attributes. This directly addresses RC2 by providing a schema suitable for knowledge-graph–based storage and querying, and by preserving expert-related relationships across system boundaries and conference editions.

3.3. Component Model and IELC Alignment

While the information model defines *what* is represented, the component model specifies *how* the system realizes the IELC phases. Figure 3 presents the UML Component diagram, which maps system components to the corresponding IELC stages. Data collection and preparation are realized through dedicated Abstract and Author Data Collector–Importer components, which ingest external PCO data,

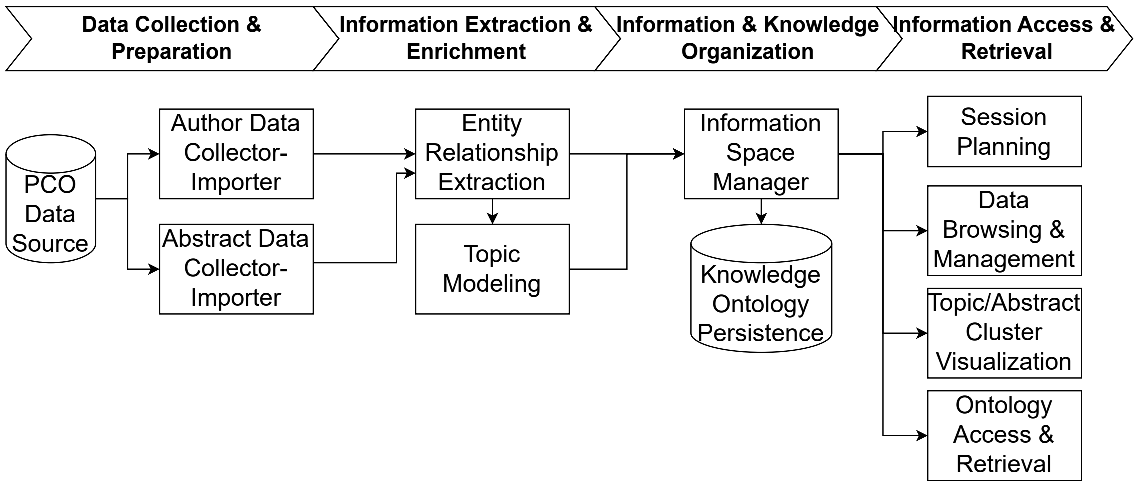

Figure 3: UML Component Diagram showing the system components mapped to their IELC phases.

perform normalization and validation, and prepare it for downstream processing. IE is implemented through two complementary components: Topic Modeling, which derives thematic structures from abstract text, and Entity Relationship Extraction, which identifies authorship and co-authorship relations. These components transform raw input into structured extraction artifacts. Information and knowledge organization are handled by the Knowledge Graph Persistence component, which stores entities and relations in an RDF-based knowledge graph. Finally, information access and retrieval are realized through the Data Browsing, Topic/Abstract Cluster Visualization, Session Planning, and Ontology Access & Retrieval components, which expose the organized knowledge to PCOs in an explorable and interpretable form. By explicitly mapping components to IELC phases, the design ensures lifecycle coverage and addresses RC1 and RC2 in a systematic manner.

3.4. MVC Architecture

To structure the interaction between persistent data, processing logic, and user interfaces, the system adopts the Model–View–Controller (MVC) architectural pattern. Figure 4 illustrates how the previously defined components are assigned to the MVC layers. The *Model* layer comprises the Knowledge Ontology Persistence component and the underlying PCO data sources, which together maintain persistent representations of authors, submissions, topics, and their relationships. The *Controller* layer includes the Abstract and Author Data Collector–Importers, the Entity Relationship Extraction component, the Topic Modeling component, and the Information Space Manager Component. These controllers orchestrate lifecycle operations such as ingestion, extraction, and access and retrieval to the ontology. The *View* layer consists of the Data Browsing and Topic/Abstract Cluster Visualization components, which provide organizer-facing interfaces for exploration and inspection. This separation of concerns improves modularity, maintainability, and extensibility, while ensuring that organizer-facing interactions remain focused on decision-making rather than technical processing. By aligning the MVC layers with the IELC phases and the modeled use cases, the architecture provides a coherent realization of the lifecycle-aware design principles derived from the Introduction and the State of the Art.

4. Implementation

This section describes the concrete realization of the system modeled in the preceding section. The implementation follows the defined use case model, information model, component model, and MVC architecture, and operationalizes the phases of the IELC in an executable prototype.

