

Article

A Methodology for Building a Medical Ontology with a Limited Domain Experts' Involvement

Sabrina Azzi 

School of Business & Creative Industries, University of the West of Scotland, High St, Paisley PA1 2BE, UK; sabrina.azzi@uws.ac.uk

Abstract: Ontology development is a multidisciplinary work involving domain experts and knowledge engineers. Bringing together such a team to develop an ontology of quality is not easy. Therefore, ontologies are often created with limited expertise either in the medical domain or in ontology engineering. Unfortunately, the existing methodologies do not provide much guidance on how the different steps of ontology development should be performed, particularly in the case of reduced involvement of domain experts. This challenge is getting more difficult when there is a multitude of medical knowledge sources and ontologies covering parts of the domain, and often, each has a different representation of the same concept, for example, as a symptom, disease, or clinical sign. This research presents a methodology for creating a medical ontology of quality with limited involvement of the domain experts. The latter are only consulted in the domain definition and evaluation phases. We combine building an ontology from codified knowledge and ontology reuse to enhance reusability and interoperability. The methodology is inspired by METHONTOLOGY, for which we make several improvements, especially in the ontology reuse phase. We provide proof of concept of the proposed methodology with a case study involving the development of the pneumonia diagnosis ontology (PNADO).

Keywords: medical ontology engineering; ontology reuse; clinical practice guidelines; OBO Foundry; competency questions; pneumonia diagnosis



Received: 20 March 2025
Revised: 29 April 2025
Accepted: 21 May 2025
Published: 28 May 2025

Citation: Azzi, S. A Methodology for Building a Medical Ontology with a Limited Domain Experts' Involvement. *Digital* **2025**, *5*, 18. <https://doi.org/10.3390/digital5020018>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The contribution of ontologies to understanding, sharing, and integrating knowledge is well established. The usage of ontologies in medicine gains ground thanks to their ability to represent the medical domain knowledge and their expressive power formally [1]. An ontology can assure the interoperability and improve the functionality of a clinical decision support system by providing a standard vocabulary to help integrate heterogeneous biomedical data sources [2].

The ontology-building process is composed of several steps. Each step requires a set of comprehensive instructions and techniques on how it should be carried out, and together, they form an ontology-building methodology. The completeness and the degree of detail of the latter are very important to build an ontology of quality.

Medical ontologies have significantly grown over the last decade, but they have quality and content problems [2–4]. This is due, among others, to the following causes: (1) Medical ontologies are created with limited expertise either in the medical domain or in ontology engineering; for example, Coronavirus Infectious Disease Ontology (<https://biportal.bioontology.org/ontologies/CIDO>, accessed on 8 May 2025) is created by a medical expert, whereas COVID-19 Ontology (<https://biportal.bioontology.org/ontologies/COVID-19>,

accessed on 8 May 2025) exhibits the issue of misclassification of signs and symptoms, a common challenge in clinical knowledge representation; (2) they are often built from scratch and do not take advantage of reusing existing ontologies that have already been validated to save time and improve the interoperability; (3) the applied ontology-building methodology lacks some important steps like evaluation or validation, or during the building process, some steps are skipped.

Many methodologies have been proposed to build an ontology from a corpus of text and/or by reusing elements of existing ontologies, which is called ontology reuse. The latter is the process of selecting existing ontologies and manipulating them in some way to meet the requirements and needs defined in the conception [5].

A key limitation of methodologies focused on building ontologies from text is their lack of clear guidance on how to execute the different steps. This can result in methodologies that are difficult to apply, time-consuming, error-prone, and costly [6], particularly in cases where domain expert involvement is limited. Although they recognize the importance of ontology reuse, most methodologies address it only briefly.

On the other hand, ontology reuse reduces the ontology development time and required effort since the knowledge engineer can focus on new content specific to the ontology being created rather than developing content that has already been created and validated [6]. Besides, the reuse of ontologies increases interoperability [7], reduces redundancy [8], and prevents the random proliferation of ontologies. However, ontology reuse has the drawback of reproducing the errors of the reused ontologies. During the reuse process, inconsistencies and conflicts between concepts represented differently across ontologies have to be fixed. Even though ontology reuse is recommended [9], to the best of our knowledge, there are currently no methodologies that provide adequate support for the ontology reuse process [10] and resolve the problems identified above. Another important problem related to the reuse is the lack of documentation explaining the conceptual model of the ontology and/or the design criteria that were followed to build and evaluate the ontology [11]. This problem concerns most of the published medical ontologies.

In this paper, we propose a methodology designed for knowledge engineers to build a medical ontology that deals with the interoperability and conflicts between concept representations, with minimal involvement of medical experts. Depending on the ontology domain, the methodology relies on the ability to access already structured and validated medical knowledge sources. Examples of such knowledge for a medical domain include clinical practice guidelines, systematic reviews, and gold standard frameworks.

Our methodology guides the knowledge engineer step by step by providing the necessary instructions. It consists of seven phases: (1) definition of ontology domain and scope; (2) building of knowledge corpus and extraction of terms; (3) building of preliminary ontology by applying ARCHONTE methodology on the terms previously identified; (4) reuse of ontologies and resolution of conflicts using competency questions; (5) evaluation of the ontology; (6) documentation; and (7) maintenance and evolution. The whole process is executed with minimal but targeted involvement of the medical experts. We illustrate our ideas with a case study of building pneumonia diagnosis ontology (PNADO). Pneumonia is one of the diseases associated with many diagnostic errors [12,13]. To the best of our knowledge, PNADO is the first comprehensive ontology of pneumonia diagnosis. Several of the existing medical ontologies, such as Human Disease Ontology (DOID: <http://www.obofoundry.org/ontology/doid.html>, accessed on 8 May 2025), SNOMED-CT (<https://browser.ihtsdotools.org/?>, accessed on 8 May 2025), and Human Phenotype Ontology (HPO, <http://www.obofoundry.org/ontology/hp.html>, accessed on 8 May 2025), contain concepts related to the diagnosis of pneumonia; however, they provide different representations for some concepts, and they do not capture the entire domain knowledge.

2. State of the Art

Most of the first ontology engineering methodologies were developed as a by-product of ontology development for some specific projects. ENTERPRISE [14,15] and Toronto Virtual Enterprise (TOVE) [16] are examples of projects that relate to the domain of enterprise modeling, developed in 1995. TOVE concentrates on modeling the internal workflow of a company. The methodology principle is that the authors address ontology completeness based on competency questions. ENTERPRISE methodology, also called Uschold and King methodology, is based on the experience of developing the Enterprise Ontology. The four phases that are recommended in the methodology are: (1) identify the purpose, scope, and scenarios for its use; (2) build the ontology, including ontology capture, coding, and integrating existing ontologies; (3) document the ontology; and (4) evaluate the ontology. A refined version of this methodology was proposed in 1996 by Uschold and Gruninger [17]. In the same year, a well-structured methodology called METHONTOLOGY [18] was developed in the Laboratory of Artificial Intelligence at the Polytechnic University of Madrid. The phases proposed in METHONTOLOGY are: (1) specification that defines the purpose and scope of the ontology; (2) knowledge acquisition; (3) conceptualization that consists in structuring the domain knowledge in a conceptual model that describes the problem and its solution according to the domain vocabulary identified in the specification phase; (4) implementation; (5) integration; (6) evaluation; and (7) documentation. ARCHONTE is another methodology [19] that was proposed for ontology design based on the principle of semantic commitment, meaning that the ontology's structure should reflect a clear, formal commitment to the meaning of its concepts; however, it lacks concrete and operational guidance on how to systematically perform its different steps.

The main shortcoming of these methodologies is that they consider texts as the only sources of knowledge, and they do not provide guidance on how the different steps should be conducted. An analysis of these methodologies for building ontologies was conducted in [11]. The study demonstrated that none of them are fully mature comparing them with IEEE standards [20]. However, METHONTOLOGY is the most mature, but some steps require details.

101 method provides guidelines to manually develop domain ontologies [21]. It provides strategies to assist in making design decisions during ontology construction. It declares the possibility of reusing ontologies without specifying how. Also, the methodology focuses on defining class hierarchies, properties of classes and instances, and omits an important step, which is the evaluation.

Since developing an ontology from text is difficult and time-consuming, frameworks and tools were also proposed to support the user in the process of building. They use machine learning, linguistic, logic, and statistics techniques [22]. ASIUM [23], TextToOnto [24], Text2Onto [25], HASTI [26], SYNthesis of DIstributed Knowledge Acquired from TEXTs (SYNDIKATE) [27], ProMine [28], OntoLearn [29], OntoLT [30], OntoGen [31], and Terminae [32] are examples of these tools. However, all these tools suffer from several shortcomings, such as the difficulty of extracting relations and axioms. For example, Text2Onto struggles with very large corpora and often extracts concepts and relationships based primarily on shallow linguistic patterns (e.g., co-occurrence and syntactic parsing) rather than deep semantic understanding. This can lead to superficial ontologies that lack the rich structure and meaning needed for complex domains, like medicine. Authors of [22] recommend using deep learning rather than shallow learning for automatic ontology construction.

While ontology reuse is supposed to facilitate and speed up the development process, it remains a very important issue, especially in the biomedical domain where large resources are available. Several methodologies have been proposed to deal with this issue:

MetRON was suggested to build a reusable ontology by reusing the existing ontologies and other resources, such as classifications, vocabularies, terminologies, and standards [33]. Top-down, bottom-up, middle-out approaches, and even a combination of them are proposed to be used for defining classes.

A methodology for biomedical ontology reuse was proposed in [8] and applied for Abdominal Ultrasound Ontology development. It provides steps and tools for finding and choosing relevant ontologies to reuse as well as concepts mapping.

A methodology was proposed in [34] to build a multilingual domain ontology by reusing termino-ontological resources and text corpora. The methodology was applied to build an Alzheimer's disease ontology.

Authors in [35] proposed seven steps to develop an ontology for guiding appropriate antibiotic prescribing intended for a clinical decision support system. These steps are: (1) define the domain ontology and scope; (2) review existing ontologies for reuse; (3) create classes and properties; (4) create a conceptual model; (5) select and implement upper ontology; (6) implement the ontology in a formal representation; and (7) evaluate the ontology. The methodology evaluates the resulting ontology using Semantic Web Rule Language (SWRL) rules.

OntoReuse, to assist ontology engineers in choosing which ontology to reuse, was proposed in [36]. It integrates lexical evaluation, structural evaluation, maturity evaluation, and fairness evaluation.

All these methodologies lack guidance in carrying out the steps and do not raise the problem of conflicts between concept representations and how to resolve them.

3. Materials and Methods

In this section, we present a methodology inspired by METHONTOLOGY [18] to which we made several improvements. Figure 1 presents the seven phases of the methodology and shows that domain experts are only involved in the definition of the ontology domain and scope phase and the evaluation phase.




Figure 1. Methodology phases.

