

# A Semi-automated Method for Domain-Specific Ontology Creation from Medical Guidelines

Omar ElAssy<sup>[0000-0001-6202-2465]</sup>, Rik de Vendt<sup>[0000-0002-2131-0425]</sup>, Fabiano Dalpiaz<sup>[0000-0003-4480-3887]</sup>, and Sjaak Brinkkemper<sup>✉[0000-0002-2977-8911]</sup>

Dept. of Information and Computing Sciences, Utrecht University, The Netherlands  
{o.omarihabelsayedelassy, r.devendt, f.dalpiaz, s.brinkkemper}@uu.nl

**Abstract.** The automated capturing and summarization of medical consultations has the potential to reduce the administrative burden in health-care. Consultations are structured conversations that broadly follow a guideline with a systematic examination of predefined observations and symptoms to diagnose and treat well-defined medical conditions. A key component in automated conversation summarization is the matching of the knowledge graph of the consultation transcript with a medical domain ontology for the interpretation of the consultation conversation. Existing general medical ontologies such as SNOMED CT provide a taxonomic view on the terminology, but they do not capture the essence of the guidelines that define consultations. As part of our research on medical conversation summarization, this paper puts forward a semi-automated method for generating an ontological representation of a medical guideline. The method, which takes as input the well-known SNOMED CT nomenclature and a medical guideline, maps the guidelines to a so-called Medical Guideline Ontology (MGO), a machine-processable version of the guideline that can be used for interpreting the conversation during a consultation. We illustrate our approach by discussing the creation of an MGO of the medical condition of ear canal inflammation (Otitis Externa) given the corresponding guideline from a Dutch medical authority.

**Keywords:** Domain ontology · method engineering · knowledge graph · SNOMED CT · Medical Guideline Ontology.

## 1 Introduction

The automated summarization of conversations may save time and cost, especially in domains where dialogues are structured based on predefined guidelines. Medical consultations are a prime example, as they broadly follow a systematic examination of predefined symptoms and observations to diagnose and treat well-defined conditions affecting fixed human anatomy [?]. The potential of automated conversation summarization depends on the readiness of the structured representations of domain-specific knowledge in a machine-processable format [?]. In particular, rule-based conversation summarization grants explainability that is unachievable with machine learning approaches, thus, ontologies can efficiently represent domain guidelines in such rule-based applications [?].

Medical consultations typically last 5–10 minutes in which about 1,000–1,500 words are spoken by care provider and patient. Medical reports contain between 20 and 40 highly standardized terms with abbreviations for frequent words (e.g., pt for patient) [?]. In our Care2Report research project [?,?], we combine *domain ontology learning* with the concept of *ontological conversation interpretation* in a two-layered architecture (Fig. 1). The semantic interpretation of the consultation conversation requires the availability of a domain-specific ontology, that binds spoken words to formally defined medical concepts. Input are medical guidelines and clinical practice guidelines for the patient’s ailment type that are available from medical professional associations. Many ontology learning techniques (linguistic analysis, inductive logic programming, and statistical learning [?]) can be applied to generate a domain-specific ontology, in our case the Medical Guideline Ontology (MGO). As medical conversations require optimal precision, we prefer linguistic parsing analyzers (see Sec. 5). The SNOMED CT medical glossary [?] serves for the medical concept identification and linking.

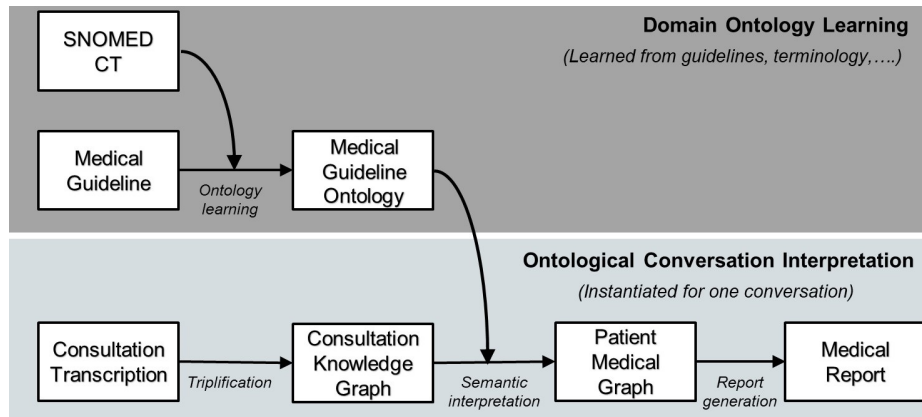


Fig. 1: Ontological Conversation Interpretation in Care2report [?], with the Medical Guideline Ontology supporting semantic interpretation.

