An Ontology-Based Approach to Decision Support for Healthcare Workflows

Cristian Cocos[*,a], Wendy MacCaull[a], Bryan Kramer[b] and Markus Latzel[b]
[a]*Centre for Logic and Information, St. Francis Xavier (StFX) University, 54B St. Mary's Street, Antigonish NS, Canada B2G 2W5, tel.: +1 (902) 867-4931, fax: +1 (902) 867-1397*
[b]*Palomino System Innovations Inc., 533 College Street, Suite 404, Toronto ON, Canada M6G 1A8, tel.:+1 (416) 964-7333, fax: +1 (866) 256-7866*

Abstract. We provide the description of a healthcare workflow management system, with emphasis on the data structures associated, and their interface with the system. The two major data structures are: (a) an access control ontology, which contains the access control policy, and (b) a palliative and seniors' care ontology, which stores palliative and seniors' care knowledge, and offers decision support for the workflow. Both these ontologies take advantage of healthcare knowledge contained in several established healthcare terminology standards (such as SNOMED-CT® and ICNP®), as well as information gathered from focus groups and interviews with clinicians and healthcare managers. Their development principles are expounded, their structure analyzed, and their interaction with the system elaborated by means of querying examples and interaction scenarios.

Keywords: Ontology, workflow, patient-centered care, access control, SNOMED-CT

1. Introduction

Researchers at the StFX Centre for Logic and Information (CLI) are currently developing a careflow management system for palliative and seniors' care, which is scheduled to be deployed at several hospitals in rural Nova Scotia in the near future ([18]). This is part of a collaborative six-year R&D project involving clinicians and managers of the local health authority (Guysborough Antigonish Strait Health Authority (GASHA)) and an industry the partner (Toronto-based Palomino System Innovations Inc.). This system aims at streamlining workflow by improving documentation and communication. It will offer several high-end features such as adaptability, flexibility, semantic web integration (in view of projected interoperability with globally-accepted biomedical controlled vocabularies), etc. In order to comply with the desideratum of patient-centered health care delivery (whereby health care data is owned by the patient), our endeavor requires a major commitment to ensuring data privacy and confidentiality. In an era marked by the advent of electronic health records (EHR), controlling and restricting access to healthcare data is of key importance. It is our belief that an ontology-structured knowledge base constitutes an ideal mechanism for the implementation/representation of a role-based access control (RBAC) policy. Here we present the design and employment of such an RBAC ontology (the Access Control Ontology (ACO)), which is intended to fulfill the commitment to privacy and confidentiality. We also introduce an ontology-based mechanism, used to offer decision support for the palliative and seniors' care workflow. While this article may be viewed as having predominantly a practical value, several issues are of certain theoretical interest, most of which deal with our treatment of the OWL version of SNOMED-CT.

The paper is structured as follows: we first cover the background of the careflow management system, and then we present the access control part by undertaking a brief recount of the particular access control

[*] Corresponding author. E-mail: ccocos@stfx.ca.

scenario to be implemented, followed by an outline of the access control ontology and its interaction with the workflow. Next we give an overview of the palliative and seniors' care ontology (PCSO), including a brief listing of its sources and development principles, proceeded by a presentation of the workflow-ontology interaction scenario via examples. We conclude by outlining some directions for future work. Some of the material on the ACO covered in section 3 is contained in [2].

2. Background

The term "palliative care" refers to the physical, psychological, spiritual and practical care given to patients and their families when dealing with issues associated with serious illness. As patients are usually part of a family when care is provided, the patient and family are treated as a unit. The main focus of palliative care is to ease patients' and their families' suffering, and to help them cope with their difficulties. The patients'/families' needs and care priorities are at the core of care provision, hence the expression "patient-centered" care. This contrasts with traditional models of care, that are typically illness focused, provider driven, or determined by the care setting. Palliative care is ideally delivered by an interdisciplinary care team with the requisite knowledge and expertise required to address the many domains of issues that are commonly experienced by such categories of patients and families. The goal of the seniors' care initiative is to enable senior individuals to continue living independently in community at their highest possible functional level and with the best quality of life through an integrated and comprehensive continuum of care in the community. Senior's care is based on several principles, out of which the most important are: (a) a holistic approach toward an aging population that promotes maximum function, independence, comfort and quality of life, (b) a continuum of care that is seen by seniors and their support networks as comprehensive, connected, and easily accessed, (c) enrichment of communities by embracing an aging in place philosophy, (d) dignity for the client and their family, including respect for ethnic, cultural, and spiritual needs.