3.1. Phase 1: Ontology Domain and Scope Definition

This phase consists of establishing the domain of interest of the ontology to model and its scope. This can be done either by consulting domain experts (physicians) or by finding a gap in the literature. First, two fundamental questions must be addressed:

1. What is the domain that the ontology will cover?
2. What is the purpose of creating this ontology?

Second, functional requirements, defined as the questions the ontology should be able to answer, are defined. To support this, competency questions (CQs) are established by physicians and/or extracted from codified knowledge sources (Section 3.2). The CQs have the form of informal queries that a particular community of users think that the ontology should answer [16]. They are used in many methodologies for ontology building to specify requirements [17,18,21]. The CQs are a sketch and do not have to be exhaustive.

At the end of this phase, the knowledge engineer has a clear vision of what the ontology will cover, for what the ontology will be used, and what information it will contain.

3.2. Phase 2: Corpus Building and Term Extraction

In this phase, the knowledge engineer builds the corpus of knowledge extracts and identifies relevant terms to the ontology domain defined in phase 1. The use of codified knowledge makes this task easier for the knowledge engineer since this type of source is more reliable and comprehensible. Since the quality of ontology depends on the quality of the knowledge corpus, we recommend that for a medical domain, the knowledge corpus be composed of the most reliable and complete sources, such as clinical practice guidelines (CPGs) or systematic reviews related to the medical domain for which the ontology is being built. Another reason to use codified knowledge is that there are no available annotated texts that could be used as a base to create a medical ontology.

Systematic reviews investigate and summarize scientific literature in a systematic and methodologically rigorous way. Each of them focuses on a specific question and uses methods to study and summarize the findings of similar but separate studies. They are considered to be the most important type of scientific review as they are fundamental to evidence-based practice [37]. They are the basis of the CPGs. The latter are medical documents resuming the state of the art in the management of pathologies and clinical cases, which we call basic recommendations. They are consulted by healthcare practitioners to guide diagnostic and management decisions. They are based on the concept of evidence-based medicine (EBM). In [38], EBM was defined as “*the conscientious, explicit, and judicious use of current best evidence in making decisions about the care of individual patients*”. The CPGs are found in libraries such as Cochrane (<https://www.cochranelibrary.com/>, accessed on 8 May 2025), National Institute for Health and Care Excellence (NICE: <https://www.nice.org.uk/guidance>, accessed on 8 May 2025) and other libraries that are more specific to a specific condition. In fact, given the high degree of completeness of the CPGs regarding the most important concepts, they can be integrated into clinical decision support systems, thus providing recommendations that can be adapted to the patient’s profile at the time of decision making [39].

We distinguish between a term and a concept. According to ISO 1087:2019 (ISO 1087:2019(en), Terminology work and terminology science—Vocabulary, <https://www.iso.org/obp/ui/#iso:std:iso:1087:ed-2:v1:en>, accessed on 23 May 2025), a term is a designation that represents a general concept by linguistic means, whereas a concept is a unit of knowledge created by a unique combination of characteristics. We represent concepts corresponding to the extracted terms in the ontology in phase 3 (Section 3.3).

Term extraction can be manual or automatic. Some of the tools presented in Section 2 could be used for automatic term extraction from a corpus of knowledge constituted of

CPGs. Unfortunately, several of these tools, such as Terminae, ProMine, OntoLT, and ASIUM, are no longer available for use. There are also powerful tools for clinical text analysis, such as cTAKES (<https://ctakes.apache.org/index.html>, accessed on 8 May 2025). They could be used for term extraction; however, their installation and usage are quite heavy. Indeed, cTakes uses full clinical natural language processing (NLP) pipelines (sentence detection, tokenization, parsing, and entity linking), making it slower and heavier but very accurate. We suggest using Text2Onto [25], which is more suitable for this task. It is available and user-friendly, and its installation and usage are simple. Text2Onto has been developed to support the acquisition of ontologies from textual documents, allowing flexible extraction of candidate concepts without heavy reliance on pre-existing medical models. It is composed of modules that extract terms, relationships (equivalence relationship, hierarchical, etc.), and instances. It offers extraction algorithms calculating the following measures: relative term frequency (RTF), term frequency inverted document frequency (TFIDF), entropy, and the C-value/NC-value. We recommend using TFIDF when there is more than one document because it evaluates how relevant a term is to a document in a collection of documents.

The extracted terms that should be analyzed are nouns, proper nouns, verbs, adjectives, adverbs, and phrases. These terms represent the concepts that will be treated in the next phase.

It is important to note that the extraction of relations and axioms is a complex task, and most existing learning systems do not infer new relations or axioms [22]. According to [22], deep learning techniques hold significant potential for effectively learning and inferring both relations and axioms.

3.3. Phase 3: Building of Preliminary Ontology

This phase is dedicated to the building of preliminary medical ontology. This preliminary ontology represents the intended models of the requirements defined in phase 1. Usually, an ontology can have many interpretations. To limit the number of possible models, it uses axioms. Nevertheless, it is usually impossible to create an ideal ontology whose models coincide with the intended ones. An intended model represents all the requirements and only them [40].

This study adopts the ARCHONTE methodology to build a preliminary ontology [19]. The strong point of this methodology is the recommendation of semantic commitment in the design process. In other words, to explain the meaning of each of the ontology concepts represented by the terms already selected in phase 2, the knowledge engineer expresses, in natural language, the similarities and differences that the concept has with those close to it. The ontology structure is similar to a tree, which facilitates the determination of the meaning that a concept has according to its position. ARCHONTE is composed of three steps: (1) semantic normalization; (2) knowledge formalization; and (3) toward a computational ontology. They are elaborated below.

3.3.1. Step 1: Semantic Normalization

The objective of this step is to reach a semantic agreement about the meaning of the concept labels. The knowledge engineer applies differential principles on the set of candidate concepts chosen previously. According to the differential paradigm proposed by Bachimont [19], the meaning of a concept (or a node, since an ontology is structured as a tree) is determined by its closest neighbors: parent and its siblings, i.e., its position in the ontological hierarchy based on terminological structure. Similarities and differences of each concept concerning its parent and its siblings are expressed in natural language. Four differential principles are distinguished: (1) similarity with each parent; (2) similarity with

each sibling; (3) difference with each sibling; and (4) difference with each parent. The result of this step is a so-called differential ontology. A differential ontology is an ontology where each concept and relationship must be defined according to its similarities and differences with its parents and siblings. At the end of this step, the ontology is ready to be enriched with axioms.

3.3.2. Step 2: Knowledge Formalization

In this step, the referential ontology is considered from an extensional semantics point of view. Each concept is linked to a set of objects that allows defining new concepts and relations through set operations. During this step, the ontology engineer has to specify the arity and domains of the relations, in correspondence with the intended models. The ontology engineer can also add some logical axioms to constrain the domains of the relations. The result of this step is a so-called referential ontology. A referential ontology is a formal ontology that covers the concepts, relations, instances, and axioms for a domain. At the end of this step, the ontology is ready to be implemented.

3.3.3. Step 3: Toward a Computational Ontology

This last step of ARCHONTE methodology is implementing the referential ontology in an operational language of knowledge representation, such as OWL. This is achieved by adopting, on the one hand, a formalism of representation (conceptual graphs and description logic) and, secondly, by adapting the representation of the ontology to the objectives fixed in phase 2. This step marks the transition to a computational ontology. For interoperability and reusability purposes, we recommend following OBO Foundry principles. An upper medical ontology that provides the most general concepts in the medical domain can be reused to enhance interoperability with other medical ontologies. Basic Formal Ontology (BFO) is widely used as an upper ontology in OBO Foundry and the biomedical domain [41]. This makes BFO the best choice for building any medical ontology to enhance interoperability. It aims to model the basic structures of reality and provides classes that categorize, at a high level, real-world entities. BFO is designed to support information retrieval, analysis, and integration in several domains. It has two basic types of entities: (1) continuants that are entities that continue or persist through time, such as objects, qualities, and functions and (2) occurrents that are events or happenings in which continuants participate [42]. Also, there is usually an upper ontology for a specific domain that includes general classes related to this domain. Ontology for General Medical Science (OGMS) is one of these ontologies [43]. OGMS enriches BFO by including the concepts of general medical science. This consists of approximately 100 terms that describe the fundamental aspects of medicine, such as *disorder*, *diagnosis*, *disease*, *disease course*, *laboratory test*, *sign*, *symptom*, and *syndrome*.

At the development level, Protégé ontology editor, which offers the possibility to visualize the different aspects of the ontology, can be used.

For further enrichment of the ontology, Unified Medical Language System (UMLS: <https://www.nlm.nih.gov/research/umls/index.html>, accessed on 8 May 2025) can be used to extract definitions, synonyms, acronyms, and other annotations. Concepts in the ontology can be mapped to the referenced resource, UMLS concept unique identifiers (CUIs), where possible. Mapping the ontology to UMLS enables the ontology to be linked to many other relevant biomedical resources such as LOINC and SNOMED-CT. MetaMap [43] is a tool designed to identify biomedical concepts from free-text and maps them into concepts from the UMLS. It primarily uses a rule-based approach, leveraging linguistic and syntactic analysis techniques. MetaMap analyzes text through lexical lookup, variant generation, and candidate evaluation using a knowledge-intensive strategy to identify

the best matching concepts. Every concept has, when possible, CUI code, preferred label, definitions, synonyms, and external references of the reused ontologies. This step overlaps with the ontology reuse phase as reuse begins in phase 3 but is fully explored and formalized in phase 4.

3.4. Phase 4: Ontology Reuse and Enrichment

In this phase, we finalize the preliminary ontology by adding and detailing concepts. For this purpose, we reuse some terminological resources that we consider relevant to the ontology domain. Reuse can be hard, i.e., one imports the complete ontology to reuse, or soft, i.e., one imports the reused ontology concepts. The ontologies to be reused must be evaluated for the suitability of their content coverage and the depth of knowledge, as well as the potential to support the inference process. The phase is composed of three steps: (1) finding ontologies; (2) choosing relevant ontologies; and (3) resolving conflicts.

3.4.1. Step 1: Finding Ontologies

In this step, we select an ontology for reuse when we conjecture that its axioms could be used to respond to some of the requirements defined in phase 1. We use four main search criteria to identify candidate ontologies for reuse:

1. Both the natural and formal languages of the candidate ontology are the same as those of the ontology we are building.
2. There is a mapping from a non-empty subset of the new ontology requirements (defined in phase 1) to a robust set of concepts of the ontology candidate.
3. The candidate ontology is accepted within its user community. This acceptance is assessed by using some metrics such as the number of community members that endorse the ontology, the number of times the ontology has been reused, etc.
4. The candidate ontology has been published in a peer-reviewed publication.

Ontology repositories, such as NCBO BioPortal (<https://bioportal.bioontology.org/> (accessed on 8 March 2025)), OBO Foundry (<http://www.obofoundry.org/> (accessed on 8 May 2025)), Protégé Wiki (https://protegewiki.stanford.edu/wiki/Main_Page, accessed on 22 May 2025), Swoogle (ontology lookup service from the European Bioinformatics Institute: <http://ebiquity.umbc.edu/project/html/id/53/Swoogle>, accessed on 22 May 2025), and UMLS that integrates many terminologies and coding standards, are used to find relevant ontologies.