The *ontological conversation interpretation pipeline* enables the consultation summarization in four steps: (1) transcribing the audio file by utilizing an automated speech recognizer (ASR), (2) generation of OWL triples (aka triplification) into a complete consultation knowledge graph, (3) semantic interpretation of these triples by means of the Medical Guideline Ontology to a significantly smaller patient medical graph [?], and (4) the generation of a medical report that is presented to the care provider for editing, approval and uploading into the Electronic medical file of this patient.

We adopt the distinction of Ehrlinger and Wöß [?]: the medical ontology acts as a schema on a generic layer serving multiple similar consultations, whereas the knowledge graph represents an instantiated fact base as spoken during the consultation. In the context of conversation summarization, the ontology repre-

sents the diverse and primarily textual domain knowledge (e.g., human anatomy and medical guidelines) in a machine-processable fabric of concepts and relationships. Manually building a domain ontology is time consuming and error prone; thus, it should be carried out using ontology learning: constructing and integrating a machine-readable semantic-rich ontology (semi-)automatically [?].

However, tools and techniques for ontology learning primarily aim to extract knowledge from general data corpora as it was, and still is, driven by the necessity of linking either openly available data (e.g., DBpedia [?]) or corporate-specific business data (e.g., Google Knowledge Graph [?]). Therefore, research is needed to develop a method that enables the systematic representation of domain guidelines in machine-processable ontologies. This paper introduces the notion of *ontological conversation interpretation* (Fig. 1) by presenting a method for developing ontologies in the medical domain to support the expansion of the automated conversation summarization system Care2Report [?].

Care2Report’s summarization pipeline [?] relies on an ontology embodying the domain’s vocabulary and guidelines that act as the structured container to be filled with the multimodal information from the medical consultation. The information extracted from the conversation populates the ontology to generate the rule-adhering report that the physician checks before uploading to the Electronic Medical Records system (EMR) [?]. This paper investigates elements of the second stage of the Care2Report pipeline: the definition of a Medical Guideline Ontology, a machine-processable version of a medical guideline.

To illustrate, we use a medical guideline for ear canal inflammation (Otitis Externa) from the Dutch College of General Practitioners [?]. The medical domain knowledge is assembled from the terminology in the Systematized Nomenclature of Medicine - Clinical Terms (SNOMED CT) [?].

This paper addresses the two aspects of the ontology learning method: the notational side represented by the Medical Guideline Ontology (MGO), and the procedural method expressed in a Process-Deliverable Diagram (PDD) model and by an algorithm. The design science research cycle, as described by Wieringa [?], is fitting to answer the research question: *How to systematically construct ontologies from the human anatomy and medical guidelines?*

The paper makes three contributions:

- We introduce and formalize the Medical Guideline Ontology (MGO), a domain ontology that constitutes a machine-readable version of a medical guideline and that describes the relevant aspects concerning the patient’s anatomy, symptoms, physician’s observations, diagnosis, and treatments.
- We define a procedural method to develop such ontologies in the form of a Process-Deliverable Diagram model, refined into an algorithm that can be at the basis of automated tooling.
- We illustrate the MGO and the procedure to the case of the external ear canal inflammation.

*Paper organization.* After reviewing the related work in Sec. 2, we introduce the MGO in Sec. 3 and its formalization in Sec. 4. We then present our method in Sec. 5 and its application in Sec. 6. Finally, we draw conclusions in Sec. 7.

## 2 Guidelines and Nomenclature in Medical Informatics

**Medical Guidelines.** A medical guideline is a document with recommendations that aim to optimize patient care based on a systematic review of evidence and on weighing the benefits and harms of alternative care options [?]. It consists of definitions and procedural instructions for executing an anamnesis, diagnosis and treatment in care provisions that aim to advance care quality, improve the patient’s health and well-being, and support medical decision-making. Many (inter)national medical authorities publish and maintain medical guidelines [?].

In the Netherlands, both the Dutch College of General Practitioners (Nederlands Huisartsen Genootschap - NHG) and the Dutch Federation of Medical Specialists (Federatie Medisch Specialisten - FMS) publish numerous guidelines [?,?], only the former is used in this paper. The guidelines include sections about prognosis, common causes and background, physical examination and diagnosis guidelines, treatment policy, consultations and referral guidelines (if any), and control and future patient check-ups.

The symptoms and observations indicating a condition and the treatments recommended for such condition by the guidelines are relevant to the construction and population of the Medical Guideline Ontology (MGO) and the related sub-ontologies, as will be detailed in the coming sections. The MGO, the consultation knowledge graph, and the consultation report aim to serve as a representational artifact, rather than supporting the physician’s decision making; thus, the reasoning behind which treatment to choose is beyond our scope. Therefore, the MGO should contain all treatment possibilities as options, while the physician’s discretion will decide which ones to use.