The high-level conceptual blueprint of the careflow management system (CfMS) in the works is shown in figure 1 ([18]):

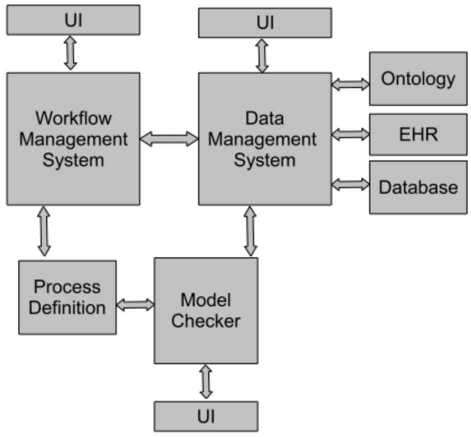

Fig. 1. A Careflow Management System ([18]).

The two major components relevant from our point of view:

- The Workflow Management System (WfMS). A WfMS [27] is a system that defines, creates and manages the execution of workflows through the use of software, running on one or more workflow engines, who is able to interpret the process definition, interact with workflow participants and, where required, invoke the use of information technology tools and applications. A process definition is the representation of a business process in a form which supports automated manipulation, such as modeling, or enactment by a WfMS.
- The Data Management System (DMS). A DMS manages data and knowledge during the execution of healthcare workflows. It includes the following sub-components:
  - the knowledge base—in our case, the ontologies. These provide an efficient way of managing medical health and organizational knowledge. The CfMS allows the ontologies to guide and inform the workflow process and offer decision support using real-time access to knowledge.
  - the patient database—a.k.a. the Electronic Health Records (EHR). An EHR is an evolving concept defined as a systematic collection of electronic health information about individual patients or populations [28]. It is a record in digital format designed to be shared across different healthcare settings. Such records may include a whole range of data in comprehensive or summary form, including

demographics, health history, medication and allergies, laboratory test results, radiology images, and billing information.

Some prominent terminological resources that occupy a central role in the development of the ontological knowledge bases are SNOMED-CT ([13]), and the International Classification for Nursing Practice (ICNP) ([11]). Both standards are currently used extensively with the aim of building a palliative and seniors' care ontology (PSCO) that will guide and offer decision support for the palliative and senior's care workflow; they also play a (slightly less emphatic) role in the development of the ACO.

The workflow management system thus interacts with (among others) both ACO and PSCO in order to obtain access control clearance for the system user and, respectively, to guide the workflow in deciding between alternative trajectories on the basis of palliative and seniors' care knowledge.

3. Access Control (ACO)

Representing the roles of a RBAC policy as classes of an access control ontology emerges as a natural choice, especially when the ontology in question accommodates roles natively. We have thus decided to adopt an upper-level ontology where roles figure prominently in the asserted hierarchy. Basic Formal Ontology (BFO) has enjoyed a certain success in semantic web and ontology circles, and lends itself quite naturally to the type of enterprise in which we have engaged. From our perspective, four BFO classes are of immediate interest to ACO: BFO:role, BFO:object, BFO:process, and BFO:generically_dependent_continuant. For a presentation of BFO the reader is invited to consult [25] and [26].

Two major types of entities ("resources") are subject to access control under the ACO scenario: informational items and actions. Database fields are paradigm examples of the former category (e.g., patient ID, patient name, primary diagnosis etc.), though reports also constitute informational content that carries access restrictions (e.g., incident reports). The latter class of entities has been assembled from actions implemented by the workflow management system, and represents actions that the system user can (or cannot) invoke, such as form/report printing, faxing, phoning, appointment scheduling, etc.

Corresponding to these two types of resources, two mechanisms have been employed, so that users' access control credentials will be checked at login time, and consequently some of these information fields and/or actions will be cleared for access. Here are the defining features of the access control policy implemented:

1. Roles are organized hierarchically, hence permissions granted at higher levels of the hierarchy are inherited by roles lower in the hierarchy ("hierarchical RBAC" in [22]);
2. Resources are organized hierarchically—e.g., allowing access to a whole form means allowing access to all of its fields;
3. Access control constraints can be provided for each form field and action individually;
4. Unlike a purely linear hierarchy, we have disjoint roles, whereby not all roles intersect with respect to the range of permissions;
5. Database fields can be accessed as both read only and write;
6. System users may have multiple roles with regard to the same patient; in such circumstances, the clearance level assigned is a union of permissions for all roles involved.