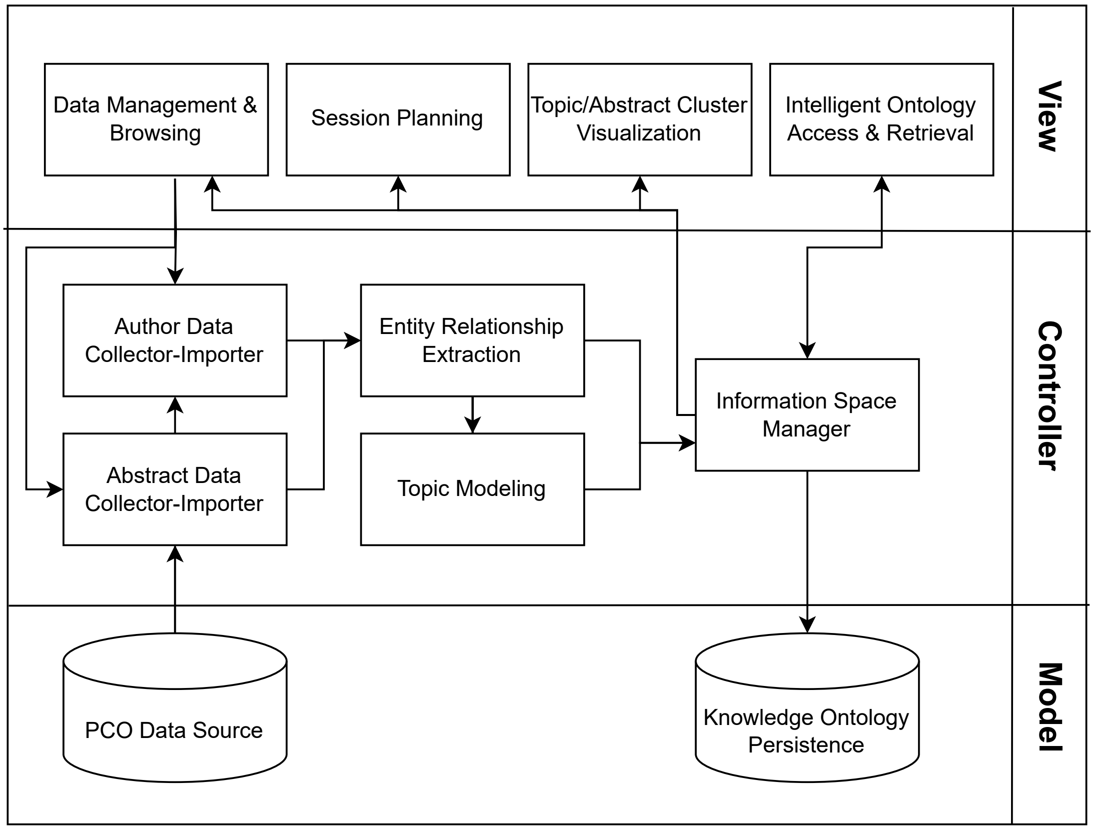

Figure 4: UML MVC diagram showing the system components attribution to the MVC layers.

The components assigned to the Controller layer are implemented as a lightweight RESTful backend using Python and Flask. Flask was selected for its simplicity, modularity, and suitability for prototyping service-oriented architectures. The backend is containerized using Docker to ensure portability, reproducibility, and cloud readiness. The Abstract Data Collector–Importer and Author Data Collector–Importer expose REST endpoints that accept structured input from external PCO data sources, currently provided as JSON exports originating from a relational database. Incoming data is validated, normalized, and transformed into an internal representation aligned with the information model presented in Section 3. During this process, identifiers are assigned and basic consistency checks are performed to prepare the data for downstream processing. The prepared entities are automatically serialized into RDF, establishing a unified and machine-interpretable representation for subsequent extraction and organization.

IE is realized through two complementary Controller-layer components. The Topic Modeling component implements conference-specific thematic extraction by applying the BERTopic library [44] to the imported abstract texts. The implementation includes standard text preprocessing steps, such as tokenization and the application of custom stop-word lists, as well as tuned UMAP and HDBSCAN parameters to balance topic granularity, coherence, and stability. To improve interpretability for PCOs, an integrated large language model is used to generate concise, human-readable summaries for each topic. The resulting topics and topic–submission assignments are persisted as explicit RDF entities and relations, ensuring that thematic structures are treated as reusable extraction artifacts rather than transient analysis outputs. In parallel, the Entity Relationship Extraction component derives authorship and co-authorship relations from the imported submission and author data and persists these relations as first-class semantic links in the knowledge graph. Together, these components transform raw textual

Figure 5: Application screenshot of session topic cluster visualizations.

and relational input into structured, decision-relevant evidence.

Knowledge organization is realized through the Knowledge Graph Persistence component in the Model layer. All entities and relations produced during ingestion and extraction—including authors, submissions, topics, and their semantic connections—are stored in an RDF-based knowledge graph. This representation preserves relationships and supports flexible querying and traversal, consistent with the information model defined in the design. The use of RDF enables schema evolution and interoperability across heterogeneous conference data sources and provides a stable semantic backbone for expert-finding–related knowledge management. The Information Space Manager facilitates communication with the RDF structure.

Organizer-facing access to the organized knowledge is provided by the View layer, which is implemented using Jinja templates. Jinja templates were selected for their declarative, component-based architecture and its ability to support dynamic and responsive user interfaces. The Data Browsing view enables PCOs to inspect imported abstracts, authors, and associated metadata in detail, supporting transparency and traceability. The Session Planning View allows the PCO user to explore the clustered abstracts for the use case of session planning recommendations. The Topic/Abstract Cluster Visualization view presents topic clusters and their associated submissions in a graphical form, enabling exploratory analysis and thematic overview, visible in figure 5. The Intelligent Ontology Access & Retrieval allows the user to retrieve information from the RDF structure using natural language queries that are transformed into SPARQL queries through LLM processing by the Information Space Manager to support naive users in accessing complex data structures and reduce dependency on technical SPARQL and KG structure knowledge. Figure 6 shows an excerpt of the interface that combines natural language

