

The Expertise Ontology: Modeling Expertise in the Context of Emergency Management^{*}

Shirly Stephen^{1,*}, Mark Schildhauer¹, Ling Cai², Yuanyuan Tian³, Kitty Currier¹, Cogan Shimizu⁴, Krzysztof Janowicz^{1,5}, Pascal Hitzler⁶, Anna Lopez-Carr⁷, Andrew Schroeder⁷, Zilong Liu^{1,5}, Rui Zhu⁸, Dean Rehberger⁹, Colby K. Fisher¹⁰ and Gengchen Mai¹¹

¹University of California, Santa Barbara, CA, USA

²IBM Research, San Jose, CA, USA

³Arizona State University, Tempe, AZ, USA

⁴Wright State University, Dayton, OH, USA

⁵University of Vienna, Vienna, Austria

⁶Kansas State University, Manhattan, KS, USA

⁷Direct Relief, Santa Barbara, CA, USA

⁸University of Bristol, Bristol, UK

⁹Michigan State University, MI, USA

¹⁰Hydronos Labs, Princeton, NJ, USA

¹¹University of Georgia, Athens, GA, USA

Abstract

It is crucial for emergency management organizations to have rapid access to relevant experts who can advise and assist following a disaster. To improve expert-mining and recommendation capabilities, creating a knowledge graph that links experts to their corresponding topics of expertise and other sources of relevant information is a natural choice to capture an integrated network of people and a rich taxonomy of expertise. In this paper, we present an ontology for modeling experts, their expertise topics and relations between them, and their spatiotemporal scoping. We go on to discuss the primary conceptual components and how they can be instantiated, then present overarching examples related to emergency management operations. The ontology synthesizes three different ways to characterize an expert, based on a) identifiable academic expertise; b) voluntary engagements, work-related responsibilities or experience; and c) organization specializations or affiliations.

Keywords

ontologies, expertise modeling, emergency management, semantic web, knowledge graphs

1. Introduction

Identifying people with knowledge and expertise in specific topic areas is critical for many purposes—for example, to find peer reviewers, to identify strategic hires, or to find consultants. While not all manner of knowledge qualifies as expertise, it implies a sense of competency

Ontology Showcase and Demonstrations Track, 9th Joint Ontology Workshops (JOWO 2023), co-located with FOIS 2023, 19-20 July, 2023, Sherbrooke, Québec, Canada.

^{*}This research has been supported by the National Science Foundation under Grant No. 2033521: “KnowWhereGraph: Enriching and Linking Cross-Domain Knowledge Graphs using Spatially-Explicit AI Technologies”.

^{*}Corresponding author.

✉ shirlystephen@ucsb.edu (S. Stephen)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

that is obtained through a combination of theoretical knowledge, practical experience, and training. Some philosophers, such as Stichter [1], suggest expertise to be a social or reputational phenomenon, while others, like Goldwin [2], propose expertise as a functional phenomenon that denotes an agent’s “capacity” to help others address issues in a distinctive manner that the latter, themselves, could not solve. While the work within this paper is not critical of any of these stances, we refrain from strictly adopting them and rather recommend a representation model that synthesizes different views of expertise and what constitutes an *expert*.

The knowledge modeling in this paper is driven by needs in the emergency management domain to quickly identify people with relevant knowledge and field experience to advise on medical and health concerns in the immediate aftermath of a disaster. For humanitarian aid organizations, it is beneficial to collaborate with individuals and organizations who have a detailed understanding of the disaster, geographic and socio-demographic characteristics of the affected areas, and who may have provided assistance during similar situations in the past.

Open Knowledge Graphs (KGs) such as the KnowWhereGraph [3] can contain a wealth of machine-readable information on events, infrastructure, environmental and health observations, places, regions, and other subjects. Their content—including concepts and their instances—can be aligned to specific topics in a structured hierarchy. By doing so, one can mine for people who have expertise in some topic but also retrieve all the connected information such as events and places in the graph that is relevant to the topic or vice-versa. Not only are these connections important, but they are absent in most existing expert-finding systems. This is because, in contrast to conventional systems, graphs provide a means to easily enrich connections of topics with their real-world instances. For example, the topic ‘Hurricane Katrina’ can be connected with instances corresponding to its trajectory, its socio-demographic impacts, damage caused to infrastructure, and the disaster declaration made by the state seeking federal financial assistance.

People having different levels of expertise in a topic can be identified manually by relying on expertise directories, or using automated approaches such as Information Extraction methods on publications, online documents, web pages, and CVs. Regardless of the method used to identify experts, all the varied information can be made searchable in a KG through a formal model that integrates the information to allow consistent and straightforward querying. The focus of this paper is on describing this model, namely the Expertise Ontology (ExO) that can be used to represent expertise-related information in a KG.