The BFO:role branch contains the main mechanism of our RBAC implementation, whereby roles relevant from the point of view of access control in a palliative/seniors care setting have been grouped on categories representing clearance and permission levels for informational items and actions respectively. As mentioned above, the "role" branch is intended to accommodate roles such as "Patient," "Caregiver," "Medical Doctor" etc. Most of the classes that populate this branch have been imported straight from SNOMED-CT: we thus saw fit to import SNOMED's "Occupation (occupation)" (SCTID_14679004) branch in its entirety, as it contains a significant portion of relevant roles. Interestingly enough, SNOMED does not comprise a properly titled "role" class, nor does it contain any class to this effect, hence most of the classes that we have assembled under the "role" branch have been culled from various (and quite heterogeneous) places in SNOMED—mostly "occupation" and "person." As such, our endeavor can also be viewed as a SNOMED-CT restructuring enterprise by bringing it in line with BFO—an undertaking that, as of this writing, is being contemplated under the auspices of the IHTSDO (the SNOMED-CT custodian). HealthCareRole (see figure 2) is another ACO-specific class that groups several SNOMED-CT classes such as "Caregiver (person)," "Clinical trial participant (person)" etc. The SNOMED "Relative (person)" class (SCTID_125677006) comprises the usual relative roles, as well as some more or less exotic degrees of human inter-relation, such as “Great-great grand-parent (person),” “Fraternal twin brother (person),” etc.

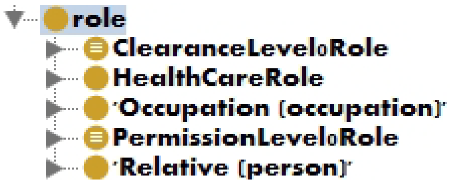

Fig. 2: Roles in ACO.

The core ACO mechanisms reside under the Clearance-Level0Role and PermissionLevel0Role branches. These defined classes contain five, respectively two, defined (sub)classes labeled Clearance-Level*x*Role, ClearanceLevel*xw*Role and, respectively, PermissionLevel*y*Role (figure 3), where *x* takes values from 1 to 4, and *y* and *w*, 1 and 2. (The number of clearance and permission levels is subject to future amendments.) Due to spatial considerations we will refrain from elaborating the labeling scheme, though the labeling principles should emerge clearly from figure 3. We will detail the content of the clearance level hierarchy, and occasionally point out the differences between it and the permission level hierarchy.

Each of the ClearanceLevel*x*Role classes is defined as a union of roles that have a certain security clearance level, which is to be dictated by real-life security and privacy considerations. ClearanceLevel2Role, for example, is a union of the following SNOMED_CT classes: “Physiotheraist/occupational therapist (occupation),” “Social worker (occupation),” “Community nurse (occupation),” “Pain management specialist (occupation)” and “Pharmacist (occupation)” (see the Manchester syntax rendering in figure 4).

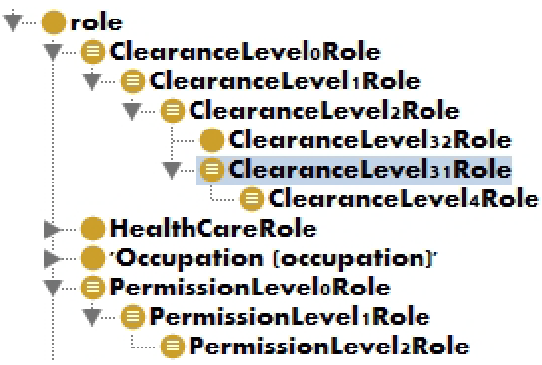

Fig. 3: ACO clearance and permission levels

A corollary of this particular chain-like structure is that all ClearanceLevel*y*Role roles have also clearance level *x* for $x \geq y$, in other words, all information that is accessible to a community nurse (say) is also accessible to a clinical oncologist (a ClearanceLevel0Role role), but not vice versa; in general, a parent inherits the clearance level of its children. Furthermore, figure 3 shows that the ClearanceLevel2Role class has two children: ClearanceLevel32Role and Clearan-ceLevel31Role, which have been stipulated as disjoint, such that the two actually constitute a partition of ClearanceLevel2Role. A similar story goes for permission levels (see also figure 3), and while we have not designed any disjoint permission levels, these can be captured in a manner completely analogous to clearance levels.

We further added classes that represent actual information and actions that are subject to access control, and devised a method to tie clearance and permission level roles with that information. We have thus opted to add six relations (“object properties”) together with reciprocal (“inverse”) relations:

1. **hasWriteAccessTo** and **hasReadAccessTo** are relations from “role” to “DatabaseField” (see below), and, respectively, to “'Report (record artifact)' or DatabaseField.” The former is a child of the latter, hence whoever has write access to some resource obviously has read access to it. Their reciprocals are **writeAccessibleTo** and **readAccessibleTo**, with the former being a child of the latter;
2. **invokableBy** is a relation from ACO:SystemProcedure to BFO:role, and plays the same role as either of **read/writeAccessibleTo** vis-à-vis action permissions;
3. **hasRole** (adapted from ICNP) is a relation whose range is BFO:role, and is intended to tie roles with their bearers; its reciprocal is **roleOf**;
4. **hasClearanceLevel** ties the SNOMED-CT class of “Homo sapiens (organism)”—a subclass of BFO:object—with a clearance level. Its reciprocal relation is **clearanceLevelOf**;
5. similarly, **permissionLevelOf** connects permission levels with elements of “Homo sapiens (organism).”