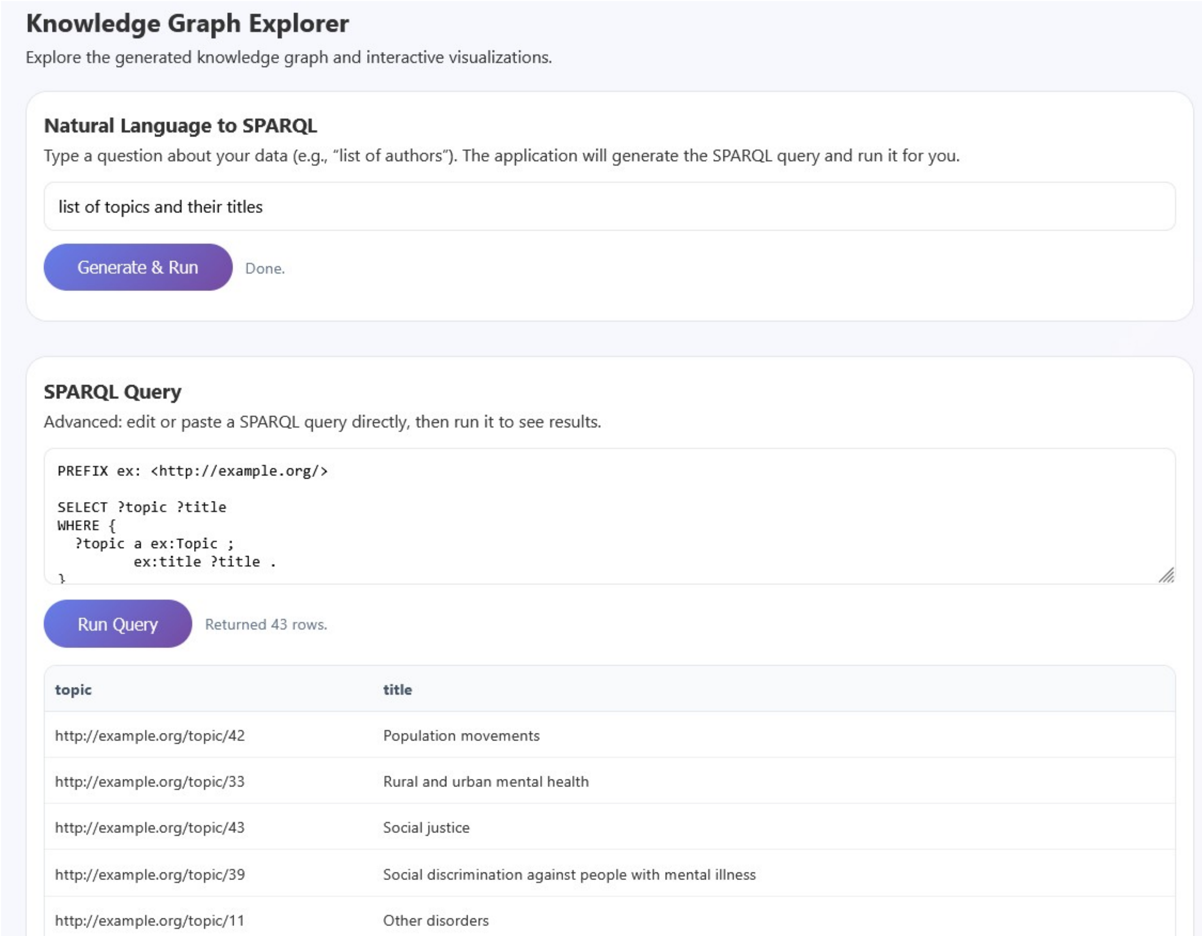

Figure 6: Application screenshot of natural language query retrieval support to interact with the RDF structure.

input, the resulting SPARQL query and log information to provide transparency, and the resulting data. The View layer communicates exclusively with the Flask-based backend via REST APIs, ensuring a strict separation of concerns between presentation and processing logic.

Taken together, the frontend and backend realize the MVC architecture described in the modeling section and provide an integrated, lifecycle-aware implementation that addresses the data ingestion, IE, knowledge organization, and organizer-facing access requirements identified in RC1–RC3.

5. Evaluation

This section evaluates the proposed system with respect to both organizer-facing utility and the technical quality of the extracted thematic structures. In alignment with the lifecycle-aware design and the research questions formulated in the Introduction, the evaluation combines qualitative and quantitative perspectives. The qualitative evaluation examines whether the system supports the intended knowledge management and session planning workflows of MICE PCOs in a transparent and understandable manner (RQ1). The quantitative evaluation assesses whether the implemented topic modeling pipeline produces coherent, diverse, stable, and reusable topic structures that can function as operational information extraction artifacts embedded in the IELC (RQ2). Where possible, the quantitative results are contextualized against published baselines that use the same NeurIPS dataset.

5.1. Qualitative Evaluation

A qualitative evaluation was conducted using a Cognitive Walkthrough (CW), a well-established inspection-based method for assessing interactive systems with respect to user goals, action sequences, and feedback [45, 46]. CW is particularly suitable for early-stage research prototypes and decision-support systems, where transparency, traceability, and workflow fit are more critical than efficiency metrics alone. Three domain experts from the Faculty of Mathematics and Computer Science at the University of Hagen participated in the walkthrough (one professional providing MICE conference management systems, one Ph.D. and one Ph.D. candidate). The experts were provided with simulated user scenarios, that contain a user persona and task description. The experts were provided with an instance of the prototype application and asked to identify potential usability issues for the persona in each scenario and rate them on a Nielsen score [47] (1 - cosmetic, 2 - minor usability problem, 3 - major usability problem, 4 - task failure). The described persona was a Experienced congress organizer, comfortable with event tools and basic analytics, working under tight deadlines and stakeholder pressure. The scenarios are outlined along the view components from the design and modeling section: Scenario 1 - Data management & Browsing (Data collection & preparation IELC phase), Scenario 2 - Session Planning (Information Extraction IELC phase), Scenario 3 - Topic/Abstract Cluster Visualization (Information & Knowledge Organization, Information Retrieval IELC phase), and Scenario 4 - Intelligent Ontology Access & Retrieval (Information & Knowledge Organization, Information Retrieval). We will describe the results of each CW scenario in the following sub-sections.