**Medical Nomenclature: SNOMED CT.** The Systematized Nomenclature of Medicine - Clinical Terms (SNOMED CT) is currently the world’s most extensive clinical terminology system [?]. This paper uses SNOMED CT as an ontology source representing human anatomy and terminology hierarchy to identify relevant medical concepts from all the potential concepts in the textual medical guidelines. The terminology structure is in hierarchical formations of concepts defined and connected to each other by relationships, with identifiers for machine use and descriptors for human readability. The top node of the SNOMED CT hierarchy is occupied by the root concept *SNOMED CT concept*, and nineteen direct subtypes of it are the *top level concepts* that provide the structure of the SNOMED CT. Different conditions and medical consultations use a subset of the available concept hierarchies.

## 3 Medical Guideline Ontology

The Medical Guideline Ontology (MGO) is a domain ontology that represents a medical guideline in a machine-processable format. In the context of the considered guideline, the MGO represents the relevant patient’s anatomy and symptoms, the physician’s observations, diagnosis and prescribed treatments. Note

that symptoms are subjective abnormalities that the patient perceives, while observations are objective abnormalities detected by the physician [?]. The schema of the MGO is illustrated in Fig. 2.

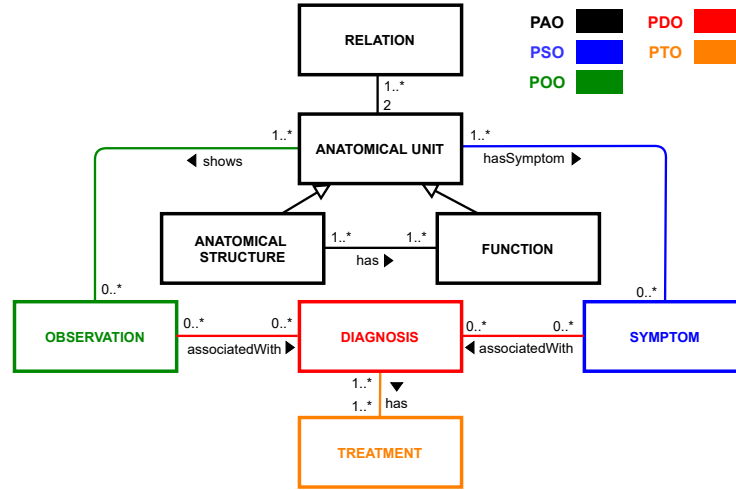


Fig. 2: Schema of the Medical Guideline Ontology

The MGO consists of five (sub)ontologies. The Patient Anatomy Ontology (PAO) depicts the human anatomical structures and functions (within the context of the guideline). The Patient Symptoms Ontology (PSO) represents the complaints and anomalies (symptoms) that a patient may report. The Patient Observations Ontology (POO) represents all the observations that the physician may make about the patient’s condition. The Patient Diagnosis Ontology (PDO) describes the physician’s diagnosis of a patient’s condition. Finally, the Patient Treatment Ontology (PTO) describes all the treatments prescribed by the physician, including medications, instructions, referrals, or additional medical tests.

The Patient Anatomy Ontology (PAO) is the foundation for the knowledge representation in the MGO. The PAO can be built based on existing resources like the Foundational Model of Anatomy [?], or it can utilize an existing hierarchical terminology structure like SNOMED CT, as in this research. Within the SNOMED CT hierarchies, the human anatomy is represented in one section: *SNOMED CT concept* ⇒ *Body structure* ⇒ *Anatomical or acquired body structure* [?]. The key concepts in the PAO are those of ANATOMICAL UNIT and RELATION. The former is specialized by ANATOMICAL UNIT (e.g., ear, left eye) and FUNCTION (e.g., hearing), with the HAS relationship linking the two, e.g., ‘an ear has the function of hearing’. The RELATION is meant to represent other types of links between anatomical units, such as part-of, next-to, etc.

The medical report produced by the automated conversation summarization system Care2Report to upload into the EMR follows a predefined structure [?]

of four sections: (i) Subjective: what the patient reports; (ii) Objective: what the physician identifies; (iii) Evaluation: assessment and diagnosis by the physician; and (iv) Plan for treatment and follow up.

Following the same arrangement, the MGO consists of a composite of four sub-ontologies reflecting the four aspects of reporting and a fifth foundational ontology for human anatomy. Each of the those (anatomy, symptoms, observations, diagnosis, and treatment) builds on the previous one(s) to include further knowledge into the resulting MGO shown in Fig. 2. The MGO provides the notational aspect of the ontology development method proposed in this paper.

The MGO and its instantiation into a consultation knowledge graph with information from a specific consultation (with a concrete patient) provide a graphical representation of ontology-based knowledge where every two concepts and the relationship between them form a triple. Examples of such triples are introduced in the coming sections.