Overview and Significance: The ExO incorporates three categories of “expert” based on their 1) explicitly identifiable expertise, 2) experience or activities they may have engaged in, and 3) job role(s) or organizational affiliation(s). This ontology encourages the representation of topics at different levels of granularity, creates mappings between experts (people and groups) and these topics, and then connects topics with relevant content in a KG. It is also structured to facilitate access to not only research- and theory-based expertise, but also experience-based expertise, which often is not reported in scientific journal articles. The ExO also recommends the use of reification to characterize qualitative and quantitative levels of expertise.

The rest of the paper is organized as follows. We briefly expand on the use-case scenario, present a list of competency questions, and outline the conceptual aspects of the ExO in Section 2. In Section 3 we discuss the modeling of the ExO. We use a few example datasets to demonstrate the use of ExO in Section 4. In Section 5 we provide a brief review of existing related ontologies, and finally, in Section 6 we conclude.

2. Approach

The ExO is developed following the ontology development approach presented by Noy & McGuinness in [4]. The steps undertaken in the development of this ontology are as follows: 1) describe the domain and scope of the ontology through a humanitarian aid use-case within the KnowWhereGraph (KWG) [3] framework; 2) identify important terms and conceptual aspects of the ontology; 3) identify existing ontologies that will be reused; 4) formally define the classes and properties; and finally 5) populate the KG with sample data, and 6) evaluate the ExO using competency questions.

The foundation upon which we are building the ExO is KWG¹, a geographic KG structured by a set of ontologies [5, 6, 7, 8, 9] that integrates over 32 different datasets and contains over 16 billion triples as of April 2023. These triples involve observations on disasters, environmental factors, transportation infrastructure, health-related statistics for populations across the US, and so forth. Envisioned as a tool for providing area briefings within seconds—for humanitarian relief and other domains—KWG is intended to assist decision-makers in industry, government, and the nonprofit sector [10, 11]. In the ExO-specific use case presented below, we focus on the capacity of KWG to help disaster management personnel discover experts with knowledge, experience, and on-the-ground skills relevant to a population’s health and medical needs in anticipation of, during, and following a disaster.

2.1. Use Case: Humanitarian Aid Scenario

In the aftermath of a hurricane, saving lives depends on a response that is timely and provides resources appropriate to the situation. This is one goal of Direct Relief², a private humanitarian nonprofit organization based in Santa Barbara, California, whose core mission is to improve the health and lives of people worldwide affected by poverty or emergencies without regard to politics, religion, or ability to pay. With a vast scope in the geographic context and type of disaster that triggers an aid response, Direct Relief responders must quickly identify experts who can advise and assist in a particular situation. This might include a researcher who models a pathogen’s spread; staff in a health center that specializes in treating a particular disease; or a local government official with experience in previous disasters. The competency questions included below pertain to this use case.

- (CQ1) Who has medical expertise relevant to health needs that may follow a hurricane?
- (CQ2) Who is familiar with the unique needs—healthcare-related and otherwise—of a population demographic comprising middle-aged adults with diabetes?
- (CQ3) Who can advise on how to respect the historical and cultural sensitivities of Asian population in Fresno, California, during our response?
- (CQ4) Who is likely to be directing local emergency response operations?
- (CQ5) Where are the nearest medical facilities with staff and equipment, including cold-chain storage, capable of handling the outbreak of a water-borne endemic?
- (CQ6) Who has expertise administering aid in regions having arid weather conditions?
- (CQ7) How recent is their experience or work on pandemic response?
- (CQ8) What geographic regions are associated with mosquito-borne diseases?

¹<https://knowwheragraph.org/>

²<https://www.directrelief.org/>

2.2. Identifying Core Terms and their Semantics for Ontology Design

The key notions of the ontology as laid out in both the introduction and the use case are Agent (“who”) and Topic (“what expertise”) concepts, expertise relations that connect the two, relations to structure and organize topics, and relations that connect them with other content in a KG. To build a comprehensive expert KG, information pertaining to expert knowledge and topic vocabularies is typically collected from varied sources, some of which are discussed later in Section 4. Quantitative evaluations to assign types and levels of expertise can be performed using different algorithms [12, 13]. Thus, it is obvious that the provenance of the data and metrics be preserved. In the paragraphs below, we identify the key conceptual aspects of the use case to elucidate the ExO’s modeling choices outlined in Section 3.