5.1.1. Data Management & Browsing

Regarding scenario 1 (Data collection & preparation), the task was to upload a provided abstract dataset into the prototype application. One major usability problem has been identified. Users could be confused about which action to take, caused by multiple upload interface options being available (automatic, manual) and these options also persisting throughout the import process. Overall, users could understand the visual feedback provided by the application about the imported abstracts. At the same time, the evaluators noted that the current prototype is limited to a single PCO dataset provided as a JSON export from a relational database. While sufficient for validating the architectural concept, this limitation constrains generality at the prototype stage and reinforces the need for ingestion as an explicit, configurable lifecycle phase (RC1). To improve usability and lifecycle transparency, the experts recommended short in-UI guidance explaining the application's purpose and expected input formats, as well as labeled loading indicators communicating progress during long-running steps such as validation, topic modeling, and persistence.

5.1.2. Session Planning

Scenario 2 (Information Extraction) tasked users to initiate topic modeling to automatically identify thematic clusters and generate a preliminary session structure. The CW revealed three minor and 1 cosmetic usability problems. First, the user could be confused by the wording of the term 'topics', which is used synonymous with 'sessions'. Second, the options of automatic discovery of topics or providing a list of topics in the same field may be confusing. It was noted that PCOs either already possess a list of pre-existing topics, therefore making automatic discovery unnecessary or are interested in discovering emerging topics from a previous conference only. Third, the user is not properly guided to the next step after triggering a topic model execution to the results page. The cosmetic problem was a notification regarding a OpenAI API key that is used to improve the titles of session clusters confuses users and should be removed.

5.1.3. Topic/Abstract Cluster Visualization

Regarding Scenario 3 (Information & Knowledge Organization, Information Retrieval), the user goal is to evaluate how well abstracts were grouped, identify potential overlaps or gaps, and decide whether the structure is trustworthy enough for a potential session plan, using the dedicated analysis page of the application. This use case identified two minor and one cosmetic usability issues. First, the visualizations of top keywords may not be useful to users. Second, the wording issue of 'topics' and 'sessions' persists in this scenario. The first cosmetic issue is that users are not informed about where the word tags on the submission, derived from the top words in each cluster, originate from. Generally the analysis page was noted to be helpful overall.

5.1.4. Intelligent Ontology Access & Retrieval

Regarding Scenario 4 (Information & Knowledge Organization, Information Retrieval), the user task is to answer specific, ad-hoc questions about the program, such as which clusters and how many abstracts exist. They want to retrieve this information without manually cross-checking abstracts or searching multiple databases, and should utilize the intelligent ontology access & retrieval feature. Scenario 4 revealed two critical usability problems. The critical identified problem is that users would have no knowledge about the SPARQL query language. This makes the chosen interface, LLM-supported SPARQL queries based on user-provided text, ineffective. The second usability problem is that users do not have enough knowledge about the underlying RDF graph structure, likely preventing effective querying.

Table 1
Summary of Cognitive Walkthrough Scenarios and Usability Issues

Scenario	Identified Usability Issues	Severity
S1: Data Management & Browsing	Multiple upload options (automatic/manual) cause confusion and persist throughout workflow	3 (Major)
	Lack of guidance on input formats and system purpose; missing progress indicators for long-running tasks	2 (Minor)
S2: Session Planning	Terminology confusion: "topics" vs. "sessions"	2 (Minor)
	Mixed input method (automatic topic discovery vs. manual topic list) unclear	2 (Minor)
	No guidance to next step after topic modeling execution	2 (Minor)
	Confusing OpenAI API key notification	1 (Cosmetic)
S3: Topic/Abstract Cluster Visualization	Limited usefulness of top keyword visualizations	2 (Minor)
	Persistent terminology issue ("topics" vs. "sessions")	2 (Minor)
	Missing explanation of origin of word tags	1 (Cosmetic)
S4: Intelligent Ontology Access & Retrieval	Users lack knowledge of SPARQL, making interface ineffective	4 (Task Failure)
	Insufficient understanding of RDF graph structure for querying	4 (Task Failure)

To summarize, the CW revealed that the prototype is functionally coherent but suffers from several usability and interaction design deficiencies that limit its practical applicability. In data management, users face action ambiguity due to multiple parallel upload options and a lack of guided ingestion, compounded by the current restriction to a single dataset format, which highlights the need to establish ingestion as a configurable lifecycle phase and to provide clearer in-UI guidance and progress feedback. In session planning, inconsistent terminology ("topics" vs. "sessions"), unclear handling of automatic versus manual topic definition, missing workflow guidance after model execution, and unnecessary technical notifications introduce avoidable friction, indicating a need for consistent domain-aligned language, separation of modeling modes, explicit navigation cues, and removal of irrelevant UI elements. The analysis view is generally useful, but limited by low-value keyword visualizations, persistent terminology issues, and insufficient explanation of generated artifacts such as word tags, suggesting improvements in interpretability and analytical relevance. The most critical shortcomings arise in the ontology access component, where reliance on SPARQL and exposure of RDF structures create a fundamental mismatch with user expertise, effectively rendering the feature unusable; this necessitates a redesign toward fully abstracted natural language querying or guided query interfaces. Overall, the key remaining tasks are to standardize terminology, simplify and guide workflows, enhance transparency and feedback, improve the usefulness and explainability of visualizations, generalize data ingestion, remove technical distractions, and, most importantly, replace technically demanding interaction paradigms with user-centered abstractions aligned to the mental models of professional conference organizers. The findings are sumamrized in table 1

5.2. Quantitative Evaluation

The quantitative evaluation assesses whether the BERTopic-based topic modeling pipeline used in the application produces thematic structures that are sufficiently coherent, diverse, and stable to function as reusable information extraction artifacts. We evaluated the pipeline on two datasets: (i) a real-world PCO dataset containing 7,604 conference abstracts and (ii) the publicly available NeurIPS papers dataset with 7,241 documents [48]. The NeurIPS dataset is widely used in topic modeling research and enables comparison with published baselines. Topic granularity was varied using `min_topic_size` $\in \{10, 20, 35, 50, 75, 100\}$. For each configuration, we measured C_v coherence, topic diversity (ratio of unique words across top-n terms), and topic stability across random seeds using average Jaccard overlap of topic word sets, following established evaluation practice [40]. Figure 7 visualizes the results, while Table 2 summarizes representative values.