The tool of BioPortal, Ontology Recommender (<https://bioportal.bioontology.org/recommender>, accessed on 8 May 2025), which uses, among others, coverage and acceptance metrics, could be used to find ontologies in this repository.

3.4.2. Step 2: Choosing Relevant Ontologies

This step involves selecting the most relevant ontologies to be reused. For each significant topic found in the requirements defined in phase 1, we are looking for ontologies that deal with this topic, so-called relevant ontologies. Examples of such topics might be symptoms or laboratory tests.

The preference criteria for choosing relevant ontologies are accuracy and precision orderings [44]. To compare the accuracy and precision of ontology candidates, two kinds of models are considered:

Superfluous models (SUP), as defined in this context, correspond to a situation where some aspect of the candidate is weaker than required, i.e., they prevent some aspects of the requirements from being satisfied (see Figure 2) [44]. Note that “superfluous” here denotes components that are present but ineffective rather than excessive; for example, Coronavirus Infectious Disease (CIDO: <https://bioportal.bioontology.org/ontologies/CIDO/>

[?p=summary](#), accessed on 8 May 2025) is weaker than PNADO intended model because it does not satisfy some ontology requirements defined in phase 1 (i.e., some concepts defined in PNADO are not covered by CIDO).




Figure 2. Ontology enrichment and resolving conflicts approach.

Omitted models (OM), as defined in this context, correspond to a situation where some aspect of the candidate is stronger than is required, i.e., this candidate ontology exceeds what was specified in the requirements, i.e., includes concepts that go beyond the requirements; for example, Radiology Lexicon (RadLex) covers several unnecessary concepts for pneumonia diagnosis.

If there are two candidate ontologies T_1 and T_2 that deal with the same topic, we choose them as follows:

Case 1: If the set of superfluous models of T_1 is included in the set of superfluous models of T_2 and the set of omitted models of T_1 is included in the set of omitted models of T_2 , then T_1 is more *accurate* than the candidate ontology T_2 . We recommend choosing T_1 . If the set of superfluous models of T_1 is empty, hard reuse should be considered.

Case 2: If T_1 and T_2 have no superfluous models, then both are *precise*, and we reuse both. If the intersection of the sets of intended models is empty, then hard reuse should be considered. If not, hard reuse can be considered only if there are no conflicts between the representations of shared concepts.

Case 3: If T_1 has superfluous models and T_2 has not, then regardless of omitted models, T_2 is more *precise*, we reuse T_2 , and hard reuse should be considered.

Case 4: If both T_1 and T_2 have superfluous models, then they are incomparable using the accuracy and precision orderings, and reuse is problematic. If so, soft reuse should be considered.

If there are more than two candidate ontologies, we apply the same algorithm by comparing two ontologies at a time.

3.4.3. Step 3: Resolving Conflicts

Once the ontologies for reuse have been selected, we proceed to manually identify the concepts that will enrich our original ontology. We are interested in adding new concepts and completing the concepts already represented in our ontology.

A concept may be differently represented in different ontologies. In this case, there is a conflict, and we have to decide which representation should be used. To resolve the conflict, we use conflict resolution questions (CRQs) defined by the knowledge engineer. These CRQs are related to the meaning and the structure of the ontology concept tree, i.e., the hierarchy of classes, and other description logic (DL) concepts like intersection and union of classes, equivalent classes, universal classes, universal and existential quantification, has-value restriction, and cardinality restriction. We apply the principle that the concept axiomatization model is semantically correct if it does include an ontology intended model. Here are a few examples of CRQs:

1. Does the concept have a definition (in the form of annotation)? If yes, does its definition correspond to an intended model of the ontology?
2. Does the concept have a superclass? If yes, is it relevant to the ontology?
3. Does the concept have children? If yes, is each of them relevant to the ontology?
4. Does the concept have synonyms?
5. Does the concept have alternative names?
6. Does the concept have related names?

The knowledge engineer should propose and use any CRQ he deems suitable to determine the most appropriate concept representation.

To complete an already existing concept with annotations or add a new concept, we proceed as follows (see Figure 2):

- If the concerned concept is not found in any ontology, then we propose some relevant annotations that can be obtained from UMLS, such as definitions, synonyms, or concept unique identifier (CUI). If the concept is not found in UMLS, no annotation is added.
- If the concept is represented in only one ontology, then we reuse it if relevant, i.e., its representation fits the requirements, with its identifier and annotations. If not, we propose a new identifier. If the concept is found in UMLS, we propose to enrich it with annotations.
- If the concept is found in more than one ontology, then there is a conflict, and we have to choose the ontology that provides the best concept representation for the domain, i.e., the one that provides answers to the CRQs and corresponds to an intended model.
- If the chosen ontology is SNOMED-CT where concepts can have multiple parents and multi-hierarchies [45] that often cause user uncertainty in the selection of concepts and a messy situation in their classification [46], then we check if this is the case for

the concerned concept. If it has more than one hierarchy, then we identify the most relevant hierarchy or combine one or more hierarchies to build a new one. If not, we reuse the concept with its hierarchy.

- If the chosen ontology is not SNOMED-CT, reuse the concept as it is represented in the ontology.

3.5. Phase 5: Ontology Evaluation

Ontology evaluation is a process of assessing the quality of an ontology involving a set of evaluation criteria and assessing its adequacy for being used in a specific context for a specific goal. It is also referred to as “quality assurance”, or “auditing”, when it is conducted by a third party. Brank et al. [47] grouped existing evaluation approaches into four broad categories: (1) approaches that use the target ontology in an application and evaluate the application results [48]; (2) approaches that compare the target ontology to a gold standard [49]; (3) those that use data sources about a specific domain [50]; and (4) approaches that recommend a manual assessment by domain experts according to a set of criteria [51].

The evaluation process consists of several tasks dealing with the evaluation of different aspects of the ontology. The tasks can be divided into verification methods that examine the structure of the ontology (they answer if the ontology was built the right way) and validation methods that examine its applicability in the real world (they answer if the right ontology was built) [52,53]. It is worth noting that Brank’s evaluation approaches correspond to validation methods. More specific evaluation criteria have been proposed by [53,54] (see Table 1). Most of them correspond to verification methods.

Table 1. Ontology evaluation criteria.

Criteria	Definition
Accuracy	Does the asserted knowledge in the ontology agree with the expert’s knowledge, which is often measured in terms of precision and recall?
Completeness	Is the domain of interest appropriately covered?
Conciseness	Does the ontology include irrelevant or redundant axioms?
Consistency	Does the ontology include or allow for contradictions?
Computational efficiency	How fast can the reasoners work with the ontology?
Adaptability	How easy or difficult is it to use the ontology in different contexts?
Clarity	Does the ontology communicate effectively the intended meaning of the defined terms?

Several authors have suggested that automated tools are needed to ensure that high-quality ontologies are developed [55,56]. Authors in [1] emphasized the lack of tools for ontology evaluation. Most of the tools discussed in the literature are prototypes or proposals. To our knowledge, OOPS! (<http://oops.linkeddata.es/>, accessed on 27 May 2025), OAF (<https://njitsaboc.github.io/>, accessed on 27 May 2025), and OntoMetrics (<https://ontometrics.informatik.uni-rostock.de/ontologymetrics/index.jsp>, accessed on 27 May 2025) are the only available evaluation tools that can be used currently.

To evaluate each criterion listed in Table 1, we suggest using additional data sources. The evaluation of some criteria can be in part automatized, and for some others, it requires the involvement of the knowledge engineer and domain experts. The following

sections address these criteria using the verification methods outlined by Brank et al. [47] mentioned earlier.

3.5.1. Consistency

Internal consistency relates to the adherence of the ontological model to the rules of the description logic. It can be checked using reasoners in PROTÉGÉ and OOPS!.

3.5.2. Accuracy and Coverage (Completeness)

The evaluation of accuracy (measured in terms of precision and recall) and completeness usually relies on access to a gold standard built by domain experts. However, building this gold standard poses significant challenges, including the rapid evolution of domain knowledge, inconsistencies in annotation standards, the heterogeneity of biomedical sources, and the high cost of expert involvement [1].

To calculate precision and recall, two different approaches can be used. Precision is evaluated by domain experts that assess how well the ontology meets the set of predefined requirements. Laddering interview technique [57] and competency questions can be used. WebProtégé (<https://webprotege.stanford.edu/>, accessed on 27 May 2025). provides extensive collaboration features that allow sharing the ontology with domain experts. We recommend writing a document to explain the ontology, its objectives, and what is expected of the experts. This document should include competency questions designed using the laddering technique.

To calculate the recall, we suggest building an evaluation corpus of knowledge composed of documents that substantially cover a given domain and include CPGs, systematic reviews, and patient health records. The goal of using CPGs is to verify that all the concepts identified in phase 3 are covered by ontology, including relationships.

To calculate the completeness that reflects how well the ontology fully captures the domain it is intended to represent, we manually annotate a number of relevant text fragments from the corpus. Then we link all concepts that concern the ontology domain to concepts in the ontology. This step allows us to find the ratio of the covered concepts.

3.5.3. Clarity

Clarity concerns the meaning of the defined concepts and their independence of any social or computational context. It should be evaluated by domain experts. This step should be preceded by checking the readability of each concept, i.e., the existence of human-readable descriptions such as labels, definitions, or synonyms. This can be done by using tools such as OntoMetrics and OOPS!.

3.5.4. Conciseness

Unnecessary concepts with regards to the domain to be covered can be reported by domain experts when evaluating. Redundancy is also evaluated at this stage. It occurs when an axiom can be inferred from already defined axioms. Two types of redundancies are defined in [52]: (1) concepts (classes, properties, or instances) having more than one subsumption relation between them and (2) identical formal definitions of concepts with possibly different labels. The two kinds of redundancy can be detected by reasoners and OOPS! [58].

3.5.5. Computational Efficiency

The size of an ontology and the complexity of its axioms can slow down the processing of this ontology by tools. Computational efficiency measures how efficiently the tools used can process the ontology in particular reasoners. Unfortunately, ontologies are often unnecessarily complex due to a lack of respect for modular structure. Excessive response

times and excessive space requirements make these ontologies unusable for the deployment and maintenance. Since the metrics of processing time and required space are difficult to evaluate, we suggest simple verifications using tools like Protégé, OOPS!, or OntoMetrics.