Who is an *Expert*? State-of-the-art AI-powered expert recommendation systems, such as Expertise Finder [14] and Elsevier’s Expert Lookup [15], generate recommendations by analyzing published journal articles, which tend to be authored by academic researchers. This would include people who have specialized knowledge in particular health and medical topics, such as the risk of a waterborne disease outbreak following a flood in an urban US setting. However, these systems would be less likely to identify individuals who lack a publication record, though they may have valuable training and experience in responding to previous disasters. Individuals in the medical, healthcare, first responder, and humanitarian relief fields, who may be affiliates of government, non-profit, and commercial organizations—and include professionals, volunteers, and people affected by the crisis—might equally qualify as experts. Online directories may provide some information, but they are frequently incomplete, out of date, or compiled at such a coarse resolution as to be of limited use in an actual emergency. For example, an online directory of county health departments is maintained by the National Association of County and City Health Officials³, but it only provides a contact information record for each department. Discovering whether a particular department has an epidemiologist on staff, for example, would require a visit to the individual department’s website, which may or may not expose the sought-after information. Since the granularity of information available through an online source, such as a web-page or directory entry, may prevent us from resolving expertise within an organization to a single person, we expand the scope of who may be considered an expert beyond an individual person to include an organization. To be useful in the context of emergency management operations, a database, and its associated schema must synthesize these different views of what constitutes an expert.

What characterizes an *Expertise Topic*? The relevance of an expert recommended by an information system depends largely on the system’s ability to characterize expertise topics. A hydrologist, while possibly an expert in flooding, may be of little help in delivering medical aid following a flood. Systems with an expert network that extends beyond academia, such as those powered by LexisNexis [16] or SEAK Experts [17], have expertise areas that are broadly scoped but follow a flat taxonomy. Even literature publishing platforms like Semantic Scholar⁴ or Academia.edu have topic vocabularies that are too general for many applications. For example, one can expect to find “Disaster” as a topic and maybe even a more granular “Hurricane” topic, but the chances of finding “Hurricane Katrina” are slim. This limits the usefulness of

³<https://www.naccho.org/membership/lhd-directory>

⁴See <https://www.semanticscholar.org/>

these platforms for disaster operation needs, where identifying people with experience in a specific geographic region and time is beneficial. A comprehensive treatment of topics from every domain is clearly out of the scope of this paper. Since many standard or authoritative vocabularies have been developed in a principled way, we choose to use these as a reference for constructing topic hierarchies using appropriate classes and relations from our ontology. Specific to the humanitarian aid use-case scenario described above is the *disease* and *disaster* themes, and we therefore specifically use the Disease Ontology (DO) [18] and the United Nations Office for Disaster Risk Reduction (UNDRR) hazard information profile (HIP) [19, 20] to demonstrate the translation process.

Spatiotemporal context Spatial location is closely tied to any disaster phenomenon, and spatiotemporal context is crucial during emergency operations in the case of both individuals with on-the-ground experience and academic researchers, in terms of where agents are located, and where their area of expertise is focused. Spatial information can also link other critical information such as the socio-demographic and public health profile of a community. Understanding the health issues faced by a population can help in formulating a plan, including identifying what specialties are needed, and narrowing down which agents are relevant for requesting assistance based on the situation expected. An aid provider may have learned through experience, for example, that high rates of diabetes in their community make insulin a valuable part of their response kit. Likewise, temporal context is important to identify the duration of an agent's connection with a topic, both to determine their level or strength of expertise and to determine if they are still relevant to the need at hand.

2.3. Identifying Existing Ontologies for Reuse

We choose to use existing ontology standards where possible for representing fundamental concepts and to provide greater interoperability and reusability. Some ontologies are obvious choices, such as using the core of PROV-O [21] for provenance, GeoSPARQL [5] for representing spatial features and geometries, and OWL-Time [6] for temporal representation. Further, we show alignments with specific agent classes in FOAF [22], and the Organization Ontology [23].

3. Description of the ExO

Here we provide a broad overview of the ExO (with the namespace *eo*), discuss key modeling choices, and provide examples where needed. The key notions of our model are the Expert and Topic classes and their relationship. Figure 1 shows the corresponding schema diagram. Note the use of spatial, temporal, and provenance information, which are indicated using a different color and border to indicate that they are described in detail in external patterns or ontology standards. The OWL file for the ExO can be found online [here](#). Detailed documentation of axioms in Description Logic syntax can be found [here](#).