## 4 Formalization of the MGO

The structure and the interaction between the various ontologies are relatively intuitive for simple conditions concerning comparatively simple body structures. However, adding the complete human anatomy and the guidelines from several medical authorities can make the ontologies less intuitive to understand and communicate. Therefore, it is essential to introduce formal descriptions of the ontologies and their components; this section introduces a few of these descriptions. Formally defining the MGO requires the definition of the following sets:

Set	Description
$B$	anatomical structures of the body.
$F$	anatomical functions of the body.
$A$	anatomical units ( $A = B \cup F$ ).
$S$	medical symptoms (reported by a patient).
$O$	medical observations (observed by a physician).
$V$	possible values to be assigned to symptoms $s \in S$ or observations $o \in O$ .
$E$	explicit diagnoses.
$I$	implicit diagnoses.
$D$	medical diagnoses: $D = E \cup I$ .
$T$	medical treatments.

The elements in these sets may vary depending on the scope of the ontologies. Defining the comprehensiveness of these sets depends on the goal of the medical reporting and on the number and complexity of the guidelines. We leave the investigation of this aspect to future research.

The definitions of the ontologies comprising the MGO are defined with respect to a patient  $p$  (the subject of the guideline, which is also the most comprehensive anatomical unit) as follows:

Ontol.	Vertices	Edges
PAO	$A = B \cup F$	$\{(b_1, b_2) \mid b_1, b_2 \in B \wedge b_1 \text{ is a direct anatomical sub-part of } b_2\}$ $\cup$ $\{(b, f) \mid b \in B \wedge f \in F \wedge b \text{ has an anatomical function } f\}$
PSO	$A \cup S \cup V$	$\{(a, s), (s, v) \mid a \in A \wedge s \in S \wedge v \in V\}$ i.e., the symptoms
POO	$A \cup O \cup V$	$\{(a, o), (o, v) \mid a \in A \wedge o \in O \wedge v \in V\}$ i.e., the observations
PDO	$\{p\} \cup D$	$\{(p, d) \mid d \in D\}$ i.e., the possible diagnoses of a generic patient
PTO	$\{p\} \cup T$	$\{(p, t) \mid t \in T\}$ i.e., the possible treatments of a generic patient

As per Sec. 3, the MGO is a domain ontology representing the guideline’s contents in terms of patient anatomy and symptoms as well as physician’s observations, diagnosis and treatments:  $MGO = PAO \cup PSO \cup POO \cup PDO \cup PTO$ .

Finally, some of the rules that need to be coded into the system to define the relationships between the entities are listed below, consulting domain experts is expected to add to this list. We provide some examples, but note that the creation of a comprehensive list of rules is domain-specific and goes beyond the purpose of this paper:

1. Each anatomical structure (b) is a part of another anatomical structure (b) unless it is the complete body structure (b\*).  
 $\forall b_1 \exists b_2 : \text{isPartOf}(b_1, b_2) \quad b_1, b_2 \in B, \neg (b_1 = b^*)$
2. Each function (f) is assigned to one or more anatomical structure (b).  
 $\forall f \exists b : \text{hasFunction}(b, f) \quad b \in B, f \in F$
3. Each symptom (s) is associated with one or more anatomical unit (a).  
 $\forall s \exists a : \text{hasSymptom}(a, s) \quad s \in S, a \in A$
4. Each observation (o) is associated with one or more anatomical unit (a).  
 $\forall o \exists a : \text{hasObservation}(a, o) \quad o \in O, a \in A$
5. Patients are diagnosed with at least diagnosis (d).  
 Given  $p, \exists d : \text{diagnosedWith}(p, d) \quad d \in D$
6. Patients are treated with a treatment (t). A treatment encompasses any physician’s prescription, including medications, instructions for the patient to follow, referral to a specialist, further tests, or any other procedure.  
 Given  $p, \exists t : \text{treatedWith}(p, t) \quad t \in T$
7. An explicit diagnosis (e) is associated with a symptom (s) or an observation (o).  
 $\forall e \exists s \exists o : \text{associatedWith}(s, e) \vee \text{associatedWith}(o, e)$   
 $e \in E, s \in S, o \in O$
8. An implicit diagnosis is not explicit. That is, an implicit diagnosis (i) is neither associated with a symptom (s) nor an observation (o).  
 $\forall i \forall e : \neg (i = e) \quad i \in I, e \in E$

## 5 Method for Systematic Creation of Medical Ontologies

While the previous section introduced the MGO for representing a medical guideline in a machine-processable format, this section introduces the procedural method perspective outlined in a Process-Deliverable Diagram (PDD) model.

The PDD illustrates the activities and artefacts of a specific process [?]. The model emphasizes the relationships between the activities and their deliverables by connecting them with dotted arrows across the diagram [?].

### 5.1 Ontology Creation PDD

Fig. 3 shows the used PDD, in which the process is broken down into eight main activities. For simplicity, activities three to seven are illustrated (and manually performed) sequentially; however, they can also be executed in parallel.