3.5.6. Adaptability

Adaptability measures: (1) how well the ontology anticipates its use in different contexts; (2) whether it constitutes a reliable foundation for other ontologies; and (3) whether it is flexible enough to react predictably to small changes and to allow extensions without any need to remove axioms. Adaptability is strongly related to the ontology modular design. “An ontology-module is a re-employable segment of a bigger or more comprehensive ontology, which is self-sufficient but holds relationships with the other modules of the ontology that also encompasses the initial non-modularized ontology” [59]. A modular design offers numerous benefits, such as simple reuse of single components in other ontologies, feasible reasoning performance, facilitating ontology understanding, reduced complexity, making the central ontology less vulnerable to changes, and others. The quality of modular design can be assessed using semantic metrics such as coupling and cohesion. Coupling refers to the interdependencies that exist between ontology modules, i.e., to the number of shared symbols between axioms in different modules, while cohesion refers to the degree to which the elements in a module belong together. During the reuse process, the ontology engineer looks for an ontology with maximal cohesion and minimal coupling. There have been many attempts to define a strategy selecting “the best modular ontology” [60]. Unfortunately, currently, no tools support the creation of modular ontologies. A possible alternative can be a maximal reduction of the ontology scope; a similar approach has been applied to the development of the CORE subset of SNOMED CT.

Adaptability is also strongly related to the respect of the community standards. It includes usage of commonly accepted upper ontologies such as BFO, following OBO principles (<http://www.obofoundry.org/principles/fp-000-summary.html>, accessed on 27 May 2025), reusing relations from RO, specifying ontology version, and providing well-written documentation.

3.6. Phase 6: Ontology Documentation

An ontology must be documented for better understanding and use by future users on the one hand; on the other hand, its maintenance and update by any knowledge engineer will be easier. The documentation process starts in the first phase when the domain and scope of the ontology are defined. At the end of each phase, a document detailing its progress and deliverables is established. The documentation of the ontology includes its classes, relations, properties, instances, axioms, and annotations. OWLDoc (<https://protegewiki.stanford.edu/wiki/OWLDoc>, accessed on 27 May 2025) and Live OWL Documentation Environment (LODE: <https://essepuntato.it/lode/>, accessed on 27 May 2025) are tools that can be used to generate readable documentation presented to the user in the form of an HTML page with embedded links for easy navigation.

3.7. Phase 7: Maintenance and Evolution

Medical knowledge is evolving, and this evolution has to be reflected during the maintenance process. It includes integrating the missing knowledge and maintaining ontology consistency and coherence.

If there are changes that occur in the ontology domain and scope definition (phase 1), all the phases of the ontology building should be carried out.

If the updates concern only a few concepts in the ontology that need to be added, then follow the steps of adding concepts and resolving conflicts approach as presented in Figure 2.

If the updates concern reused concepts that were modified in the reused ontologies, then reflect the modification in the concept if relevant. If not, keep the old version.

The changes must be consistent with the domain and scope of the ontology. Once the changes are applied to the ontology, the ontology must be evaluated for consistency.

Reused ontologies are usually updated by their authors. Currently, ontology updates must be tracked manually, as no tools exist to automate this process.

4. Results

We present in this section the results of applying the proposed methodology to building pneumonia diagnosis ontology (PNADO).

4.1. Ontology Domain and Scope Definition

We defined the ontology domain and scope after meetings with physicians from the Gatineau Hospital, who were interested in reducing erroneous pneumonia diagnoses. Functional requirements are defined by a set of CQs established by physicians and come from the CPGs covering pneumonia diagnosis. The following questions, among others, have to be answered by physicians to diagnose pneumonia correctly:

1. What are the symptoms and clinical signs of pneumonia?
2. What are the types of pneumonia?
3. How is pneumonia diagnosed?
4. What are the pathogens of pneumonia?
5. What is the clinical history of the patient?
6. What are the laboratory tests to diagnose pneumonia?
7. What are the results of the physical examination of the patient?
8. What is the result of lung imaging of the patient?
9. What are the results of the lab tests of the patient?
10. Once the diagnosis of pneumonia has been made, what complications should the physician look for?

4.2. Corpus Building and Term Extraction

The corpus of knowledge is constituted of 13 publicly available CPGs from national and international repositories that cover most of the pneumonia diagnosis knowledge. They are Cochrane (<https://www.cochranelibrary.com/>, accessed on 27 May 2025), NICE (<https://www.nice.org.uk/>, last access on 27 May 2025), Infectious Diseases Society of America (IDSA: <https://www.idsociety.org/>, accessed on 27 May 2025), European Society of Clinical Microbiology and Infectious Diseases (ESCMID: <https://www.escmid.org/>, accessed on 27 May 2025), Australian Society for Infectious Diseases (ASID: <https://www.asid.net.au/>, accessed on 27 May 2025), Canadian Respiratory Guidelines (CTS: <https://cts-sct.ca/>, accessed on 27 May 2025), Pulmonary & Critical Care Medicine (PulmCCM: <https://pulmccm.org/>, accessed on 27 May 2025), American Thoracic Society (ATS: <https://www.thoracic.org/>, accessed on 27 May 2025), and British Thoracic Society (BTS: <https://www.brit-thoracic.org.uk/>, accessed on 27 May 2025). The language of all the CPGs is English.

To build our ontology, we select the terms that are relevant to pneumonia diagnosis. These terms include symptoms, clinical signs, imaging, laboratory tests and their results, pathogens (agents and others), antecedents, types of pneumonia, differential diagnosis, and complications. Then, in the set of found terms, we examine adjectives, adverbs, nouns and proper nouns, verbs, and phrases because we believe that they convey the knowledge necessary to identify concepts in the pneumonia diagnosis domain.

4.3. Building of Preliminary Ontology

4.3.1. Semantic Normalization

To develop a differential ontology, we have to articulate the candidate terms selected previously by specifying the differential principles that characterize them. For example, bacterial pneumonia, fungal pneumonia, infective pneumonia acquired prenatally, pneumonia due to parasitic infestation, and viral pneumonia are sibling concepts because they all represent types of infectious diseases. This step is the most time-consuming.

4.3.2. Knowledge Formalization

The following example illustrates adding an axiom. In the following text, ontological concepts are written in italic. *Sign* (also called *clinical sign*) is a concept in PNADO. It is defined as the “quality of a patient, a material entity that is part of a patient, or a processual entity that a patient participates in, any one of which is observed in a physical examination and is deemed by the clinician to be of clinical significance”. *Symptom* is another concept in PNADO; it is defined as “quality of a patient that is observed by the patient or a processual entity experienced by the patient, either of which is hypothesized by the patient to be a realization of a disease”. We add an axiom to indicate that the two concepts are disjoint.

4.3.3. Toward a Computational Ontology

During that process, using Protégé, we operationalize the ontology in the OWL language because it meets our needs in terms of expressiveness and manageability. We use the upper medical ontology OGMS. The choice of OGMS will be explained in the next phase. For further enrichment of PNADO, we extract definitions, synonyms, acronyms, and other annotations from UMLS by using MetaMap. At the end of this step, we obtain the preliminary PNADO.

4.4. Ontology Reuse and Enrichment

4.4.1. Finding Ontologies

We focus on two open content repositories of biomedical ontologies: Open Biomedical Ontology (OBO) Foundry and BioPortal. The English language and the OWL format are required. Using the search criteria, we found 26 candidate ontologies that could be reused.

4.4.2. Choosing Relevant Ontologies

We identify ontologies to be reused by using accuracy and precision orderings. No precise ontology has been found. Choosing between Basic Formal Ontology (BFO: <http://www.obofoundry.org/ontology/bfo.html>, accessed on 27 May 2025) and General Formal Ontology (GFO: <https://biportal.bioontology.org/ontologies/GFO>, accessed on 27 May 2025): both are designed to be upper ontologies. BFO has fewer superfluous and fewer omitted models than GFO; then it is more accurate. We reuse BFO as an upper ontology (hard reuse).

Choosing between Ontology of General Medical Science (OGMS: <https://biportal.bioontology.org/ontologies/OGMS>, accessed on 27 May 2025) and DIAGONT (<https://biportal.bioontology.org/ontologies/DIAGONT>, accessed on 27 May 2025): these ontologies cover the most important concepts for diagnostic. Both have superfluous and omitted models that are incomparable. Since OGMS is built upon BFO, we choose to reuse it (hard reuse). Hence, PNADO follows OGMS representation.

Choosing between SYMP (<http://www.obofoundry.org/ontology/symp.html>, accessed on 27 May 2025) and Clinical Signs and Symptoms Ontology (CSSO: <https://biportal.bioontology.org/ontologies/CSSO>, accessed on 27 May 2025): SYMP has fewer superfluous and fewer omitted models than CSSO; it is more accurate. We reuse SYMP, which covers mostly symptoms and clinical signs.

Choosing between SNOMED-CT (<https://browser.ihtsdotools.org/?>, accessed on 27 May 2025) and NCBITaxon (<http://www.obofoundry.org/ontology/ncbitaxon.html>, accessed on 27 May 2025): SNOMED-CT is the most widely used healthcare terminology and has the most comprehensive coverage of concepts for representing clinical knowledge [61,62]. SNOMED-CT is used as an ontology rather than terminology, it is based on DL, and it can be reasoned over [63]. NCBITaxon is a classification and nomenclature of all the organisms in the public sequence databases. Both have superfluous and omitted models that are not comparable. Both are interesting to make soft reuse. SNOMED-CT covers other clinical aspects in addition to organisms that also can be reused.

Choosing between GAMUTS (<https://bioportal.bioontology.org/ontologies/GAMUTS>, accessed on 27 May 2025) and Radiological Lexicon (RadLex: <http://radlex.org/>, accessed on 27 May 2025): both are comprehensive terminologies for radiology, and both have superfluous and omitted models, but Radlex has fewer. We have decided to reuse Radlex (soft reuse). Choosing between Logical Observation Identifiers Names and Codes (LOINC: <https://loinc.org/>, accessed on 27 May 2025) for medical laboratory observations identification and clinical Laboratory Ontology (LABO: <https://bioportal.bioontology.org/ontologies/LABO>, accessed on 27 May 2025) both have superfluous and omitted models that are not comparable. We choose to reuse LOINC because it is the most used by clinicians.

Choosing between the International Classification of Diseases (ICD10: <https://bioportal.bioontology.org/ontologies/ICD10>, accessed on 27 May 2025) and Human Disease Ontology (DOID: <http://www.obofoundry.org/ontology/doid.html>, accessed on 27 May 2025), which is a representation of human diseases organized by etiology: both have superfluous and omitted models that are not comparable. DOID provides more concepts related to PNADO than ICD10. We reuse DOID (soft reuse).

We also reuse the following four ontologies dealing with important topics that are not covered by other ontologies chosen in the previous step: Relation Ontology (RO: <http://www.obofoundry.org/ontology/ro.html>, accessed on 27 May 2025) containing relations between entities; Computer-Based Patient Record Ontology (CPRO: <https://bioportal.bioontology.org/ontologies/CPRO>, accessed on 27 May 2025) for patient profile description; Human Phenotype Ontology (HPO: <http://www.obofoundry.org/ontology/hp.html>, accessed on 27 May 2025) providing a structured and controlled vocabulary for the phenotypic features in human hereditary and other diseases; Infectious Disease Ontology (IDO: <http://www.obofoundry.org/ontology/ido.html>, accessed on 27 May 2025) covers the infectious disease domain.

The reused ontologies use different naming conventions. In SNOMED-CT, the name of each concept begins with a capital letter, whereas in SYMP, it begins with a lowercase letter. In PNADO, names begin with a capital letter. In this paper, each concept is written according to the convention used by the ontology in question.