3.1. Topics and Modeling Topic Hierarchy

The general Topic class refers to themes or subject areas. Topics can range in scope from broad areas of knowledge like Science and Nanoscience to fine-grained areas of knowledge like the Impact of Hurricane Ida in Louisiana. Many domain ontologies organize concepts referring to specific domain topics such as environmental science (e.g., the Environment Ontology [24]) or




Figure 1: A simplified schema diagram of the ExO. Orange boxes represent concepts central to ontology. The blue box with a dashed border represents an interface to external classes/patterns that are left unmodeled in the ExO. Yellow boxes are actual concepts from external ontologies that may have more semantics not covered here. Black-filled arrows are object or data properties, and open arrows represent subclass relationships.




Figure 2: Schema diagram for the Topic class and inter-relationships.

biomedical science (e.g., the DO [18]) using a class–subclass hierarchy. Applications outside the knowledge representation realm might reuse terms from these ontologies as topics, rather than with their more formal axiomatization as, e.g., OWL classes. Ontologically speaking, then, an instance of the Topic class is distinct from the class that it might refer to in a domain ontology. For example, Disease is an instance of the Topic class in the ExO, as opposed to a Disease class in the DO, but a mapping can be made between the two. Classes differ from specific (instances of) Topic in that they represent a bag of instances, they may have properties describing their various features and attributes, and they can have subclasses that represent concepts more specific than the superclass. Any specific topic represents an area of discourse, useful for annotating distributed content (instances and even classes) in a KG for search and summarization purposes. Topics can also be related to narrower ‘sub-topics’ or broader ‘super-topics’ without the logical constraints implied by more formal class–subclass models. Moreover, in the ExO we model specific topics as instances of the very general Topic class, and therefore one can assert taxonomic relations between topics, or relations between topics and other instances in the graph (e.g., people, events, observations) without punning.

Developing a topic hierarchy from scratch can be tedious. Tools for publishing and browsing scientific literature, such as Semantic Scholar and PubMed, adopt a bottom-up approach to find and organize topics based on clustering algorithms run on text segments. We build a topic hierarchy by reusing existing domain vocabularies, taxonomies, and ontologies that are driven by a diverse community of expertise (e.g., the DO, the UNDRR HIP classification), and some even constructed by government institutions and considered more authoritative (e.g., Medical Subject Headings (MeSH) [25]). For the use case discussed in Section 2.1, the DO and the UNDRR HIP classification are used to construct the initial hierarchy of topics. At the next stage, compound topics that span multiple domains (e.g., *DiabetesDisasterResponse*) are arranged manually within the constructed hierarchy.




Figure 3: Example: Construction of topics (pink boxes) for KWG using a subset of the DO concepts (green boxes). The entire topic hierarchy in turtle format can be found [here](#).

The schema diagram in Figure 2 shows how topics are organized and related through three relations: `hasSubTopic`, its inverse `isSubTopicOf`, and `hasRelatedTopic`. The `hasSubTopic` relationship is transitive and denotes that one topic is the parent of another. Figure 3 shows a subset of classes and their relations from the DO, and corresponding topics and their relations as implemented in KWG. In this example, `topic.Disease` is the parent topic of `topic.Lipoma` as derived from their subclass relationship in the DO. A parent topic need not completely encompass all the relevant information of a child topic. In other words, unlike with a more formal class-subclass hierarchy, parent topics need not be strict super-sets of their sub-topics. For instance, there may be a compound topic that is a sub-topic of two separate parent topics—the topic *DiabetesDisasterResponse* can be represented as a sub-topic of both *Disease* and *DisasterManagement* topics. Specialized relationships between topics are denoted using `hasRelatedTopic`, a symmetric relation. For instance, relationships between diseases and anatomical features in the Open Biological and Biomedical Ontology (OBO) are made using the specific `obo:derivesFrom` predicate, as seen in Figure 3. This specific semantic relation is captured using the general `hasRelatedTopic` relation between `topic.Lipoma` and `topic.FatCell`. Because we only have two formal types of relations, denoting hierarchical and relatedness relationships between topics, we follow standard reification using the `TopicConnectednessDescription` class to encode other semantics such as specific semantics of relationships, provenance, reference to external classes, etc. Topic instances can also be linked to their corresponding class in an external ontology using metadata properties such as `rdfs:label`, `dc:references`, `prov:hadPrimarySource`, etc. Simple Knowledge Organization System (SKOS) and the Scientific Taxonomy Pattern [26] are other representation schemas that can be followed along with the ExO to translate informal vocabularies/taxonomies such as the UNDRR HIP classification into a formal topic hierarchy. The `referencesConcept` property can be used to relate a topic with a concept in an external ontology.