The classes that represent controlled *information* have been grouped under the BFO:generically_dependent_continuant branch, which is the BFO entity designed to account for in-

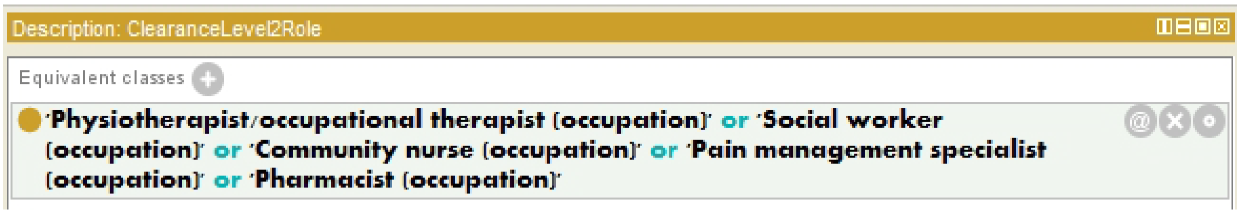

Fig. 4: ClearanceLevel2Role definition

formational and other abstract entities. ACO imports the Informational Artifact Ontology (IAO) ([7]) under the BFO:generically_dependent_continuant class, and entities of the IAO are used to define ACO's metadata fields. The root of the IAO is "information content entity" (IAO_0000030). We have added the following chain of ACO-specific subclasses as a child of IAO_0000030: "Proposition," "Declarative-Proposition," and "DatabaseField" (see figure 5). And while the latter inclusion might strike some as strange (at least *prima-facie*), we are confident that this choice will emerge as ontologically compelling once one contemplates the concrete case of a numerical value (126, say), housed in the "Blood Pressure" field of a medical form—hence the "BloodPressure-Field" contains a *proposition* about (the value of) blood pressure, namely "This patient's blood pressure has value 126 mmHg."

The type of information that makes the subject of access control policies has been, at this point, limited to community palliative and seniors' care records. Our focus has been to extract information fields from GASHA's palliative and seniors' care programs. The GASHA workflow collects patient information using several forms. Given the provisional status of most of the forms as of this writing, we chose to represent in the ACO taxonomy six of the more stable ones—see below the children of DatabaseField—while the rest of them will be added as they reach a reasonably stable status.

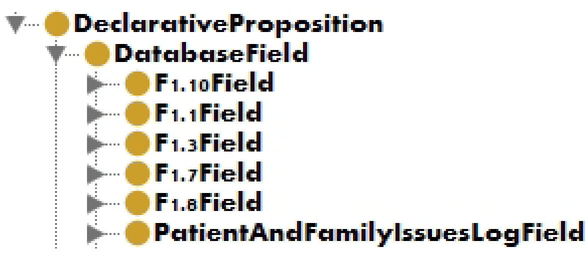

Fig. 5: Database fields.

Similarly, classes that represent controlled *actions* comprise the ACO:SystemProcedure class, which is a child of BFO:process (figure 6).

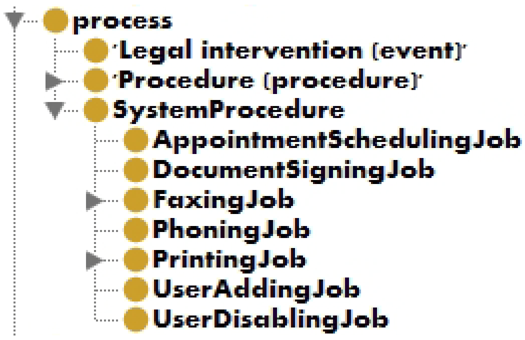

Fig. 6: System procedures.

Finally, all the classes that make the subject of access control policies have been endowed with restrictions that outline their corresponding clearance/permission level. Here are a few examples:

The HospitalizationHistoryField (ACO_0000101), a child of the F1.8Field class (i.e., a field on GASHA form F1.8) has been designed so as to not be accessible to a clearance level 2 role—which would obviously make it inaccessible to any role in a higher clearance level category (i.e., 31, 32 and 4) as well. Figure 7 shows this:

Fig. 7: Hospitalization History field restrictions.

F1.10 form will not be accessible to a clearance level 3 role, which means that all fields on this form will bear the same restriction. This is a case where the restriction has been implemented at the parent class level (F1.10Field (ACO_0000129)), instead of adding it to all component fields individually (see figure 8).

Fig. 8: Form F1.10 restrictions.