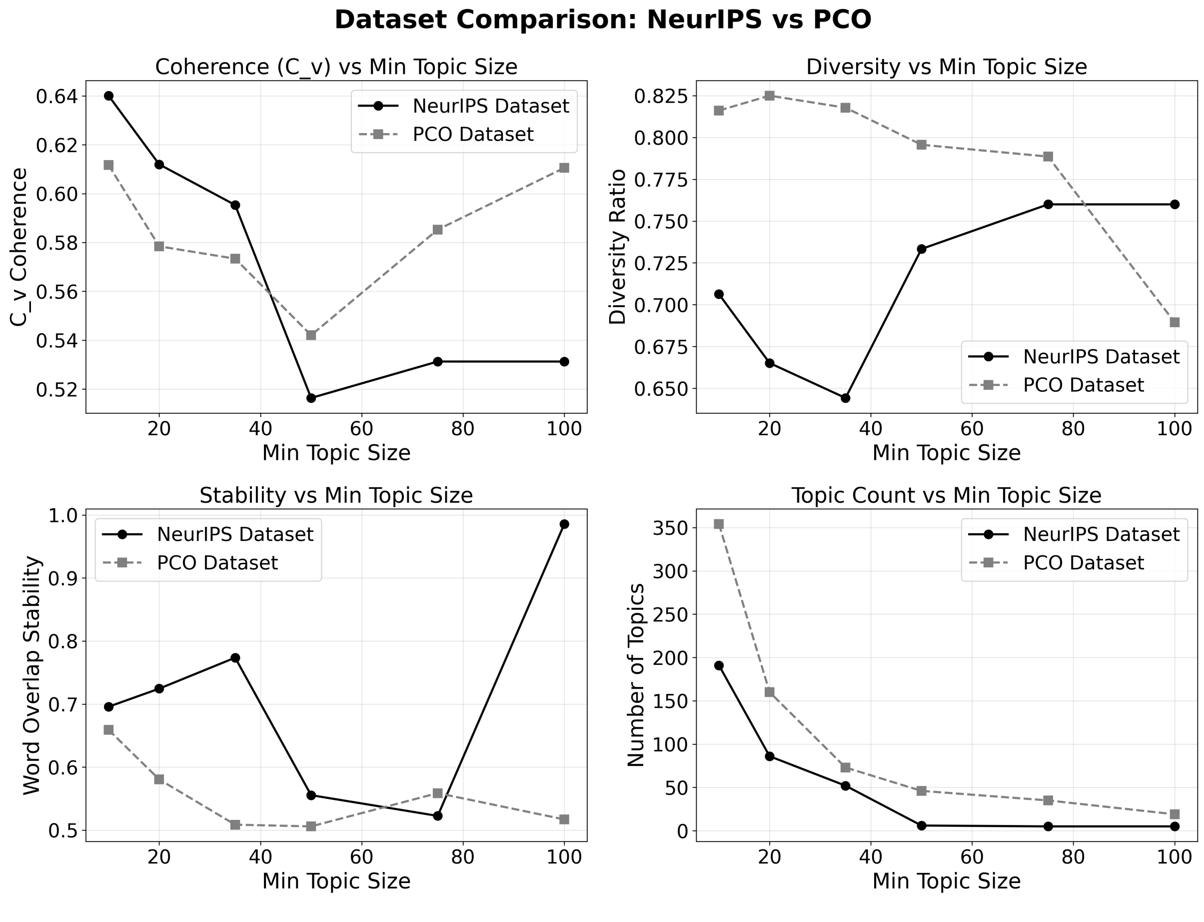

Figure 7: Comparison of BERTopic evaluation runs for the NeurIPS [48] and PCO datasets. The figure reports coherence (C_v), topic diversity, topic stability, and the resulting number of topics across different `min_topic_size` configurations.

On the NeurIPS dataset, C_v coherence ranges from approximately 0.52 to 0.64, peaking at

Table 2
Summary of quantitative topic modeling results across representative configurations. Values are approximated from Figure 7.

Dataset	min_topic_size	#Topics	C_v	Diversity	Stability
NeurIPS	10	~190	~0.64	~0.71	~0.69
NeurIPS	35	~50	~0.59	~0.65	~0.77
NeurIPS	100	8–10	~0.53	~0.76	~0.99
PCO	10	~350	~0.61	~0.82	~0.65
PCO	35	~70	~0.57	~0.82	~0.51
PCO	100	~20	~0.61	~0.69	~0.52