3.2. Asserting Expertise

The `Expert` class is meant to include individuals, and groups as needed, and is denoted as a subclass of `foaf:Person` or `foaf:Group` in Figure 4. The `hasExpertise` property indicates the relation between an expert and a topic. `ExpertiseRelation` is a reification of this property so that provenance and different measures of an expert's expertise and how they change over time can be attached. A numerical attribute to assert an expert's degree of expertise on a topic is represented using the data property `quantitativeExpertiseLevel`. The object property




Figure 4: Schema diagram of the ExO representing three different views of an expert’s expertise. Color and shape usage is the same as in the previous diagram. In addition, grey boxes with a dashed border represent controlled vocabularies (i.e., classes that have been defined as a set of individuals). Orange ellipses are literals.

qualitativeExpertiseLevel can be used to assert a categorical assignment of expertise (e.g., *academic* or *field*).

The ExO provides three objective views of asserting an agent’s expertise. They are denoted in the schema diagram in Figure 4 and are as follows:

1. The first category of experts is primarily from the academic or broader scientific community whose expertise can be determined from their publication history. Their expertise on a topic can be represented quantitatively and/or as a qualitative description. For instance, in pilot work, we are developing a similarity metric algorithm that adopts embedding-based natural language processing methods to represent expertise topics and experts as vectors derived from their publications. Then cosine similarity is computed over such vectors, which can be viewed as an expert’s degree of expertise in a topic.
2. The second category of experts is asserted based on their job role or organization affiliation using the reified Affiliation class, for example, a health professional’s affiliation as a program director or a trauma surgeon within a hospital. Organizations may also have an associated specialty. Hospitals, for example, specialize as surgical centers, or cardiac facilities. Both affiliations and specialties can be related to relevant topics using the fallsUnderTopic property.
3. The third category of experts is asserted based on activity-oriented facts, such as job-performance assessments or volunteer activities, using the engagedInActivity property




Figure 5: Schema diagram showing the different ways in which expertise is spatially scoped.

or the reified ActivityRelation class. The fallsUnderTopic property relates activities to topics.

3.3. Annotating Knowledge Graph Content with Topics

The fallsUnderTopic relation can be used to annotate instances in the graph, such as those pertaining to events, activities, organization specialties, and affiliations, with relevant topics. These connections can be identified through document embeddings to determine the semantic similarity between entities. Through punning, classes can also be linked to specific topics using this property. The property chain axiom described over fallsUnderTopic and isSubTopicOf infers relations between these instances/classes with topics above in the hierarchy.

3.4. Spatial Scoping

Searching for people or organizations for disaster relief purposes may follow one of the following schemes as illustrated in Figure 5:

1. Start from a disaster and its spatially defined region, and then look for experts located there who have expertise in topics of interest; can facilitate fast, on-site response; or have local knowledge of logistical challenges. This knowledge is represented in two ways: through modeling a) the geographic location of the agent using the hasGeographicLocation relation, and b) the geographic location of the organization that a person or team is affiliated with using the locatedIn relation.
2. Search for experts with expertise in topics that pertain to an area of interest, which may be different from their actual geographic location. For example, a person may have experience or familiarity with hurricanes in Puerto Rico even though they reside in Arizona. In this scenario, spatial context is attributed to a topic using the hasSpatialAssociation relation.
3. The spatial scope of an expert's expertise in a phenomenon/topic can be different from the actual geographical coverage of a phenomenon. For example, agent-A may only have experience studying the coastal impacts of hurricanes along the Gulf of Mexico, which may cover only a fraction of the full spatial extent of a hurricane event. The spatial scope of the disaster phenomenon does not necessarily reflect the areas of secondary or tertiary impacts where assistance is often needed. This sort of segmented spatial scoping is made possible through the hasSpatialAssociation relation on the reified ExpertiseRelation class.




Figure 6: Schema diagram showing the different ways in which expertise is temporally scoped.

3.5. Temporal Scoping

The ExO includes object properties as illustrated in Figure 6 to temporally scope 1) the assertion about an expert’s expertise, and 2) the expert’s affiliation with an organization. An expert’s expertise area may shift over time, as may their level of experience in a certain topic. For instance, Robert was engaged in researching the interaction between hurricane impacts and public health from 2001–2010 but then shifted his research focus to wildfire disaster response after 2010. Or Robert may have had theoretical knowledge of wildfire response processes up until 2010, but since then he has developed advanced experience due to engaging in response activities in the field. The relation `particularToTemporalPeriod` can be used to temporally scope the reified `ExpertiseRelation` class in order to accurately describe an expert’s focus across time. On the other hand, if one were to locate an organization and look up specific operational staff, e.g., an epidemiologist, this can be identified through their active period—represented using the `activePeriod` relation.