4.4.3. Resolving Conflicts

Here, we illustrate resolving concept conflicts with a few examples.

The concept of *hypoxemia* is found in HPO, SYMP, and SNOMED-CT. To choose the best representation of this concept, we proceed as follows:

- CRQ 1: “Does *hypoxemia* have a definition in each ontology?” Answer: Only HPO provides a definition.
- CRQ 2: “Does its definition correspond to the PNADO requirements?” Answer: The definition given in HPO is relevant to PNADO.
- CRQ 3: “Does *hypoxemia* have children?” Answer: *Hypoxemia* has children in HPO and SNOMED-CT but none in SYMP.
- CRQ 4: “Is each found child relevant to the PNADO requirements?” Answer: *Hypoxemia*’s children found in HPO are relevant to the PNADO requirements.

- Analyze: *Hypoxemia* presented in HPO seems to have the best representation for PNADO since it provides a relevant definition and relevant children that correspond to the PNADO requirements.

The concept of *patient* is represented differently in CPRO and SNOMED-CT. We use the following CRQs to choose the most relevant representation for this concept:

- CRQ 1: “Does *patient* have a definition in each ontology?” Answer: No ontology provides a definition.
- CRQ 2: “What is the parent concept of *patient*?” Answer: In CPRO, the parent concept is *organism* that is provided by OGMS. In SNOMED-CT, the parent is *social context* that does not fit the PNADO requirements.
- CRQ 3: “Is any person a patient?” Answer in CPRO is the axiom “*Human/Person* and (*Plays Role* **some** *Patient* role) and (*Participates_in* **some** *Clinical* act).” No answer is found in SNOMED-CT.
- Analyze: CPRO provides a better representation for *patient* than SNOMED-CT. It specifies with an axiom in which case a *person* can be a *patient*, and it is a subclass of *organism* defined by OGMS.

Two other examples of resolving conflicts are presented in Table 2.

Table 2. Examples of inconsistencies resolved with CRQs.

Concept	Ontologies	CRQs	CRQs Responses
<i>Hypotension</i>	HPO SNOMED-CT SYM	Is <i>hypotension</i> a symptom or a vital sign?	- HPO: <i>Hypotension</i> is a phenotypic abnormality. - SNOMED-CT: <i>Hypotension</i> (low blood pressure) is considered a disorder of the cardiovascular system. - SYMP: <i>Hypotension</i> is considered a hemic system symptom.
		What is the synonym of <i>hypotension</i> ?	- <i>Hypotension</i> has two synonyms in HPO, three synonyms in SNOMED-CT, and no synonym in SYMP.
		Are there any <i>hypotension</i> ’s children?	- <i>Hypotension</i> has three children in HPO, twelve children in SNOMED-CT, and one child in SYMP.
		Is each child relevant to the diagnosis of pneumonia?	- Only <i>hypotension</i> ’s children provided by HPO and SYMP are relevant for the diagnosis of pneumonia.
		Analyze: <i>Blood pressure</i> is a measurement of the pressure in arteries. We consider it in PNADO as a <i>vital sign</i> . <i>Hypotension</i> ’s children in HPO combined with SYMP’s child correspond to PNADO intended model. We reuse <i>hypotension</i> ’s synonyms of SNOMED-CT.	
<i>Respiratory rates</i>	SNOMED-CT LOINC	Is <i>respiratory rate</i> a symptom or a vital sign?	- In both SNOMED-CT and LOINC, it is mentioned that the scale type of <i>respiratory rate</i> is quantitative.
		Are there any <i>respiratory rate</i> ’s children?	- SNOMED-CT: The <i>respiratory rate</i> has two children. - LOINC: The <i>respiratory rate</i> has no children.
		Is each child relevant to the diagnosis of pneumonia?	- Yes, the <i>respiratory rate</i> ’s children provided by SNOMED-CT are relevant for the diagnosis of pneumonia.
		Is there any measurement unit?	- No measurement for <i>respiratory rate</i> is given in SNOMED-CT. - LOINC provides “[4]/min” as an example of the <i>respiratory rate</i> measurement unit.
		Are there any synonyms? Are there any related names?	- SNOMED-CT provides two synonyms and no related names. - LOINC doesn’t provide any synonym for <i>respiratory rate</i> but provides 13 related names.
		Analyze: The <i>respiratory rate</i> is represented as a <i>vital sign</i> in PNADO. We reuse SNOMED-CT’s children because they respond to the PNADO scope and its intended model. We reuse the unit measurement of LOINC and some of the related names. We also reuse the synonyms of SNOMED-CT.	

4.5. Ontology Evaluation

For the evaluation of PNADO, we used tools, a clinical dataset that contains multiple cases of pneumonia, and a clinical corpus of knowledge captured in the CPGs and systematic reviews. Domain experts were also involved in the process.

4.5.1. Consistency

We used the Pellet reasoner integrated with PROTÉGÉ and the tool OOPS! to check the internal consistency of PNADO. OOPS! reported minor pitfalls regarding annotations. The missing license for PNADO was reported as an important pitfall. No critical pitfalls were reported.

4.5.2. Accuracy and Coverage (Completeness)

To calculate recall, we constructed an evaluation knowledge corpus constituted of (1) the 13 CPGs chosen in Section 4.2; (2) 43 systematic reviews that treat pneumonia diagnosis from national and international repositories containing evidence-based medical documents; and (3) the clinical dataset Multi-Parameter Intelligent Monitoring in Intensive Care III (MIMIC-III: <http://physionet.org/>, accessed on 27 May 2025) [64].

First, we treated the CPGs and systematic reviews. We manually extracted 710 different terms related to pneumonia diagnosis from 336 text fragments (each fragment had 20 lines on average) and tried to link them to concepts in PNADO. Linking extracted concepts to concepts in PNADO, we found a match for 570 concepts giving a value of recall of 80%. The remaining 140 concepts that could not be matched were added to PNADO.

Here is an example (see Table 3) of annotation of a fragment of text from the clinical guideline “Management of Community-Acquired Pneumonia in Adults” [65].

Table 3. Example of annotation of clinical guideline paragraph with PNADO.

Concepts from the Text	Equivalences in the PNADO			
	Class	Object Property	ID	Parent Entity in PNADO
Chest radiograph	<i>Chest radiography</i>		PNADO:0000783	<i>Imaging of lung</i>
Patients likely to have pneumonia		<i>Has a diagnosis (patients, pneumonia)</i>	PNADO:0001291	
Diagnosis	<i>Diagnosis</i>		OGMS:0000073	<i>Data item</i>
Differentiating CAP from ... acute bronchitis		<i>Differential diagnosis of (CAP, Acute bronchitis)</i>	PNADO:0001157	
CAP	<i>Community-acquired pneumonia</i>		SNOMEDCT_US:385093006	<i>Pneumonia</i>
Cough	<i>Cough</i>		SNOMEDCT_US:49727002	<i>Respiratory system and chest symptoms</i>
Fever	<i>Fever</i>		SYMP:0000613	<i>Neurological and physiological symptoms</i>

“A *chest radiograph* is required for the routine evaluation of *patients* who are *likely to have pneumonia*, to establish the *diagnosis* and to aid in *differentiating CAP* from other *common causes* of *cough* and *fever*, such as *acute bronchitis*.”

Second, we treated records from MIMIC III that consists of 38,597 distinct and de-identified adult patients and 49,785 hospital admissions. It includes, among others, laboratory data, radiology reports, vital signs, therapeutic intervention profiles, ventilator settings, nursing notes, diagnostic codes, discharge summaries, and provider order entry data. It comprises 26 tables linked by identifiers such as HADM_ID referring to unique hospital admission and SUBJECT_ID referring to a unique patient. Among those tables, we considered the “diagnoses_icd” table that contains ICD-9 diagnoses for patients and

the “note events” table that contains all notes about patients from their admissions to their discharges, including nursing and physician notes, ECG reports, radiology reports, and discharge summaries.

We found 7702 cases of pneumonia diagnosis in the “diagnoses_icd” table. As shown in Figure 3, there were 32 types of pneumonia. We were interested in making sure that PNADO covered all the cases captured in MIMIC-III. The number of pneumonia cases with code 860 was more than half of all cases. Consequently, we used the logarithm function to reduce the number of cases to include in the evaluation for each type of pneumonia, 43 cases in total.




Figure 3. MIMIC-III pneumonia cases.

The next task was the extraction of the records from the “noteevents” table by using HADM_ID already obtained in the previous task. We randomly selected cases according to the cardinality of each type and completed the annotation. We manually annotated terms related to pneumonia diagnosis and verified whether they were covered by PNADO or not. For those not covered, we analyzed if they were synonyms for the existing concepts or new concepts.

At the end of the evaluation of PNADO with the MIMIC III database, 36 concepts of 988 concepts found in the electronic health records were not covered by PNADO. The analysis of each concept revealed that 16 concepts were synonyms, and 20 concepts were subclasses of the existing concepts. Here are some examples of synonyms: pulmonary embolus, rhonchi sound, rhonchus sound, pneumococci, pneumonic infiltrates, MSSA PNA, MSSA pneumonia, MRSA PNA, MRSA pneumonia, and trouble breathing. The new subclasses include types of pneumonia such as *polymicrobial pneumonia*; physical examination findings such as *respiratory failure*; symptoms such as *recurrent cough* and *occasional cough*; and clinical histories such as *tracheostomy*. The recall ratio was 96%.

To calculate the precision of PNADO, we involved a pneumologist from Charles-Le Moyne Hospital (<https://santemonteregie.qc.ca/installations/hopital-charles-le-moyne>, accessed on 27 May 2025), an emergency physician from the Gatineau Hospital (<https://cisss-outaouais.gouv.qc.ca/>, accessed on 27 May 2025), and a family doctor from the University of Quebec in Outaouais medical clinic (<https://ssuqo.ca/>, accessed on 27 May 2025). We prepared a document to explain the ontology and its objective. We also prepared a document containing the CQs used for the determination of the domain and scope of PNADO (Section 4.1 Ontology Domain and Scope Definition) and questions using the laddering technique [57]. Here are a few examples of these questions: What are the symptoms of pneumonia? Is the classification of symptoms according to the different systems (digestive system, cardiovascular system, etc.) correct? Is each symptom well classified?

The participating physicians created accounts in WebProtégé, and we shared PNADO with them. Once they had familiarized themselves with the environment, they evaluated each concept and commented directly in WebProtégé. They verified the relevance of each concept in PNADO to pneumonia diagnosis, its position, and its relations with the other concepts. For each concept, they provided a mention of evaluation (relevance and suggestions) in WebProtégé. They suggested several changes including adding radiological signs of consolidation: *air bronchograms*, *ill-defined*, *fluffy opacities*, *air alveologram*, *patchy opacities*, *acinar*, *preserved lung volume*, *extension to pleural surface*, *CT angiogram sign*, *silhouette sign*; moving concepts of *fever* to *general symptom*, *purulent tracheobronchial secretions* to *respiratory system and chest symptom*, *headache* to *neurological and physical symptom*, *arterial blood gas* to *procedure*; and removing the concepts *transitory tachypnea of the newborn*, *sputum eosinophilia*, and *lithoptysis*. The precision ratio was 96%.