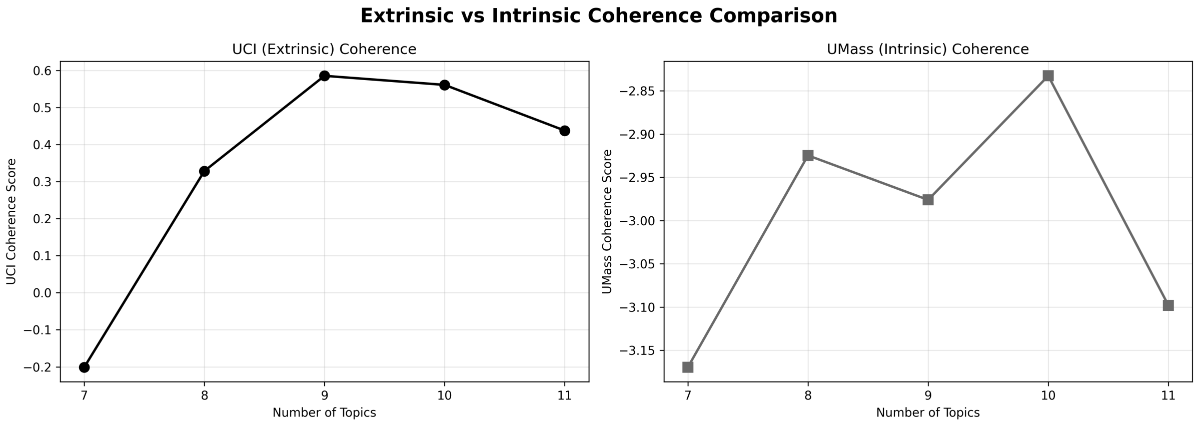

Figure 8: Reference UCI (extrinsic) and UMass (intrinsic) coherence values computed for the NeurIPS dataset.

min_topic_size = 10 and decreasing for coarser configurations. Topic stability increases strongly with granularity, reaching values close to 0.99 at min_topic_size = 100, indicating highly reproducible but very coarse topic structures. The PCO dataset exhibits comparable coherence levels but consistently lower stability (approximately 0.50–0.66), reflecting the greater heterogeneity and noise of real-world conference submissions. Topic diversity is consistently higher on the PCO dataset at fine granularities (up to ~0.83), suggesting that finer topic modeling better captures the breadth of emerging and interdisciplinary themes relevant for session planning.

To contextualize these results, we relate our NeurIPS-based findings to published work using the same dataset. Terko et al. [49] apply LDA to the NeurIPS corpus and select ten topics based on coherence-driven model selection. Using Gensim, we computed UCI and UMass coherence values for comparable topic counts and observed peaks at nine to ten topics (UCI ~0.58, UMass ~ -2.84; cf. Figure 8), closely matching the reference behavior reported by Terko et al. This confirms that our evaluation pipeline produces coherence trends consistent with established baselines on the same corpus.

More recent neural topic modeling work further supports this positioning. Chen et al. [50] report BERTopic as a baseline on the NeurIPS dataset and demonstrate improvements through their FASTopic model. While absolute coherence values differ due to preprocessing and evaluation choices, the relative position of BERTopic in our experiments aligns with its role as a strong but non-optimized neural baseline. Table 3 summarizes these dataset-consistent reference points.

Overall, the quantitative results demonstrate that the proposed system achieves topic quality comparable to established baselines on a widely used benchmark dataset, while explicitly exposing the coherence–diversity–stability trade-offs that matter for organizer-facing workflows. Rather than optimizing for a single coherence optimum, the system enables configurations that balance interpretability, granularity, and stability according to session planning and expert-finding needs.

Table 3
Comparison of topic modeling reference points on the NeurIPS dataset (7,241 papers). Reported values are taken from the cited works where available; differences in preprocessing and evaluation pipelines can affect absolute coherence values. The table serves to contextualize orders of magnitude and modeling regimes on a shared corpus.

Work	Model	#Topics	Coherence	Main focus
This work	BERTopic	~190	C_v ~0.64	Fine granularity, session planning
This work	BERTopic	8–10	C_v ~0.53	High stability, coarse overview
Terko et al. [49]	LDA	10	UCI ~0.58	Coherence-optimized topic count
Chen et al. [50]	FASTopic	25–50	C_v > BERTopic baseline	Neural TM improvement
Chen et al. [50]	BERTopic (baseline)	25–50	C_v ~0.38	Baseline reference

5.3. Summary of Evaluation Results

The evaluation combines qualitative and quantitative analyses to assess both user-facing utility and technical performance of the system. The qualitative Cognitive Walkthrough with domain experts indicates that the prototype is conceptually coherent and supports key workflows for professional conference organizers, but suffers from notable usability issues, including ambiguous actions in data ingestion, inconsistent terminology (e.g., "topics" vs. "sessions"), insufficient workflow guidance, and limited transparency of system feedback. While the analysis and visualization components are generally useful, their interpretability and relevance require improvement. The most critical limitation lies in the ontology access feature, where reliance on SPARQL and RDF concepts creates a clear mismatch with user expertise, rendering it impractical without significant abstraction. The quantitative evaluation shows that the BERTopic-based pipeline produces topic structures with coherence, diversity, and stability comparable to established baselines on the NeurIPS dataset, while also performing robustly on a real-world PCO dataset. Results highlight expected trade-offs: finer topic granularity improves coherence and diversity but reduces stability, whereas coarser configurations yield highly stable but less detailed structures. Overall, the system demonstrates technically sound topic modeling capabilities and flexible configuration for different use cases, but requires substantial improvements in usability, interaction design, and abstraction to achieve practical applicability.

These findings reinforce the central design claim of this work: topic modeling should be treated as a configurable, inspectable information extraction artifact embedded in a lifecycle-aware system, rather than as a one-off analytic step. Overall, the evaluation provides empirical support for RQ1 and RQ2 by demonstrating that Transformer-based topic modeling and knowledge-graph–based organization can jointly support evidence-grounded knowledge management and session planning in MICE settings.