4. Evaluating ExO for the Humanitarian Aid Scenario

Finally, we return to the use-case scenario and competency questions described in Section 2.1. We demonstrate how the ExO can be used to *represent* the required information and to *answer* the competency questions with SPARQL queries to verify the ontology. To illustrate, we explicitly used two datasets, of which a subset is shown in Figure 7. The two datasets represented with the ExO are made available in an example KG along with the CQ SPARQL queries [here](#).

The first dataset contains experts from the academic viewpoint, who have relevant expertise in disaster response. This is shown using green boxes in Figure 7. For instance, `expert.101` refers to ‘Robert Olshansky’, who has expertise in the topics of Disaster Response and Natural Hazards. We mined this information from publications in Google Scholar, where the quantitative measure of expertise was determined using the expert-similarity algorithm from [11] on a set of their publications. Details about their employment and affiliation are manually extracted from the organization’s home page. The process of automatically extracting spatiotemporal expertise scope is a future direction of this work.

The second dataset consists of US Federally Qualified Health Centers (FQHC)⁵ with FTE (full-time equivalent) information and medical specialties of affiliates—e.g., dentists, clinicians, and other medical staff—who are potentially relevant experts needed for humanitarian aid responses. This is shown using orange boxes in Figure 7. For example, `fqhc.010220`, which refers to ‘Generations Family Health Center, Inc.’, specializes in dentistry and is therefore linked to corresponding medical specialty concepts from the MESH ontology [25]. This specialty concept is then related to topics such as Oral Medicine. The process of automatically linking a

⁵<https://data.hrsa.gov/tools/data-reporting/program-data>




Figure 7: Example of some instance data that populates a portion of the ontology.

specialty (e.g., Dentistry) or a specific role (e.g., Dental Hygienist) of an expert at this location with all the relevant topics in the topic hierarchy is a future direction of this paper.

Descriptions outlining a person’s professional or volunteer experience could help characterize their level of expertise. However, this information is often unavailable, mainly due to privacy concerns. Therefore, we have not yet been able to represent such information using the ExO.

5. Related Work

Existing expert lookup systems [14, 16, 27, 28] only help search for people with academic, professional, or scientific expertise on general topics or academic fields. Textual contents from scientific publications have been used as the primary source for generating learned representations of expertise [29, 30, 31] for topics that are fine granular (e.g., disaster recovery related to Hurricane Katrina), topics spanning domains (e.g., humanitarian aid for disaster victims afflicted by water-borne diseases), or spatially or temporally scoped topics (e.g., disaster response for hurricanes along the Pacific coast). [32] went a step further, enriching expert profiles with social background information. However, topics in all these works still follow a flat hierarchy, which makes it difficult to infer experts who could potentially be of interest when data are sparse.

Existing ontologies that model relationships between experts and their expertise topics are inadequate to model the variety of people who can be considered experts as reviewed in Section 2.2 and elsewhere in this paper. For example, the Semantically-Interlinked Online Communities ontology [33] describes online communities and the people, content, and activities associated with them. It includes properties such as names, emails, affiliations, and expertise, but focuses on modeling relationships between experts and the communities they participate in. Therefore, it has no properties to comprehensively connect topics. The Human Resources

Management ontology [34] includes classes such as job roles, qualifications, and competencies, and is meant to model the expertise and qualifications of job candidates and employees, but it is intended specifically for representing human resources information. The Information Artifact Ontology [35] models expert profiles and resumes. It includes concepts such as expertise-statement, experience-statement, and education-statement as well as properties for describing the degree of confidence and the temporal validity of the statements. However, this ontology is specifically for modeling information artifacts and is not geared toward expert profiles. The Ontology for Competence Management [36] includes concepts such as job position, skill-category, and competence level but is intended for human resources management in terms of describing the assessment and the development of competencies. Moreover, none of these ontologies model the spatiotemporal aspect of expertise, topics, or an expert's affiliation.