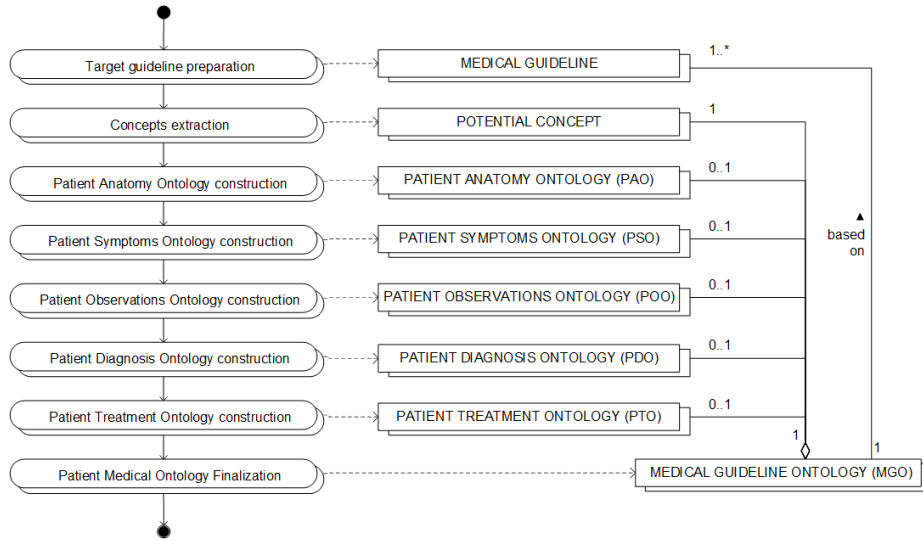


Fig. 3: PDD of the Ontology Creation Process

1. **Target guideline preparation:** The guideline of the medical condition is selected from the medical authority's website, translated (if necessary), scraped, and prepared for the following concept extraction activity.
2. **Concept extraction:** The relevant sections of the guidelines are identified, including sections describing symptoms, physical examination, and treatment plans. As nouns are the natural language representation of things, ideas and notions, they identify the concepts to be extracted. Thus, the potential concepts to extract are all nouns and noun phrases in the relevant sections that will be mapped in the next steps against the SNOMED CT to identify the constituent concepts of the ontology (e.g., anatomical units, symptoms, observations, and treatments). Some potential concepts will not be used in the ontology as they represent general nouns used in the text.
3. **Patient Anatomy Ontology (PAO) construction:** The guideline concepts corresponding to SNOMED CT *anatomical concepts* are identified and converted into a hierarchy from which the PAO is constructed.



4. **Patient Symptoms Ontology (PSO) construction:** The concepts identified in the medical guideline sections describing symptoms are mapped against the corresponding SNOMED CT hierarchies (*findings* and *disorders*) to build the PSO.
5. **Patient Observations Ontology (POO) construction:** Similar to the previous activity except dealing with physician observations instead of patient-described symptoms. Thus, the relevant guideline sections are different.
6. **Patient Diagnosis Ontology (PDO) construction:** The medical condition or disease discussed in the guideline is the diagnosis associated with symptoms and observations to construct the PDO.
7. **Patient Treatment Ontology (PTO) construction:** The concepts identified in the treatment-related sections of the guideline are mapped against the corresponding SNOMED CT hierarchies (*procedure*, *substance*, *dose form*, and *physical object*) to build the PTO.
8. **Medical Guideline Ontology Finalization:** All the previous (sub)ontologies are combined along with needed information (e.g., prefixes) to construct the complete MGO. This activity also includes checks to validity by confirming the lack of disjoint concepts.

### 5.2 A Detailed Look on Ontology Creation

Fig. 4 expands the PAO construction activity to show the various subactivities and the resulting deliverables. The PDD activities constructing the remaining

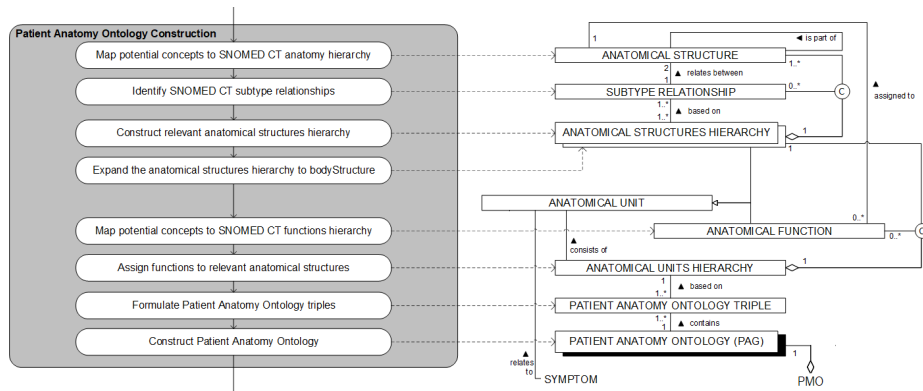


Fig. 4: Patient Anatomy Ontology (PAO) construction

ontologies follow the same general high-level pattern to build the relevant ontology based on the list of extracted potential concepts from the guidelines. Potential concepts are mapped at each stage to both the relevant sections of the guidelines (e.g., physical examination and treatment policy) and the relevant hierarchies of the SNOMED CT (e.g., findings and disorders). Thus, the concepts