4.5.3. Clarity

We used OntoMetrics and OOPS! tools to evaluate readability. OntoMetrics was not able to check the readability of all the concepts. OOPS! detected that some of the new concepts added to PNADO did not have definitions. In fact, these concepts have no definitions in UMLS. A total of 663 concepts without definition (mostly UMLS concepts) were reported.

4.5.4. Conciseness

OOPS! and Pellet reasoner were used to detect redundancies in PNADO. OOPS! reported two classes that contain the same labels. We found that the two classes have labels that only differ by one character (*legionella pneumophila serogroup 1* and *legionella pneumophila serogroup 2*).

4.5.5. Computational Efficiency

Time processing of PNADO using Pellet reasoner, OOPS!, and OntoMetrics was short (1–2 s) with one exception: OntoMetrics was not able to check the criterion of readability when providing the entire ontology.

4.5.6. Adaptability

PNADO uses the BFO upper ontology that enhances reusability and interoperability. It also follows OBO Foundry principles and reuses relations from RO. We specify the PNADO version and provide documentation (see Section 4.5.6 Adaptability). PNADO does not have a modular structure.

4.5.7. Characteristics of PNADO

PNADO is published in BioPortal repository and can be viewed and explored at <https://bioportal.bioontology.org/ontologies/PNADO>, accessed on 23 May 2025 (see the Supplementary Materials).

We present in this section the results obtained by applying the proposed methodology to the PNADO building.

PNADO contains 1598 classes (1448 are reused and 150 are new classes), 42 object properties (27 are reused and 15 are new), 1591 logical axioms, and 83 annotation properties. Table 4 presents statistics on reused concepts and resolved conflicts. We noticed that SNOMED-CT was frequently involved in conflict resolution, as it overlaps with most existing medical ontologies. PNADO is built according to the principles of OBO Foundry, using OGMS that follows the BFO paradigm. OGMS provides a set of general reference classes related to diseases and diagnoses. PNADO is only focused on diagnosis. Therefore, only the OGMS concepts related to diagnosis are used. Other high-level OGMS terms may be used in future extensions.

Table 4. Ontology reuse in PNADO.

	Ontology	Usage in PNADO	Classes	Relations	Total	Conflicts
Hard reuse						
1	BFO	Upper ontology	35	5	40	0
2	OGMS	Upper domain ontology	149	0	149	0
Soft reuse						
1	SYMP	Symptoms	57	0	57	22
2	NCBITAXON	Virus and bacteria	184	0	184	0
3	RO	Relations	0	22	22	0
4	CPRO	Roles	4	0	4	2
5	LOINC	Laboratory	26	0	26	1
6	SNOMED-CT	Diseases, symptoms, and clinical signs	796	0	796	115
7	RADLEX	Imaging	55	0	55	5
8	DOID	Disease	35	0	35	13
9	HPO	Phenotype	96	0	96	73
10	IDO	Pathogens	11	0	11	0

Figure 4 shows an example from PNADO demonstrating the hierarchy of the class *infective pneumonia* and, in particular, the class *viral pneumonia*. Figure 5 shows the representation of the *healthcare process assay* class.

4.6. Ontology Documentation

We annotated every concept in PNADO and used OWLDoc to document the ontology. Every phase of the building process is described in this work.

4.7. Maintenance and Evolution

Within the COVID pandemic context, two CPGs treating pneumonia diagnosis were found in NICE and the Centre for Evidence-Based Medicine (CEBM: <https://www.cebm.net/>, accessed on 27 May 2025) repositories. We extracted 16 new concepts relevant to PNADO. We also frequently check the parts of the reused ontologies. Recently, OGMS has

been subject to changes that affect PNADO significantly. For example, *symptom* was moved to *process* class, *vital sign* was moved to *material entity* class, *sign* class was deleted, etc.

Currently, we manage ontology versioning through manual tracking, which involves carefully documenting changes and updates to the ontology. We ensure that these changes are consistent and aligned with the requirements of the system. However, when reused ontologies evolve independently, it requires close monitoring to detect changes or updates in those ontologies.




Figure 4. The disease class with infective pneumonia as an example of the subclass.




Figure 5. Representation of *healthcare process assay* class in Protégé.

5. Discussion

5.1. Generality of the Methodology

The techniques proposed in each phase are domain-independent and can be applied across various fields. However, the corpus of knowledge must be customized to align with the specific domain of the targeted ontology. Similarly, any ontologies to be reused must be relevant to the same field. Using codified knowledge is particularly beneficial for knowledge engineers, as it provides a clear understanding of the domain and minimizes the risk of confusion.

5.2. Ontology Reuse

Ontology reuse in the methodology is used to enrich the preliminary ontology with new concepts and annotations and to enhance ontology interoperability and reusability. Since the methodology is designed for a knowledge engineer and the involvement of the domain experts is limited, ontology reuse cannot be performed at the beginning of the building process because the knowledge engineer does not know which concepts are relevant to the ontology domain.

Ontology reuse is the trickiest phase in the methodology due to the interoperability issues among the reused ontologies. Performing this phase during the PNADO building revealed some findings:

- Choosing relevant ontologies to reuse is very important and determinant for the rest of the reuse process. Some candidate ontologies can be easily excluded while others require careful analysis.
- Even if the chosen ontology is of good quality, the representation of a concept may not be relevant to the requirements of the ontology that is being built. We also found out during the PNADO building that some ontologies have bad representations for some concepts. For example, SYMP represents diseases as symptoms. An example of a relation found in RO that is not relevant for PNADO is discussed in Section 5.3.
- Each ontology is chosen according to a significant topic, as discussed in Section 3.4.2, but in the process of reuse, the chosen ontology may contain concepts that deal with other topics. For example, DOID was chosen for the topic of diseases, but it was involved in conflicts regarding concepts related to symptoms and clinical signs.
- A concept can be differently presented in different ontologies, and choosing the best representation to reuse requires reflection about the relevant CRQs.

5.3. Upper Domain Ontologies

As mentioned in Section 5.2, there are ontologies designed to cover a specific domain that also cover concepts belonging to other domains. These ontologies do not respect the principles of modular design [9]. The unwanted concepts should be usually covered by upper domain ontologies. It is observed that it is often useful to design several levels of upper domain ontologies to get a better module coupling and cohesion and, at the same time, to reduce the reuse and interoperability issues.

OGMS, which covers diagnosis and treatment of disease, was used as an upper domain ontology for PNADO. Many concepts (classes and relations) required for the diagnostic process are missing in OGMS. PNADO aims to cover the concepts related to pneumonia diagnosis, and the concepts related to the diagnostic process should be available in an upper domain ontology.

A different example would be an upper domain ontology describing patient. Currently, domain ontologies differ in their modeling of the patient concept. Having one upper ontology representing patient would be very useful for reuse and enhancing interoperability.

5.4. PNADO New Concepts

In the following paragraphs, we discuss some of the reused object properties and classes that required major modifications.

Complication of

A *complication* in medicine is a medical problem that occurs during or after a disease, procedure, treatment, or function (pregnancy, for example). This concept is represented as a class in many ontologies, such as SNOMED-CT or Radlex. In SNOMED-CT, *complication* is a subclass of *disease* and represents complications of disease, such as *Complications due to diabetes mellitus* or *Complication due to Crohn's disease*. *Diabetes mellitus* and *Crohn's disease* are subclasses of *disease*. The diseases concerned by complications are also represented as subclasses of the class *disease*. Such a representation creates an inheritance chain *disease* → *complication* → *disease* and does not cover all cases of complications due to *disease*, *procedure*, *treatment*, or *function*. We propose to represent *complication* as an object property named *complication of* with four sub-properties: (1) *complication of disease* that associates *disease* to *disease*; (2) *complication of procedure* that associates *procedure* to *disease*; (3) *complication of treatment* that associates *treatment* to *disease*, and (4) *complication of function* that associates *function* to *disease*. Axioms are progressively added when needed to represent complications. For example, in PNADO, axioms are added to illustrate pneumonia complications. This way of representing *complication of* enhances interoperability and reusability.

Differential Diagnosis of

Differential diagnosis is the distinguishing of disease from others that present similar clinical features (signs and symptoms). It is represented in SNOMED-CT and LOINC as a class. It is observed that there is no representation of pneumonia differential diagnosis in both ontologies. The presence of differential diagnosis could help to avoid diagnostic errors. We propose a new object property *Differential diagnosis of* that holds between a *diagnosis* "A" and a *diagnosis* "B" if they have some similar clinical features. Recall that a *diagnosis* is the representation of a conclusion of a diagnostic process that can be a disease or a syndrome. Axioms are added to define the relation. An example of such an axiom: *pneumonia* "Differential diagnosis of" **some** *bronchitis*.

Has Symptom

Has symptom is an object property of RO where it is defined as "a relation that holds between a disease or an organism and a phenotype." We notice that it is a sub-property of the object property that *has phenotype*. The latter associates the domains *generically dependent continuant* or *material anatomical entity* or *disease* to the range *phenotype*. The class *phenotype*, which is reused by RO from Combined Phenotype Ontology (UPHENO), is defined "as a defect or loss of some anatomical structure or a biological process to wild-type." The class *phenotype* in OGMS and reused in PNADO is defined as "a (combination of) quality(ies) of an organism determined by the interaction of its genetic make-up and environment that differentiates specific instances of a species from other instances of the same species." It is obvious that the two definitions are different; the one provided by UPHENO does not correspond to PNADO intended model, and subsequently, we do not reuse an object property associating this class with any other class. We create a new object property *Has a symptom* that associates *patient* to *symptom*.

Symptom

Symptom is defined in OGMS as "a quality of a patient that is observed by the patient or a processual entity experienced by the patient, either of which is hypothesized by the patient to be a realization of a disease." It is also defined in SYMP as "a perceived change in function, sensation, loss, disturbance or appearance reported by a patient indicative

of a disease.” Nevertheless, during the reuse of SYMP, we noticed that it also contains diseases and clinical signs when it was intended to contain only symptoms. For instance, *hypotension*, *arrhythmia*, *bronchitis*, *endocarditis*, and *shock* are defined in SYMP as symptoms, whereas they are diseases. Indeed, there is often confusion between clinical signs, diseases, and symptoms in SYMP and other ontologies. In PNADO, the class *symptom*, thanks to the ontological commitment, includes only what is reported by a patient.