6. Conclusion

A critical but time-consuming task during emergency management, specifically for humanitarian aid organizations, is to reach out to experts with the appropriate knowledge and skills to respond efficiently to a dynamic situation. Existing expert lookup systems only recommend academic experts, which is too narrow for disaster response as recommendations should also include affiliates of non-governmental organizations that are central in the emergency response network (e.g., Red Cross), people who have engaged in relief work in a particular geographic area, liaisons to the affected communities, agencies (e.g., public health departments) with relevant expertise in or responsibility for a particular geographic area, and so on. However, there are no structured datasets that provide this kind of information. A manual approach would require searching through web pages and interpreting their contents, identifying individuals' fields of expertise, and then aligning them with some predefined topic hierarchy. This inefficient approach would limit the scalability and timeliness of an expert-expertise network to be used during a disaster response. As such, this paper describes the Expertise Ontology as an initial step in addressing this larger agenda of constructing a scalable expert-expertise knowledge graph. The ontology provides for a broad scope of what constitutes an expert, i.e., based on an agent's a) academic expertise; b) work- or volunteer-related responsibilities or experience; and c) organizational specialization(s) or affiliation(s). Further, we propose qualitative and quantitative ways of ascribing expertise to an agent alongside other details such as the provenance of how it was derived, and the spatial and temporal scope of the expertise.

While this paper provides an ontology to represent experts and their expertise in a knowledge graph, tackling some challenges as future work is critical to building a scalable expert-expertise network. A major challenge we envision going forward is a method to uniformly represent both academic and non-academic experts in the graph that will allow consistently retrieving them—that is, by enabling the search of experts from a singular starting point such as a topic of interest. As mentioned, due to the non-standardized nature of data sources for non-academic expertise, materializing the links between topics and organizations or their affiliates (using the `fallsUnderTopic` relation mentioned in Section 3.2) requires us to explore ontology alignment and KG embedding techniques further. It is also clear that “expertise” is reflected differently online between academic and non-academic cases; therefore, experts from each must be identified differently, e.g., mining scientific literature for academic experts vs. mining a host of heterogeneous data for non-academic experts. While we have some clarity on how to ascribe qualitative

levels and quantitative metrics to academic experts, e.g., by mining scientific literature, the process of doing the same for non-academic experts is our next work step. Developing an algorithm that can calculate non-academic expertise based on professional credentialing and experience is needed. However, since all of the concepts are modeled in the ontology and can be relatively easily represented in a KG, a potential avenue we hope to explore would be to implement rules, e.g., in SWRL to attribute “expertise” in the non-academic context based upon attributes of institutions, institutional roles, populations, and areas. Finally, another element of future work is to explore how to determine the spatial extent and temporal duration of a person’s expertise from their publications.

References

- [1] M. Stichter, Philosophical and psychological accounts of expertise and experts, *Humana. Mente Journal of Philosophical Studies* 8 (2015) 105–128.
- [2] A. I. Goldman, Expertise, *Topoi* 37 (2018) 3–10.
- [3] K. Janowicz, P. Hitzler, W. Li, D. Rehberger, M. Schildhauer, et al., Know, Know Where, KnowWhereGraph: A densely connected, cross-domain knowledge graph and geo-enrichment service stack for applications in environmental intelligence, *AI Magazine* 43 (2022).
- [4] N. F. Noy, D. L. McGuinness, et al., *Ontology development 101: A guide to creating your first ontology*, 2001.
- [5] R. Battle, D. Kolas, GeoSPARQL: enabling a geospatial semantic web, *Semantic Web Journal* 3 (2011).
- [6] J. R. Hobbs, F. Pan, Time ontology in OWL, W3C working draft 27 (2006) 3–36.
- [7] K. Janowicz, A. Haller, S. J. Cox, D. Le Phuoc, M. Lefrançois, SOSA: A lightweight ontology for sensors, observations, samples, and actuators, *Journal of Web Semantics* 56 (2019).
- [8] R. Zhu, C. Shimizu, S. Stephen, L. Zhou, L. Cai, G. Mai, K. Janowicz, M. Schildhauer, P. Hitzler, SOSA-SHACL: Shapes Constraint for the Sensor, Observation, Sample, and Actuator Ontology, in: *The 10th International Joint Conference on KGs*, 2021.
- [9] R. Zhu, S. Stephen, L. Zhou, C. Shimizu, L. Cai, G. Mai, K. Janowicz, P. Hitzler, M. Schildhauer, Environmental observations in knowledge graphs., in: *DaMaLOS*, 2021, pp. 1–11.
- [10] Z. Liu, Z. Gu, T. Thelen, S. G. Estrecha, et al., Knowledge explorer: exploring the 12-billion-statement knowwheregraph using faceted search (demo paper), in: *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, 2022.
- [11] R. Zhu, L. Cai, G. Mai, C. Shimizu, C. K. Fisher, K. Janowicz, A. Lopez-Carr, A. Schroeder, M. Schildhauer, Y. Tian, et al., Providing humanitarian relief support through knowledge graphs, in: *Proceedings of the 11th on Knowledge Capture Conference*, 2021, pp. 285–288.
- [12] S. Ghosh, N. Sharma, F. Benevenuto, N. Ganguly, K. Gummadi, Cognos: crowdsourcing search for topic experts in microblogs, in: *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, 2012, pp. 575–590.
- [13] H. Jung, M. Lee, I.-S. Kang, S. Lee, W.-K. Sung, Finding topic-centric identified experts based on full text analysis., *FEWS* 290 (2007) 56–63.
- [14] Expertise Finder Corp., ExpertiseFinder, <https://expertisefinder.com/>, accessed 2023.
- [15] Elsevier, Expert Lookup, <https://www.elsevier.com/solutions/>, accessed 2023.