that show in both relevant modules are the appropriate candidate concepts for the ontology at hand (e.g., symptom concept, observation concept or treatment concept). An anatomical unit (that can have symptoms, observations, diagnosis, or treatment) can be either an anatomical structure or a function. An *anatomical structure* is “a physical anatomical entity and a physical object, ... it consists of parts that are themselves anatomical structures”, anatomical structures are “...localized to a specific area or combine and carry out one or more specialized [anatomical] *functions* of an organism.” [?].

For a detailed illustration of the complete PDD and a description of all the (sub)activities and concepts, refer to the technical report [?].

### 5.3 Algorithm

Algorithm 1 refines the PDD in Fig. 3 by explaining the derivation of an MGO from a medical guideline  $MG$  for a disease  $D$ , and from a human anatomy graph  $HA$ . This algorithm is at the basis of our current Care2Report pipeline.

---

#### Algorithm 1 Medical Guideline Ontology Generation

---

**Input:**  $MG$  a Medical Guideline for disease  $D$ ,  
 $HA$  the standard complete human anatomy graph,  
**Output:** a generic ontology  $MGO$  for this medical guideline

```

1: function BUILDONTOLOGYBYTRIPLES( $MG, D, HA$ )
2:   for all  $sent \in MG$  do
3:      $P_{sent} \leftarrow \text{EXTRACTNOUNPHRASES}(sent)$ 
4:     for all  $cs \in P_{sent}$  do
5:       if  $cs \in HA$  then
6:          $AnatD \leftarrow AnatD \cup \{cs\}$ 
7:        $PAO \leftarrow \text{SUBGRAPH}(AnatD, HA)$   $\triangleright$  The concepts in  $AnatD$  and their links
8:        $MGO \leftarrow \emptyset, PSO \leftarrow \emptyset, POO \leftarrow \emptyset, PDO \leftarrow \emptyset, PTO \leftarrow \emptyset,$ 
9:       for all  $ae \in PAO$  do
10:        for all  $sent \in MG$  do
11:           $P_{sent} \leftarrow \text{EXTRACTNOUNPHRASES}(sent)$ 
12:          for all  $cs \in P_{sent}$  do
13:            if  $cs \in \text{SNOMED.symptom}$  then
14:               $PSO \leftarrow PSO \cup \{\langle ae, hasSymptom, cs \rangle, \langle cs, associatedWith, D \rangle\}$ 
15:            if  $cs \in \text{SNOMED.observation}$  then
16:               $POO \leftarrow POO \cup \{\langle ae, hasObservation, cs \rangle, \langle cs, associatedWith, D \rangle\}$ 
17:            if  $cs \in \text{SNOMED.diagnosis}$  then
18:               $PDO \leftarrow PDO \cup \{\langle patient, diagnosedWith, cs \rangle\}$ 
19:            if  $cs \in \text{SNOMED.treatment}$  then
20:               $PTO \leftarrow PTO \cup \{\langle patient, treatedWith, cs \rangle, \langle D, hasTreatment, cs \rangle\}$ 
21:           $MGO \leftarrow PAO \cup PSO \cup POO \cup PDO \cup PTO$ 
22:        return  $MGO$ 

```

---

The algorithm iterates over all sentences in the medical guideline (lines 2–6). Each sentence is parsed to extract the noun phrases (e.g., Otitis, ear canal) that are mentioned in the text (line 3). These noun phrases are analyzed (lines 4–6) so that only those corresponding to elements of the human anatomy graph are retained and stored into the variable *AnatD*. The *PAO* is defined (line 7) as the subgraph of the *HA* that includes only the concepts in *AnatD*, their descendants in the part-of hierarchy, and the relationships between the retained concepts.