Pathogen

Before discussing *pathogen*, we will briefly review the *organism* since both concepts are connected. *Organism* is an interesting concept needed in PNADO, and it can be found in other ontologies such as CPRO, IDO, or SNOMED-CT. In SNOMED-CT, it is represented as an independent entity that includes *bacteria*, *virus*, *archaea*, *eukarya*, and *prion*. Since no definition is given to this concept, its meaning (not precise) can only be deduced from its sub-classes. CPRO represents *organism* as a subclass of *object*, whereas IDO represents it as a sibling to *object*. An *object* according to the definition given by BFO “is a material entity that is: (1) spatially extended in three dimensions; (2) causally unified, meaning its parts are tied together by relations of connection in such a way that if one part of the object is moved in space, then its other parts will likely be moved also; and (3) maximally self-connected.” The definitions given by the three ontologies are almost similar, and they mean that an organism is an individual living system that may be unicellular or made up of many billions of cells like humans. This leads us to conclude that an *organism* is definitely an *object*, contrary to the IDO representation. *Pathogen* in CPRO is a subclass of *organism*, and it is defined as “any virus, microorganism, or other substance causing disease.” This definition is limited because, in the broadest sense, *pathogen* is anything that can produce disease. SNOMED-CT provides *pathogenic organism* rather than *organism*, as a subclass of *navigational concept*, separate branch in the class hierarchy. IDO provides a more accurate definition and representation of *pathogen*. *Pathogen* is a subclass of *material entity* with a *pathogenic disposition*. The latter is a subclass of *disposition* defined in IDO as “a disposition to initiate processes that result in a disorder.” It means that some organisms in some circumstances of their lives play the role of pathogens toward other organisms and cause diseases. We create *pathogen* as a subclass of *material entity* and *pathogen role* as a subclass of *role*, which is already defined in BFO. We add the axiom *pathogen* is a *material entity* **and** (“has disposition” **some** *pathogenic disposition*) **and** (“has role” **some** *pathogen role*). *Has disposition* is an object property defined in RO as “a relation between an independent continuant (the bearer) and a disposition, in which the disposition specifically depends on the bearer for its existence.” *Has role* is also an object property that is defined in RO as “a relation between an independent continuant (the bearer) and a role, in which the role specifically depends on the bearer for its existence.”

6. Conclusions

The main objective of this work was to demonstrate how to build a high-quality medical ontology. Involving domain experts throughout the development process is crucial to achieving the desired quality; however, assembling such a team is often challenging. Providing the knowledge engineer with clear construction guidelines can help compensate for the limited availability of domain experts during certain phases. This approach allows their involvement to be focused primarily on the domain and scope definition, as well as the final evaluation phase.

We propose a methodology inspired by METHONTOLOGY [19], with the following key enhancements: (1) During the corpus building and term extraction phase, we recommend constructing a knowledge corpus using the most reliable, codified sources; (2) In the preliminary ontology building phase, we suggest applying differential principles to define

each ontology concept, whereby the knowledge engineer articulates the relationships by expressing, in natural language, the similarities and differences between a concept and its neighboring concepts in the hierarchy; (3) In the ontology reuse and enrichment phase, we recommend that the knowledge engineer employ conflict resolution questions to address inconsistencies between concept representations; (4) Finally, the methodology provides explicit guidance for documentation (annotations and conceptual modeling), maintenance and evolution (versioning and dependency management), and evaluation.

The methodology integrates ontology engineering from scratch using codified knowledge with ontology reuse to enhance interoperability and reusability. It is adaptable to other medical domains; for example, gold standard frameworks could be utilized as sources of codified knowledge to represent a prognosis domain.

The proposed methodology can be effectively applied to the development of ontologies in various domains beyond healthcare. In such cases, it is essential to identify a reliable source of codified knowledge. Examples of such sources include best practice guidelines for pollution prevention and waste minimization, industry billing standards, and various industrial norms.

We validated our methodology by applying it to pneumonia diagnosis through the development of PNADO. PNADO was evaluated by physicians, and the results demonstrated the effectiveness of our approach. To the best of our knowledge, PNADO is the first reported ontology specifically designed to represent the various aspects of pneumonia diagnosis.

During the PNADO building, in the ontology reuse phase, we suggested some improvements to the OGMS, SYMP, and CPRO ontologies.

In future work, we will consider the modularity aspect in the building process of medical ontology.

Supplementary Materials: The following supporting information can be found at: <https://bioportal.bioontology.org/ontologies/PNADO>, accessed on 23 May 2025.

Funding: The author declares that she has not received project funding for this work.

Data Availability Statement: The ontology developed and evaluated in this study is publicly available on the NCBO BioPortal: Pneumonia Diagnosis Ontology | NCBO BioPortal.

Acknowledgments: We are indebted to the Centre Intégré de la Santé et des Services Sociaux de l'Outaouais (CISSSO) and specifically Sylvain Croteau, emergency physician at Gatineau, Qc. Hospital; Yasmine Lisa Rebaine, a pneumonologist in Charles-Le Moyne Hospital; and Serge Chartrand, a family doctor at UQO medical clinic. Last but not least, we would like to express our gratitude to Wojtek Michalowski, at Ottawa University, for his support. Last but not least, my heartfelt gratitude to Michal Iglwieski—may his soul rest in peace—for his invaluable support in the completion of this work.

Conflicts of Interest: The author declares that there is no conflicts of interest statements.

Abbreviations

BFO: Basic Formal Ontology; CIDO: Coronavirus Infectious Disease; CPRO: Computer-Based Patient Record Ontology; CPGs: clinical practice guidelines; CRQs: conflict resolution questions; CSSO: Clinical Signs and Symptoms Ontology; CUI: concept unique identifiers; DL: description logic; DOID: Human Disease Ontology; GFO: General Formal Ontology; HPO: Human Phenotype Ontology; ICD10: International Classification of Diseases; LOINC: Logical Observation Identifiers Names and Code; NCBO: National Center for Biomedical Ontology; OGMS: Ontology for General Medical Science; OM: omitted models; OWL: Web Ontology Language; PNADO: Pneumonia Diagnostic Ontology; RadLex: Radiology

Lexicon; RO: Relation Ontology; RTF: relative term frequency; SNOMED-CT: Systematized Nomenclature of Medicine Clinical Terms; SUP: superfluous models; TFIDF: term frequency inverted document frequency; UMLS: Unified Medical Language System.

References

- Patel, A.; Debnath, N.C. A Comprehensive Overview of Ontology: Fundamental and Research Directions. *Curr. Mater. Sci. Former. Recent. Pat. Mater. Sci.* **2024**, *17*, 2–20. [\[CrossRef\]](#)
- Amith, M.; He, Z.; Bian, J.; Lossio-Ventura, J.A.; Tao, C. Assessing the practice of biomedical ontology evaluation: Gaps and opportunities. *J. Biomed. Inf.* **2018**, *80*, 1–13. [\[CrossRef\]](#) [\[PubMed\]](#)
- Amith, M.; Manion, F.; Liang, C.; Harris, M.; Wang, D.; He, Y.; Tao, C. Architecture and usability of OntoKeeper, an ontology evaluation tool. *BMC Med. Inform. Decis. Mak.* **2019**, *19*, 152. [\[CrossRef\]](#)
- Fung, K.W.; Bodenreider, O. Knowledge Representation and Ontologies. In *Clinical Research Informatics*; Springer: Berlin/Heidelberg, Germany, 2023; pp. 367–388.
- Katsumi, M.; Grüninger, M. *What is Ontology Reuse?* FOIS: Kalyan, India, 2016.
- Pathak, J.; Johnson, T.M.; Chute, C.G. Survey of modular ontology techniques and their applications in the biomedical domain. *Integr. Comput. Aided Eng.* **2009**, *16*, 225–242. [\[CrossRef\]](#) [\[PubMed\]](#)
- Simperl, E. Reusing ontologies on the Semantic Web: A feasibility study. *Data Knowl. Eng.* **2009**, *68*, 905–925. [\[CrossRef\]](#)
- Zulkarnain, N.Z.; Meziane, F.; Crofts, G. A Methodology for Biomedical Ontology Reuse. In Proceedings of the International Conference on Applications of Natural Language to Information Systems, Salford, UK, 22–24 June 2016; Springer: Berlin/Heidelberg, Germany, 2016.
- Ochs, C.; Perl, Y.; Geller, J.; Arabandi, S.; Tudorache, T.; Musen, M.A. An empirical analysis of ontology reuse in BioPortal. *J. Biomed. Inform.* **2017**, *71*, 165–177. [\[CrossRef\]](#)
- Alharbi, R.; Tamma, V.; Grasso, F. Requirement-Based Methodological Steps to Identify Ontologies for Reuse. In Proceedings of the International Conference on Advanced Information Systems Engineering, Limassol, Cyprus, 3–7 June 2024; Springer: Berlin/Heidelberg, Germany, 2024.
- Fernández-López, M.; Gómez-Pérez, A. Overview and analysis of methodologies for building ontologies. *Knowl. Eng. Rev.* **2002**, *17*, 129. [\[CrossRef\]](#)
- Brendish, N.J.; Malachira, A.K.; Beard, K.R.; Armstrong, L.; Lillie, P.J.; Clark, T.W. Hospitalised adults with pneumonia are frequently misclassified as another diagnosis. *Respir. Med.* **2019**, *150*, 81–84. [\[CrossRef\]](#)
- Kanwal, K.; Asif, M.; Khalid, S.G.; Liu, H.; Qurashi, A.G.; Abdullah, S. Current Diagnostic Techniques for Pneumonia: A Scoping Review. *Sensors* **2024**, *24*, 4291. [\[CrossRef\]](#)
- Uschold, M.; King, M. Towards a Methodology for Building Ontologies. In Proceedings of the IJCAI'95 Workshop on Basic Ontological Issues in Knowledge Sharing, Montreal, Canada, 13 April 1995.
- Uschold, M.; King, M.; Moralee, S.; Zorgios, Y. The enterprise ontology. *Knowl. Eng. Rev.* **1998**, *13*, 31–89. [\[CrossRef\]](#)
- Grüniger, M.; Fox, M.S. Methodology for the Design and Evaluation of Ontologies. In Proceedings of the IJCAI'95 Workshop on Basic Ontological Issues in Knowledge Sharing, Montreal, Canada, 13 April 1995.
- Uschold, M.; Grüniger, M. Ontologies: Principles, methods and applications. *Knowl. Eng. Rev.* **1996**, *11*, 93–136. [\[CrossRef\]](#)
- Fernández-López, M.; Gómez-Pérez, A.; Juristo, N. Methontology: From Ontological Art Towards Ontological Engineering. In Proceedings of the Ontological Engineering AAAI-97 Spring Symposium Series, Stanford, CA, USA, 24–26 March 1997.
- Bachimont, B.; Isaac, A.; Troncy, R. Semantic Commitment for Designing Ontologies: A Proposal. In Proceedings of the International Conference on Knowledge Engineering and Knowledge Management, Sigüenza, Spain, 1–4 October 2002; Springer: Berlin/Heidelberg, Germany, 2002.
- Schultz, D.J. *IEEE Standard for Developing Software Life Cycle Processes*; IEEE: Piscataway, NJ, USA, 1997; pp. 1074–1997.
- Noy, N.F.; McGuinness, D.L. *Ontology Development 101: A Guide to Creating Your First Ontology*; Stanford Knowledge Systems Laboratory Technical Report KSL-01-05 and Stanford Medical Informatics Technical Report SMI-2001-0880; Stanford University: Stanford, CA, USA, 2001.
- Al-Aswadi, F.N.; Chan, H.Y.; Gan, K.H. Automatic ontology construction from text: A review from shallow to deep learning trend. *Artif. Intell. Rev.* **2020**, *53*, 3901–3928. [\[CrossRef\]](#)
- Faure, D.; Nédellec, C. A Corpus-Based Conceptual Clustering Method for Verb Frames and Ontology Acquisition. In Proceedings of the LREC Workshop on Adapting Lexical and Corpus Resources to Sublanguages and Applications, Granada, Spain, 26 May 1998.
- Maedche, A.; Staab, S. Ontology Learning. In *Handbook on Ontologies*; Springer: Berlin/Heidelberg, Germany, 2004; pp. 173–190.
- Cimiano, P.; Völker, J. text2onto. In Proceedings of the International Conference on Application of Natural Language to Information Systems, Alicante, Spain, 15–17 June 2005; Springer: Berlin/Heidelberg, Germany, 2005.