PhoningJob will only be invokable by a PermissionLevel2Role:

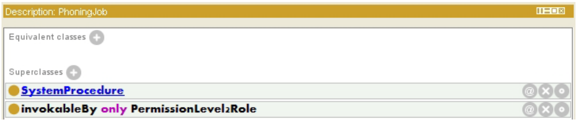

Fig 9: PhoningJob restrictions.

The last of the relevant BFO classes is BFO:object, whose only direct child is the “Organism (organism)” SNOMED class (SCTID_410607006). This one, in turn, contains “Homo sapiens (organism)” as an only subclass (SCTID_337915000), which represents the main ACO role-bearer.

ACO currently contains approximately 10,000 classes, and its level of DL expressivity is $\mathcal{SROIF}$, which is N2ExpTime-hard ([24]). Reasoning and querying the ACO using widely available free or open source OWL-DL reasoners (Pellet, FaCT++, etc.) currently poses no tractability problems, nor do we envisage any in the future.

According to the basic interaction scenario, access control clearance is checked at login time, by querying ACO upon user login. ACO provides in return a list of GASHA form fields and reports whose access is forbidden to the user, and a list of system actions permitted. The workflow system—that, by design, is in possession of the list of all actions and informational items—acts accordingly, by blocking access to the requisite actions and information entities.

An example of a DL query is the following: say Individual_1 is an instance of the “Homo sapiens (organism)” class that has been defined as a psychotherapist—i.e., is an instance of “**hasRole** some 'Psychotherapist (occupation)'.” A query such as “not (**accessibleTo** some (**roleOf** value Individual_1))” will reveal all the form fields that are *not* accessible to Individual_1. Another way of defining an individual is to assign it a clearance level outright, without going through the detour of specifying its role, thus sparing the reasoner the effort of figuring out the individual’s clearance level. If Individual_2 is defined as “**hasClearanceLevel** some ClearanceLevel0Role,” the query “not (**accessibleTo** some (**roleOf** value Individual_2))” will return all the form fields not accessible to Individual_2.

4. Palliative and Seniors’ Care Ontology (PCSO)

The main requirements of a knowledge base associated with the workflow management system to be built by the CLI lab is to represent palliative and seniors’ care knowledge as exhaustively and accurately as possible, and in a way that lends itself to querying in real time. More specifically, the project promises as a core functionality the pairing of the palliative and seniors’ care workflow with a knowledge base that will interface with the workflow and provide domain-knowledge-based decision support, so that the ontology/knowledge base can be regarded as guiding and informing the workflow process using real-time access to dynamic knowledge. It is not hard to imagine that the requirements of exhaustiveness and real-time tractability are in a state of latent tension, though a discussion of this matter will be postponed for later. In view of this conflict, we have lent preference to construing and developing the workflow-associated palliative and seniors’ care knowledge base as a *general* theory of the palliative and seniors’ care domain, therefore avoiding any reference to individual patients, and individual patient records. It is not only that a lumping together of general and particular knowledge would have considerably enlarged the ontology—thus burying any prospects of real-time reasoning—but ontologies as knowledge representation tools do not particularly lend themselves to a dynamic treatment such as the one imposed by the fluid character of patient data. Thus, by itself, the ontology will be unable to interact in any meaningful way with the workflow; in order to offer guidance to, and direct the workflow, the ontology will have to be supplemented by a proper patient database, containing the electronic health records (EHR) for each and every patient subject to the workflow.

The palliative and seniors’ care ontology will, therefore, have to adequately represent and capture the domain of hospice (or hospice-type) palliative and seniors’ health care; more concretely, sentences (ideally all) purportedly about palliative care should be expressible in some DL syntax (Manchester syntax for example) involving ontology terms. As a corollary, the ontology must employ a rich vocabulary, covering a wide spectrum of terms, from human anatomy and clinical practice to nursing routines and administrative aspects, etc. Other constraints that have been imposed are:

- The ontology should be either developed natively in OWL DL, or should be easily translatable in OWL DL;
- The ontology should comply with the norms of ontological decorum established by the OBO Foundry ([14]);
- It should, at the very least, reflect two results of Canada-wide consensus: (a) the Canadian Hospice Palliative Care Association's Model to Guide Hospice Palliative Care: Based on National Principles and Norms of Practice ([5]), and (b) the Canadian Council on Health Services Accreditation Hospice Palliative Care Standards ([4]).

The venture of building such a controlled vocabulary not only that it cannot ignore the previous related work undertaken by biomedical ontologists and standardization bodies worldwide, but due to the sheer size of the domain to be covered, it cannot realistically be carried out without appeal to such work. The ideal scenario would be to proceed at a straight import (via the owl:imports mechanism) of a relatively stable, and sufficiently rich biomedical ontology (the "background" ontology), to which specific hospice palliative terms would have to be added—either manually, or by some automated means or other.