- [16] LexisNexis, Expert Research On-Demand, <https://www.lexisnexis.com/experts-on-demand/order-searches/expert-witness-search>, accessed 2023.
- [17] SEAK, Expert Witness Directory, <https://www.seakexperts.com/>, accessed 2023.
- [18] L. M. Schriml, C. Arze, S. Nadendla, Y.-W. W. Chang, M. Mazaitis, et al., Disease ontology: a backbone for disease semantic integration, *Nucleic acids research* 40 (2012) D940–D946.
- [19] UNDRR Hazard definition and classification review (Technical Report), <https://www.undrr.org/publication/hazard-definition-and-classification-review>, 2021.
- [20] UNDRR - Hazard Information Profiles: Supplement to UNDRR-ISC hazard definition & classification review - technical report, <https://www.undrr.org/publication/hazard-information-profiles-supplement-undrr-isc-hazard-definition-classification>, 2021.
- [21] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, *Prov-o: The prov ontology* (2013).
- [22] D. Brickley, L. Miller, *Foaf vocabulary specification 0.91*, 2007.
- [23] W3C Recommendation, *The Organization Ontology*, <https://www.w3.org/TR/vocab-org/>, accessed 2023.
- [24] P. L. Buttigieg, N. Morrison, B. Smith, C. J. Mungall, S. E. Lewis, The environment ontology: contextualizing biological and biomedical entities, *Journal of biomedical semantics* 4 (2013).
- [25] C. E. Lipscomb, Medical subject headings (mesh), *Bulletin of the Medical Library Association* 88 (2000) 265.
- [26] S. Stephen, C. Shimizu, M. Schildhauer, R. Zhu, K. Janowicz, P. Hitzler, A pattern for representing scientific taxonomies, *Proceedings of the 13th Workshop on Ontology Design and Patterns*, co-located with the 21st International Semantic Web Conference (2022).
- [27] D. Choi, H. Lee, K. Bok, J. Yoo, Design and implementation of an academic expert system through big data analysis, *The Journal of Supercomputing* 77 (2021) 7854–7878.
- [28] C. Albusac, L. M. d. Campos, J. M. Fernández-Luna, J. F. Huete, Pmsc-ugr: A test collection for expert recommendation based on pubmed and scopus, in: *Conference of the Spanish Association for Artificial Intelligence*, Springer, 2018, pp. 34–43.
- [29] N. Nikzad-Khasmakhi, M. Balafar, M. R. Feizi-Derakhshi, The state-of-the-art in expert recommendation systems, *Engineering Applications of Artificial Intelligence* 82 (2019).
- [30] M. Fazel-Zarandi, M. S. Fox, An ontology for skill and competency management, in: *Formal Ontology in Information Systems*, IOS Press, 2012, pp. 89–102.
- [31] M. Gruza, Classification of european skills/competences, qualifications and occupations (esco) and its link with information about occupations, *Polityka Społeczna* 20 (2018) 10–15.
- [32] E. Davoodi, K. Kianmehr, M. Afsharchi, A semantic social network-based expert recommender system, *Applied intelligence* 39 (2013) 1–13.
- [33] J. G. Breslin, U. Bojārs, A. Passant, S. Fernández, S. Decker, SIOC: Content Exchange and Semantic Interoperability Between Social Networks, <https://www.w3.org/2008/09/msnws/papers/sioc.html>, 2009.
- [34] J. Dorn, T. Naz, M. Pichlmair, Ontology development for human resource management, in: *Knowledge Management: Innovation, Technology and Cultures*, World Scientific, 2007.
- [35] W. Ceusters, B. Smith, *Aboutness: Towards foundations for the information artifact ontology* (2015).
- [36] S. Miranda, F. Orciuoli, V. Loia, D. Sampson, An ontology-based model for competence management, *Data & Knowledge Engineering* 107 (2017) 51–66.