26. Shamsfard, M.; Barforoush, A.A. Learning ontologies from natural language texts. *Int. J. Hum. Comput. Stud.* **2004**, *60*, 17–63. [\[CrossRef\]](#)
27. Hahn, U.; Romacker, M. The SYNDIKATE Text Knowledge Base Generator. In Proceedings of the First International Conference on Human Language Technology Research, San Diego, California, 18–21 March 2001.
28. Gillani Andleeb, S. From Text Mining to Knowledge Mining: An Integrated Framework of Concept Extraction and Categorization for Domain Ontology. Ph.D. Thesis, Budapesti Corvinus Egyetem, Budapest, Hungary, 2015.
29. Navigli, R.; Velardi, P.; Gangemi, A. Ontology learning and its application to automated terminology translation. *IEEE Intell. Syst.* **2003**, *18*, 22–31. [\[CrossRef\]](#)
30. Buitelaar, P.; Olejnik, D.; Sintek, M. A Protégé Plug-In for Ontology Extraction from Text Based on Linguistic Analysis. In *European Semantic Web Symposium*; Springer: Berlin/Heidelberg, Germany, 2004.
31. Fortuna, B.; Grobelnik, M.; Mladenic, D. Semi-Automatic Data-Driven Ontology Construction System. In Proceedings of the 9th International Multi-Conference Information Society IS-2006, Ljubljana, Slovenia, 9 October 2006.
32. Biébow, B.; Szulman, S.; Clement, A.J.B. TERMINAE: A Linguistics-Based Tool for the Building of a Domain Ontology. In Proceedings of the International Conference on Knowledge Engineering and Knowledge Management, Sigüenza, Spain, 1–4 October 2002; pp. 49–66.
33. Trokanas, N.; Koo, L.; Cecelja, F. Towards a Methodology for Reusable Ontology Engineering: Application to the Process Engineering Domain. In *Computer Aided Chemical Engineering*; Elsevier: Amsterdam, The Netherlands, 2018; pp. 471–476.
34. Dramé, K.; Diallo, G.; Delva, F.; Dartigues, J.F.; Mouillet, E.; Salamon, R.; Mougin, F. Reuse of termino-ontological resources and text corpora for building a multilingual domain ontology: An application to Alzheimer’s disease. *J. Biomed. Inform.* **2014**, *48*, 171–182. [\[CrossRef\]](#) [\[PubMed\]](#)
35. Bright, T.J.; Furuya, E.Y.; Kuperman, G.J.; Cimino, J.J.; Bakken, S. Development and evaluation of an ontology for guiding appropriate antibiotic prescribing. *J. Biomed. Inform.* **2012**, *45*, 120–128. [\[CrossRef\]](#)
36. Nachabe, L.; Jahan, N. A Proposed Ontology Evaluation Tool to Assist Ontology Engineers in Selecting Ontologies During the Reuse Phase. In Proceedings of the International Knowledge Graph and Semantic Web Conference, Leipzig, Germany, 26–28 November 2024; Springer: Berlin/Heidelberg, Germany, 2024.
37. Hulshof, C.T. *1710d Systematic Reviews and Evidence-Based Guidelines, Two of a Different Kind?* BMJ Publishing Group Ltd.: London, UK, 2018.
38. Sackett, D.L.; Rosenberg, W.M.; Gray, J.A.; Haynes, R.B.; Richardson, W.S. Evidence based medicine: What it is and what it isn’t. *BMJ* **1996**, *312*, 71–72. [\[CrossRef\]](#)
39. Gillois, P.; Chatellier, G.; Jaulent, M.C.; Colombet, I.; Fieschi, M.; Degoulet, P. From paper-based to electronic guidelines: Application to French guidelines. *Stud. Health Technol. Inform.* **2001**, *84 Pt 1*, 196–200.
40. Guarino, N.; Oberle, D.; Staab, S. What is an Ontology? In *Handbook on Ontologies*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 1–17.
41. Grenon, P.; Smith, B.; Goldberg, L. Biodynamic ontology: Applying BFO in the biomedical domain. *Stud. Health Technol. Inform.* **2004**, *102*, 20–38.
42. Arp, R.; Smith, B. Function, role and disposition in basic formal ontology. *Nat. Preced.* **2008**, *1941*, 1. [\[CrossRef\]](#)
43. Aronson, A.R. *Metamap: Mapping Text to the Umls Metathesaurus*; NLM; NIH; DHHS: Bethesda, MD, USA, 2006; pp. 1–26.
44. Katsumi, M.; Grüninger, M. Choosing ontologies for reuse. *Appl. Ontol.* **2017**, *12*, 195–221. [\[CrossRef\]](#)
45. Cui, L.; Zhu, W.; Tao, S.; Case, J.T.; Bodenreider, O.; Zhang, G.-Q. Mining non-lattice subgraphs for detecting missing hierarchical relations and concepts in SNOMED CT. *J. Am. Med. Inform. Assoc.* **2017**, *24*, 788–798. [\[CrossRef\]](#)
46. Yamagata, Y.; Kozaki, K.; Imai, T.; Ohe, K.; Mizoguchi, R. An ontological modeling approach for abnormal states and its application in the medical domain. *J. Biomed. Semant.* **2014**, *5*, 23. [\[CrossRef\]](#)
47. Brank, J.; Grobelnik, M.; Mladenic, D. A Survey of Ontology Evaluation Techniques. In Proceedings of the Conference on Data Mining and Data Warehouses (SiKDD 2005), Ljubljana, Slovenia, 17 October 2005.
48. Porzel, R.; Malaka, R. A Task-Based Approach for Ontology Evaluation. In Proceedings of the ECAI Workshop on Ontology Learning and Population, Valencia, Spain, 22–24 August 2004; Citeseer: Princeton, NJ, USA, 2004.
49. Maedche, A.; Staab, S. Measuring Similarity Between Ontologies. In Proceedings of the International Conference on Knowledge Engineering and Knowledge Management, Sigüenza, Spain, 1–4 October 2002; Springer: Berlin/Heidelberg, Germany, 2002.
50. Brewster, C.; Alani, H.; Dasmahapatra, S.; Wilks, Y. Data Driven Ontology Evaluation. In Proceedings of the International Conference on Language Resources and Evaluation, Lisbon, Portugal, 24–30 May 2004.
51. Lozano-Tello, A.; Gómez-Pérez, A. Ontometric: A method to choose the appropriate ontology. *J. Database Manag. (JDM)* **2004**, *15*, 1–18. [\[CrossRef\]](#)
52. Gómez-Pérez, A. Evaluation of ontologies. *Int. J. Intell. Syst.* **2001**, *16*, 391–409. [\[CrossRef\]](#)
53. Vrandečić, D. Ontology Evaluation. In *Handbook on Ontologies*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 293–313.

54. Zhu, X.; Fan, J.-W.; Baorto, D.M.; Weng, C.; Cimino, J.J. A review of auditing methods applied to the content of controlled biomedical terminologies. *J. Biomed. Inform.* **2009**, *42*, 413–425. [[CrossRef](#)] [[PubMed](#)]
55. Burton-Jones, A.; Storey, V.C.; Sugumaran, V.; Ahluwalia, P. A semiotic metrics suite for assessing the quality of ontologies. *Data Knowl. Eng.* **2005**, *55*, 84–102. [[CrossRef](#)]
56. Aruna, T.; Saranya, K.; Bhandari, C. A Survey on Ontology Evaluation Tools. In Proceedings of the 2011 International Conference on Process Automation, Control and Computing, Coimbatore, India, 20–22 July 2011.
57. Corbridge, C.; Rugg, G.; Major, N.; Shadbolt, N.; Burton, A. Laddering: Technique and tool use in knowledge acquisition. *Knowl. Acquis.* **1994**, *6*, 315–341. [[CrossRef](#)]
58. Poveda-Villalón, M.; Gómez-Pérez, A.; Suárez-Figueroa, M.C. Oops! (ontology pitfall scanner!): An on-line tool for ontology evaluation. *Int. J. Semant. Web Inf. Syst. (IJSWIS)* **2014**, *10*, 7–34. [[CrossRef](#)]
59. Doran, P.; Tamma, V.; Iannone, L. Ontology Module Extraction for Ontology Reuse: An Ontology Engineering Perspective. In Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management, Lisbon, Portugal, 6–10 November 2007.
60. Kumar, S.; Baliyan, N. Quality Evaluation of Ontologies. In *Semantic Web-Based Systems*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 19–50.
61. Halland, K.; Britz, K. Investigations into the use of SNOMED CT to enhance an OpenMRS health information system. *S. Afr. Comput. J.* **2011**, *47*, 33–45. [[CrossRef](#)]
62. El-Sappagh, S.; Franda, F.; Ali, F.; Kwak, K.-S. SNOMED CT standard ontology based on the ontology for general medical science. *BMC Med. Inform. Decis. Mak.* **2018**, *18*, 76. [[CrossRef](#)] [[PubMed](#)]
63. Shah, T.; Rabhi, F.; Ray, P.; Taylor, K. A guiding framework for ontology reuse in the biomedical domain. In Proceedings of the 2014 47th Hawaii International Conference on System Sciences, Waikoloa, HI, USA, 6–9 January 2014.
64. Johnson, A.E.; Pollard, T.J.; Shen, L.; Lehman, L.-W.H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Celi, L.A.; Mark, R.G. MIMIC-III, a freely accessible critical care database. *Sci. Data* **2016**, *3*, 160035. [[CrossRef](#)] [[PubMed](#)]
65. Metlay, J.P.; Waterer, G.W.; Long, A.C.; Anzueto, A.; Brozek, J.; Crothers, K.; Cooley, L.A.; Dean, N.C.; Fine, M.J.; Flanders, S.A.; et al. Diagnosis and Treatment of Adults with Community-acquired Pneumonia. An Official Clinical Practice Guideline of the American Thoracic Society and Infectious Diseases Society of America. *Am. J. Respir. Crit. Care Med.* **2019**, *200*, e45–e67. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.