In light of this, we have decided for PCSO purposes to use SNOMED-CT almost entirely, as this is currently the only serious option as far as a medical terminology is concerned. For ontology development, the OBI MIREOT ("minimal information to reference external ontology terms") import mechanism will also be employed ([9]), via its web-based OntoFox implementation ([10]).

As in the case of ACO, we have adopted BFO as upper-level ontology, and we imported the widely used, and universally acclaimed, BFO-based Relation Ontology (RO) [29], which will provide a kernel of relational terms. Besides using the terms collected manually by CLI researchers from various palliative and seniors' care resources and clinical documents, components of the term pool used to construct PCSO (some of which have also been mentioned above) include:

1. The Ontology for Biomedical Investigations (OBI) ([8]): The OBI project is developing an integrated ontology for the description of biological and clinical investigations. The ontology targets the design of an investigation, the protocols, instrumentation and material used, the data generated and the type analysis performed on it. Currently OBI is being built under the Basic Formal Ontology (BFO), and is available in OWL format;
2. The International Classification for Nursing Practice (ICNP®), which is a unified nursing language system. It is a compositional terminology for nursing practice that facilitates the development of and the cross-mapping among local terms and existing terminologies. ICNP version 2 consists of only one root axis – Nursing Phenomena, which in turn has seven axes – Client, Focus, Location, Judgment, Means, Time, and Action. ICNP can be provided, upon request, in OWL format;
3. ACGT Master Ontology (ACGT MO) ([12]): The ACGT MO is an ontology that aims to represent the domain of cancer research and management, with special emphasis on mammary carcinoma ("breast cancer"), Wilms' tumor (nephroblastoma) and rhabdoid tumor. The development of the MO was guided and reviewed by researchers from two pre-existing clinical trials, namely a breast cancer related trial on Topoisomerase II Alpha Gene Amplification and Protein Overexpression Predicting Efficacy of Epirubicin (TOP) and "Nephroblastoma (Wilms Tumor) - Clinical Trial and Study SIOP 2001" by the International Society of Pediatric Oncology. In order to achieve the aim of supporting unified data annotation for these trials, the developers had to shape the MO as a cross-section of a multitude of sub-domains, all of which are vitally important to clinical cancer management and research. In effect, the outcome of this effort is best seen not as a comprehensive domain ontology, but rather as an application ontology tailored to the needs of the ACGT software system. Available in OWL format;
4. C. Kuziemsky's Environmental Scan on the Minimum Data Set (MDS) for Palliative Care ([3]);
5. Canadian Hospice Palliative Care Association's Model to Guide Hospice Palliative Care: Based on National Principles and Norms of Practice ("the National Model");

6. Canadian Council on Health Services Accreditation Hospice Palliative Care Standards (“the Standards”).

These elements of the term pool were carefully inspected, and terms extracted and examined with an eye to their proper place in a taxonomy of the palliative and seniors’ care domain. Work in analyzing these pools of data is, nevertheless, ongoing, and terms are extracted and processed in accordance with OBO Foundry and other general ontological principles, in order to secure a clean taxonomy.

GASHA forms have also been thoroughly analyzed, and terms extracted with the additional aim of building the patient database (EHR)—which is primarily a task of our industry partner.

Ontology development proceeds along two major development paths: terminology and reasoning, corresponding to two potential uses of the PCSO: (a) terminology alignment and standardization, and (b) workflow guidance. The former treats the PCSO as a controlled vocabulary, while the latter employs its full resources via the assertive knowledge (axioms, restrictions, taxonomy) contained within.

The desiderata of terminology alignment and biomedical knowledge wealth have led to the adoption of SNOMED-CT as a centerpiece of PCSO, and, to a lower extent, of ACO. In order to ensure a proper level of integration of SNOMED, we have initiated a re-grouping of SNOMED terms and branches under the BFO upper-level categories. The process of building the PCSO has thus placed a major emphasis on the (ongoing) enterprise of re-shaping and re-arranging of SNOMED-CT. To our knowledge, the only previous attempt in this regard is [23], which has been examined closely, and built upon: we have, for example, followed Hogan in subsuming SNOMED-CT’s “Procedure” concept to a BFO “process” (see #3 in the list below), and “Organism” to BFO:object. Briefly, some of the major moves that have been undertaken in this respect are the following:

1. Moved several classes from SNOMED’s “Observable entity (observable entity)” top-level branch to BFO’s “quality”;
2. Moved several classes from SNOMED’s “Function (observable entity)” to BFO’s “function”;
3. Moved the “Procedure (procedure)” class to BFO’s “process”;
4. Moved “Legal intervention (event)” to BFO’s “process”;
5. Moved “Death (event)” to BFO’s “process_boundary”;
6. Populated BFO:material_entity with classes from SNOMED-CT (e.g. “Specimen (specimen),” “Substance (substance),” etc.);
7. Populated PCSO’s “DeclarativeProposition” branch with classes from SNOMED-CT.