6. Conclusion

This application explored a system to apply TM on systematically sourced PCO abstract data in order to support MICE PCOs in session planning for conferences. The Data Collection challenge to import data from PCO source systems for efficient retrieval and persist it including relationship structures was addressed through realizing import from a PCO data export, however the implementation is basic and does not allow for automated imports from multiple sources and is not fully integrated. Future work should explore a wrapper-mediator architecture to allow for seamless imports from multiple systems.

Regarding IE, TM was applied to generate session topic clusters to support PCOs in session planning. The modeled clusters from the dataset were evaluated using manual inspection and coherence, diversity, and stability metrics, to indicate cluster and topic quality. However, this is a challenging task due to the limitations of defining and measuring optimal topic clusters. Future work should evaluate the TM approaches on datasets and include user feedback or compare against pre-existing classifications. Additionally, further TM techniques, e.g. taxonomy-based classification should be considered for implementation and evaluation.

Regarding Knowledge Organization, a KGs was applied to represent abstracts, authors, topics, and their relationships. Future work should include and evaluate further entity types and relationships into this structure to expand application use cases.

Regarding Information Access & Retrieval, the visualizations to support PCOs in session planning were not exhaustively explored, as only a prototype visualization was created using the underlying KG relationships. Future work should include further interface and visualization and retrieval implementations, conduct user studies to gather evidence on usability of such visualizations regarding the use case of session planning and further PCO tasks. The intelligent KG access & retrieval should be evaluated thoroughly and we observe a lack in comparable evaluation frameworks for natural language query support.

Declaration on Generative AI

During the preparation of this work, the author(s) used OpenAI GPT-5.3 in order to: Drafting content, Paraphrase and reword, Grammar and spelling check. Further, the author(s) used Anthropic Claude Sonnet 4.2 for figure 2 in order to: Generate images. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] B. Fowler, P. W. Cardon, B. Marshall, K. Elder, What do conference attendees want from academic presentations? a study of an information systems professional organization., Issues in Information Systems 22 (2021).

[2] D. Severt, Y. Wang, P.-J. Chen, D. Breiter, Examining the motivation, perceived performance, and behavioral intentions of convention attendees: Evidence from a regional conference, Tourism management 28 (2007) 399–408.

[3] C. T. A. (CTA), Ces - consumer electronics show, 2024. URL: https://www.ces.tech/, accessed: 2024-12-07.

[4] S. M. Tinnish, S. M. Mangal, Sustainable event marketing in the mice industry: A theoretical framework, in: Journal of Convention & Event Tourism, volume 13, Taylor & Francis, 2012, pp. 227–249.

[5] Custom Market Insights, Global mice market report - 2023 to 2032, 2024. URL: https://www.custommarketinsights.com/report/mice-market/, accessed: 2024-09-03.

[6] Die Zeit, Wissenschaftler auf reisen: Wie wichtig sind flüge für die forschung?, Zeit Online (2019). URL: https://www.zeit.de/2019/35/wissenschaftler-berufsreisen-fluege-konferenz-vortrag-feldforschung, accessed: 2024-09-03.

[7] V. Pavluković, M. Vujičić, J. Kennell, M. Cimbaljević, Sustainability in the meetings industry (2020).

[8] GCB German Convention Bureau e.V., Tagung und kongress der zukunft: Die studie, 2013. URL: http://www.gcb.de/de/tagung-und-kongress-der-zukunft/tagung-und-kongress-der-zukunft---die-studie, accessed: YYYY-MM-DD.

[9] Y. Yuan, J. Enerio, W. Zhang, A guide to sustainable events: Singapore mice carbon calculator (2024).

[10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017).

[11] O. Team, Hello, gpt-4o, 2024. URL: https://openai.com/index/hello-gpt-4o/.

[12] M. Stankovic, C. Wagner, J. Jovanovic, P. Laublet, Looking for experts? what can linked data do for you?, in: LDOW, 2010.

[13] O. Husain, N. Salim, R. A. Alias, S. Abdelsalam, A. Hassan, Expert finding systems: A systematic review, Applied Sciences 9 (2019) 4250.

[14] S. Mundry, Designing successful professional meetings and conferences in education: Planning, implementation, and evaluation, Corwin Press, 2000.