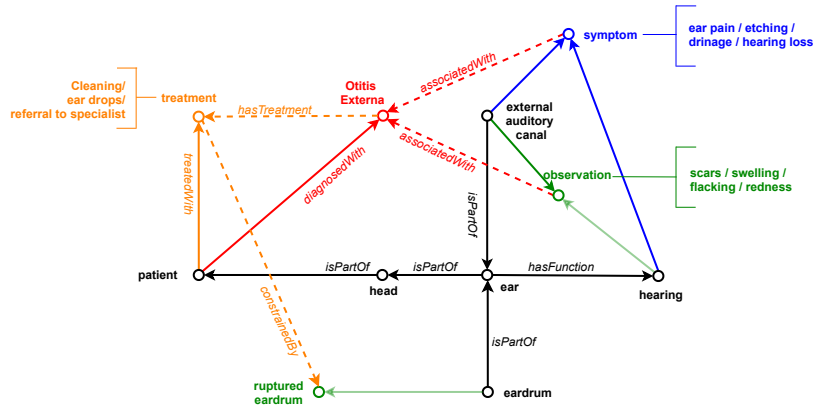
After their initialization as empty sets (line 8), the remaining ontologies *MGO*, *PSO*, *POO*, *PDO*, *PTO* are populated in lines 9–21. The elements added to the *PAO* are iterated over (lines 9–20): the sub-cycle starting in line 10 iterates again over the medical guidelines, which are again parsed by extracting the noun phrases (line 11). Each of these noun phrases contributes to populating the various ontologies, depending on whether the noun phrase is a symptom (lines 13–14), observation (lines 15–16), diagnosis (lines 17–18), or treatment (lines 19–20). Finally, the *MGO* is defined (line 21) as the union of the five ontologies.

## 6 Application to the Otitis Externa case

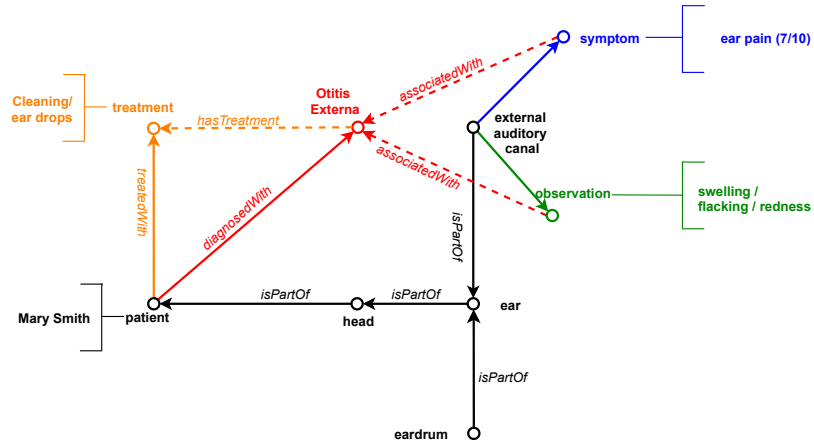
To illustrate the *MGO* and apply the derived procedure to create an ontology, we use the NHG guideline for the inflammation of the external ear canal (Otitis Externa) [?]. The condition is chosen as an example for the relative simplicity of the associated guidelines and the lack of complicated medical procedures, terminology, or differential diagnosis. Otitis Externa is an inflammation caused by a disturbance in the acidic environment of the ear canal and it is usually associated with swimming. The symptoms reported by the patient may include ear pain, ear itching, fluid drainage from the ear, and hearing loss. In addition, the physician examines both the complaint-free ear and the affected ear for signs of scars, swelling, flaking, redness, the state of the eardrum, etc. [?].

Some symptoms (e.g., hearing loss) are not required for an Otitis Externa diagnosis, indicating that while the *MGO* representation of the condition should include it as a possible symptom, some specific consultations might have this symptom present while others do not. Also, the guidelines indicate that the physician should check the eardrum; however, Otitis Externa is not associated with any observations regarding the eardrum, suggesting that any observation of the eardrum (e.g., rupture) might indicate a different diagnosis. As for the treatment, the guideline recommends that the physician should instruct the patient on how to clean the infected ear properly, and prescribe ear drops. Also, referral to a specialist is recommended if the condition does not improve promptly or if the patient is from a specific vulnerable group (e.g., elderly and diabetic).

As an example, we consider a fictitious consultation for an Otitis Externa patient suffering from ear pain (no indication of itching, fluid drainage, or hearing loss), and the examination shows swelling, redness, and skin flaking in the external ear canal. Figure (5a) represents the general *MGO* of the Otitis Externa condition with indication of the guideline sections forming it, while figure (5b) is the consultation-specific knowledge graph for the fictitious patient.



(a) MGO of the Otitis Externa Condition [?]



(b) Example Consultation Knowledge Graph of an Otitis Externa patient

Fig. 5: Otitis Externa representation

The MGO can be represented in triples for machine processing; for example, some of the triples portrayed in figure (5b) are explained below:

- The human anatomy is represented in a hierarchy of anatomical structures connected to their parent structures using *isPartOf* relationships. For example,  $\langle \text{externalAuditoryCanal}, \text{isPartOf}, \text{ear} \rangle$ .
- The second part of the PAO is the assignment of anatomical functions to the anatomical structures performing them as in  $\langle \text{ear}, \text{hasFunction}, \text{hearing} \rangle$ .
- Symptoms and observations triples define the PSO and the POO. For example, the following triples refer to a pain symptom:  $\langle \text{externalAuditoryCanal}, \text{hasSymptom}, \text{symp\_2} \rangle$ ,  $\langle \text{symp\_2}, \text{symptom}, \text{earPain} \rangle$ ,  $\langle \text{symp\_2}, \text{hasValue}, 7/10 \rangle$ .
- The patient diagnosis and treatments are expressed in triples; for example:  $\langle \text{patient}, \text{diagnosedWith}, \text{OtitisExterna} \rangle$ ,  $\langle \text{patient}, \text{treatedWith}, \text{earDrops} \rangle$ .