It should be emphasized that the above list is just a sampling of developmental moves, as countless others involve manipulating single classes or minor branches instead of whole upper-level branches. Moreover, even though entire SNOMED branches have been moved under supposedly proper BFO classes, we do not exclude at this point that many classes may have been miss-assigned as a consequence, due to the heterogeneous character of SNOMED-CT groupings and classification. While we have gone to great lengths to avoid such outcomes, the sheer terminological wealth manipulated made it impossible to verify each and every single class. Finally, even though SNOMED contains an impressive number of relations (“object properties”), given the obscure character of many of them, we have seen ourselves compelled to add many such relations manually: **admittedTo**, **dischargedFrom**, **undertakes**, **undergoes**, etc.

The effort to integrate ICNP deserves a special mention. Its almost 3,000 classes and object properties have been carefully scrutinized; to our surprise, a great deal of overlap between ICNP and SNOMED-CT has been uncovered, such that the great majority of ICNP terms are comfortably covered by SNOMED-CT equivalents. In the very few instances where an equivalent SNOMED-CT term has not been found, the corresponding ICNP term has been added to the ontology.

A great deal of time and energy has been invested in examining some of the most glaringly PC-specific term lists included in the term pool, and ensuring that ontology-worthy terms are captured in the PCSO either under the very same label, or under a slightly modified sticker. In the very few instances where an equivalent term has not been found, we have added it to the ontology. Nevertheless, some very resilient terms are still under consideration: ontology development is currently ongoing both by re-analyzing ontologically-problematic terms, and by expanding the range of our scrutiny to the remainder of the term pool items.

PCSO interacts with the workflow in tandem with the patient database (EHR), given that the ontology only contains generic information, but no information pertaining to individual patients. As stressed above, PCSO will provide logic-based guidance for the workflow at the so-called decision points, which are to be suitably chosen from those points in the workflow where it branches, and where palliative and seniors' care knowledge is involved in the decision. The sequence of actions is the following: (a) the workflow reaches a decision point; (b) the ontology is queried with the relevant patient data contained in the EHR, and furnishes information regarding the workflow branch that the process is supposed to follow; (c) the process follows the path indicated by the ontology query. Most of the time, however, the ontology output will be used to make *recommendations* for the system users, instead of simply prompting a certain mandatory conduct. Some of the reasons for involving an ontology in the decision process—rather than simply building the knowledge directly in the workflow via the WfMS—are: (i) as opposed to rewriting the workflow (hence tinkering with the workflow engine), updating an ontology is a relatively simple task, (ii) the criteria involved in the decisions are, most of the times, cumbersome, knowledge intensive, and intellectually demanding, hence capturing them via workflow is quite costly in terms of the number of (repetitive) workflow branches required and, ultimately, system resources; representing these criteria as OWL rules in PCSO makes the utmost sense, especially since the they are comfortably captured by the underlying logic of OWL 2.

We will provide in the following two concrete scenarios where such interaction takes place, accompanied by the corresponding queries.

Based on the information contained in one of the GASHA forms (the "Issues Log"), a patient may be referred to a number of professionals (e.g. Social Worker, Physiotherapist, Occupational Therapist etc.) (figure 10).

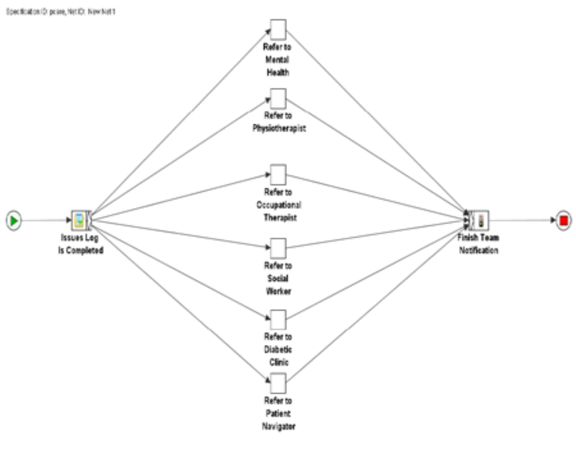

Fig. 10: Issues Log decision point.