[15] P. Tamla Mbobda, Towards a reference model for an information extraction lifecycle supporting data collection, data preparation, information extraction, information organization, knowledge organization, information access and retrieval, 2025. URL: https://ub-deposit.fernuni-hagen.de/receive/mir_mods_00002257.
[16] R. Grishman, B. Sundheim, Message understanding conference-6: A brief history, in: 16th Conference on Computational Linguistics, 1996, pp. 466–471.
[17] A. Y. Chua, D. H. Goh, R. P. Ang, Web 2.0 applications in government web sites: Prevalence, use and correlations with perceived web site quality, Online information review 36 (2012) 175–195.
[18] D. Hoe-Lian Goh, A. Chua, C. Sian Lee, K. Razikin, Resource discovery through social tagging: a classification and content analytic approach, Online Information Review 33 (2009) 568–583.
[19] G. Chowdhury, Introduction to modern information retrieval, Facet Publishing, 2010.
[20] W. H. W. Ishak, F. M. Yamin, M. F. M. Mohsin, M. F. Mansor, A comparative review of conference management system, Journal of Technology and Operations Management 18 (2023) 87–93.
[21] I. Vayansky, S. A. Kumar, A review of topic modeling methods, Information Systems 94 (2020) 101582.
[22] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, Knowledge graphs, ACM Computing Surveys (2021). URL: https://dl.acm.org/doi/abs/10.1145/3447772.
[23] A. Somani, D. Choy, J. Kleewein, Bringing together content and data management systems: Challenges and opportunities, IBM Systems Journal 41 (2002) 686–696.
[24] S. Andreo, J. Bosch, Api management challenges in ecosystems, in: International Conference on Software Business, Springer International Publishing, Cham, 2019.
[25] P. Tamla, M. Bornschlegl, M. Hemmje, KGen4MICE: Knowledge Graph–Based Decision Support for the MICE Domain, Technical Report, FernUniversität in Hagen, 2024.
[26] P. Tamla, Supporting Access to Textual Resources Using Named Entity Recognition and Document Classification, Ph.D. thesis, Hagen, 2022. URL: https://ub-deposit.fernuni-hagen.de/receive/mir_mods_00001782.
[27] B. Hartmann, P. Tamla, M. Hemmje, Supporting deep learning-based named entity recognition using cloud resource management, in: H. Degen, S. Ntoa, A. Moallem (Eds.), HCI International 2023 – Late Breaking Papers, Springer Nature Switzerland, Cham, 2023, pp. 84–100.
[28] B. Hartmann, P. Tamla, F. Freund, M. Hemmje, Fine-tune it like i'm five: Supporting medical domain experts in training ner models using cloud, llm, and auto fine-tuning, in: 2023 31st Irish Conference on Artificial Intelligence and Cognitive Science (AICS), 2023, pp. 1–8. doi:10.1109/AICS60730.2023.10470654.
[29] J. F. Nunamaker Jr, M. Chen, T. D. Purdin, Systems development in information systems research, Journal of management information systems 7 (1990) 89–106.
[30] W. Ishak, F. Yamin, M. Mohsin, M. Mansor, A. Husin, Journal of digital system development, Journal of Digital System Development: Vol 2 (2024) 126–135.
[31] A. Abdelrazek, Y. Eid, E. Gawish, W. Medhat, A. Hassan, Topic modeling algorithms and applications: A survey, Information Systems 112 (2023) 102131. doi:10.1016/j.is.2022.102131.
[32] D. M. Blei, A. Y. Ng, M. I. Jordan, Latent dirichlet allocation, Journal of Machine Learning Research 3 (2003) 993–1022.
[33] N. Moore, S. Mayer, K. Molloy, P. W. Woźniak, W. Lovo, M. Stewart, Post: A machine learning based paper organization and scheduling tool, Sub Human-Centered Computing 1 (2020).
[34] M. Burke, D. Sabatta, Topic models for conference session assignment: Organising pr as a 2014(5), in: 2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech), 2015, pp. 94–99.
[35] H. C. Lau, A. Gunawan, W. Wang, et al., Press: Personalized event scheduling recommender system (2016).
[36] K. Sweeney, Unsupervised machine learning for conference scheduling: a natural language processing approach based on latent dirichlet allocation 1 (2020).
[37] T. Škvorc, N. Lavrač, M. Robnik-Šikonja, Nesychair: Automatic conference scheduling combining neuro-symbolic representations and constrained clustering, IEEE Access 10 (2022) 10880–10897.

[38] J. B. González López, J. P. Ferreira, T. Baron, Indico central – events organisation, ergonomics and collaboration tools integration, Journal of Physics: Conference Series 219 (2010) 082002.
[39] J. D. Romero, M. A. Feijoo-Garcia, G. Nanda, B. Newell, A. J. Magana, Evaluating the performance of topic modeling techniques with human validation to support qualitative analysis, Big Data and Cognitive Computing 8 (2024) 132.
[40] C. Meaney, T. A. Stukel, P. C. Austin, R. Moineddin, M. Greiver, M. Escobar, Quality indices for topic model selection and evaluation: a literature review and case study, BMC Medical Informatics and Decision Making 23 (2023) 132.
[41] J. Pan, Resource description framework, in: Handbook on ontologies, Springer, 2009. URL: https://link.springer.com/chapter/10.1007/978-3-540-92673-3_3.
[42] G. Antoniou, F. Harmelen, Web ontology language: Owl, in: Handbook on ontologies, Springer, 2009. URL: https://link.springer.com/chapter/10.1007/978-3-540-92673-3_4.
[43] D. Li, S. Zamani, J. Zhang, P. Li, Integration of knowledge graph embedding into topic modeling with hierarchical Dirichlet process, in: J. Burstein, C. Doran, T. Solorio (Eds.), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 940–950. URL: https://aclanthology.org/N19-1099/. doi:10.18653/v1/N19-1099.
[44] M. Grootendorst, Bertopic: Neural topic modeling with a class-based tf-idf procedure, arXiv preprint arXiv:2203.05794 (2022).
[45] C. Wharton, J. Rieman, C. Lewis, P. Polson, The cognitive walkthrough method: A practitioner's guide, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), ACM, 1994, pp. 105–110. doi:10.1145/191666.191704, indexed in Google Scholar and ACM Digital Library.
[46] C. Lewis, J. Rieman, Task-Centered User Interface Design: A Practical Introduction, University of Colorado, Boulder, 1997. Widely cited technical report/book; indexed in Google Scholar.
[47] J. Nielsen, Usability engineering, Morgan Kaufmann, 1994.
[48] B. Hamner, Nips papers: Titles, authors, abstracts, and extracted text for all nips papers (1987-2017), 2017. URL: https://www.kaggle.com/datasets/benhamner/nips-papers, dataset on Kaggle.
[49] A. Terko, E. Žunić, D. Đonko, Neurips conference papers classification based on topic modeling, in: 2019 XXVII International Conference on Information, Communication and Automation Technologies (ICAT), IEEE, 2019, pp. 1–5.
[50] J. Chen, W. Li, Y. Shi, J. Wang, Fastopic: A fast, adaptive, stable, and transferable neural topic model, arXiv preprint arXiv:2210.14632 (2022). ArXiv paper; indexed in Google Scholar; reports BERTopic baseline on NeurIPS.