Applying the procedure in the PDD to Otitis Externa produces the ontology shown in Fig. 6 using the WebVOWL web application [?]. For conciseness, the ontology does not visualize the top-level classes (e.g., anatomy, symptom) and the is-a relationship linking the other concepts to them. More details on ontology construction and the resulting ontology can be found in the technical report [?].

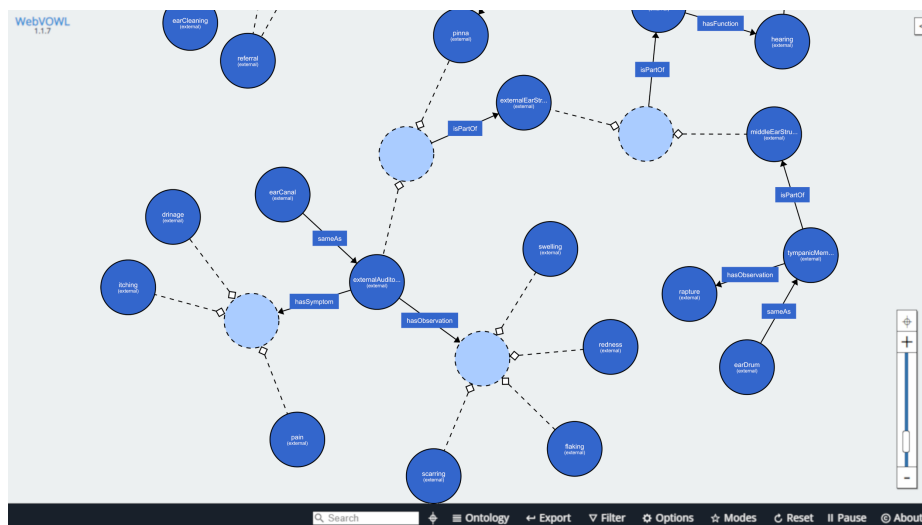


Fig. 6: Partial WebVOWL Visualization of the Otitis Externa Ontology

The method was applied manually, requiring substantial time investment. The full potential of the Care2Report system relies on developing an automatic process to generate ontologies of all medical conditions using guidelines from various countries and in different languages. The ontology development relies on data from two sources: SNOMED CT, which can be easily mapped to an OWL ontology; and the text-based medical guidelines that pose the real challenge for knowledge extraction and structuring to be tackled in future research.

## 7 Conclusion

We introduced *ontological conversation interpretation* by creating a method for the systematic creation of medical ontologies from the SNOMED CT terminology and anatomical hierarchies, combined with the medical guidelines of different medical conditions. The resulting MGO is a formalized notation, and the PDD is a procedural guide to systematically create ontologies and enrich them by connecting new guidelines to the existing definitions of anatomy, symptoms, observations, and treatments; and adding more if the existing concepts are insufficient to represent the guideline fully. Furthermore, we have applied our approach to the case of Otitis Externa to illustrate its application and feasibility.

*Limitations and Future Directions.* The method needs to be validated on guidelines of more complex and varying conditions. For example, the NHG guideline for “Non-traumatic knee complaints” details a method for differentiating between various similar conditions and thus does not follow the typical sections in most NHG guidelines. This may require an evolution of our MGO. Moreover, some guidelines only point the physician towards the tests and measurements to monitor without detailing the results or values to look for, presumably because the results are hard to detail in the text while the medical professionals understand them. For example, the guidelines advise performing an electrocardiogram test in several cardiovascular conditions without detailing the expected outcomes. This is a challenge as the source information is incomplete; thus, the degree of possible automation is limited. The Care2Report multi-modal input architecture [?] aims to eventually allow the integration of (some of) the measurement data, but an ontology to represent this knowledge still needs to be developed. Automation is another research challenge. We have applied the method manually, as explained in Sec. 6, and the correct and complete automated interpretation of textual medical guidelines is a far-fetched goal [?]. Research is necessary toward the creation of assisted, interactive methods that support our approach. This includes research towards automating the creation of ontologies of the human anatomy and medical guidelines from different medical authorities.

This research aims to foster research in a societally-relevant field: increasing the quality of healthcare via semi-automated methods that may relieve medical professionals from their administrative burden. We make a step in this direction by laying down the formal foundations for the construction of semi-automated systems that support the interpretation of conversations using ontologies.

## References