The criteria according to which such professionals are chosen can be expressed as Manchester syntax restrictions; one such example is the following criterion: "Adult patients manifesting psychological symptoms such as fears and independence concerns, and whose primary diagnosis is not cancer, should be referred to a social worker, unless they are supposed to undergo emergency kidney surgery." The corresponding Manchester syntax expression/restriction is: Patient and (**hasAge** some (Age and (**hasValue** some int[$\geq$"18"integer])) and (**hasSymptom** some (Fear or IndependenceConcern)) and not (**hasPrimaryDiagnosis** some Cancer) and not (**hasPlannedOperation** some EmergencyKidneySurgery). Consequently, a defined class ReferredToSocialWork is defined as equivalent to the above restriction. Upon querying, the workflow system launches a query to the ontology containing the relevant data collected via the Issues Log: the patient age, the list of symptoms, patient's primary diagnostic, and the list of planned surgeries. In effect, the query amounts to a satisfaction check, whereby the patient data is tested as to whether it complies with the above criterion or not. The ontology returns a Boolean value, according to which the patient in question belongs to the ReferredToSocialWork class or not. This will be treated by the workflow as a suggestion, since the decision ultimately belongs to the clinician currently interacting with the patient.

The current workflow draft contains a step that checks whether a patient that has been referred for palliative/seniors' care is, in fact, justified in receiving attention in this respect—see figure 11.

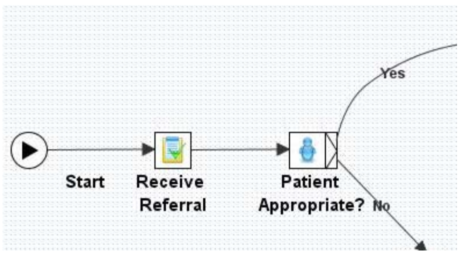

Fig. 11. Admission decision point.

The criteria involved in such a decision are multiple; for reasons of explanation simplicity we will exemplify a simplified criterion, according to which a patient is accepted only if he/she has been diagnosed with one of the diseases listed as eligible for care under the current palliative and seniors' care guidelines, and if he/she is a permanent resident, not currently suffering from an infectious disease, and if his/her life expectancy is no longer than three months. Concretely, this can be done by defining a class PCDisease as equivalent to the disjunction of all eligible diseases (Disease$_1$ or Disease$_2$ or … or Disease$_n$), and then defining another class EligiblePatient as "Patient and (**hasDisease** some PCDisease) and (**hasResidentialStatus** some PermanentResident) and not (**hasDisease** some InfectiousDisease) and (**hasLifeExpectancy** some (temporal_interval (and **hasValue** some int[$\leq$"3"integer])))." Querying the ontology is done by reading the content of the "Primary Diagnosis," "Disease," "Residency Status," and "Life Expectancy" fields from the patient's EHR, and then building a SPARQL query that checks whether the primary diagnosis is in the Disease$_1$ to Disease$_n$ list, whether any of the patient's diseases is an infectious disease, whether his/her residential status is "PermanentResident," and whether his/her life expectancy is less or equal to three months. In effect, the query will return an answer to the question of whether the patient in question is (or is not) a member of the EligiblePatient class. The answer to the query will again be treated by the workflow as a suggestion to be presented to the deciding factor.

One final issue to be discussed in connection with the PCSO is the tractability aspect. Given the size of the ontology, checking its consistency currently takes around three hours on a run-of-the-mill desktop PC[2] using the latest Pellet version. Querying is faster, though still prohibitive from the point of view of real time action. At this point in time, PCSO contains approx. 310,000 classes, 100 object properties, and its DL expressivity level is $\mathcal{SROIF}$.

We are presently considering several venues in order to handle querying in real time:

- the relational database detour—that is, exporting/storing the ontology in relational tables, and employing database techniques to carry out reasoning/inferencing; such techniques have yielded promising results for some OWL 2 subsets such as OWLPrime ([15]);
- ontology modularization, i.e. extracting a "logically correct" module ([6], p. 15) capable of tackling the requisite queries, and using it "as a way to keep performance of ontology services at an acceptable level" (ibid, p. 7);
- parallel and distributed methodologies such as the MaRVIN platform described in [20], to get the advantages afforded by multithreaded and distributed reasoning.

The second option is also what prompted the spin-off of the ACO from the initially-projected all-encompassing ontology. It is, hence, important to make the following distinction between two developmental strategies, corresponding to the two development paths announced above: (a) build a PCSO for terminological and mapping purposes and (b) build a highly customized PCSO for integration with the palliative care workflow system currently developed by CLI. Building such a highly customized knowledge base requires the existence of a patient database and of a reasonably stable and detailed workflow.

5. Conclusions and Future Work

We have (a) outlined the access control policies and implementation principles that lay the foundation of the RBAC ontology built as part of a palliative and seniors' care workflow project presently developed in Nova Scotia, and (b) given an overview of the palliative and seniors' care ontology that guides the same workflow. A pilot involving clinicians and patients in real-life palliative and seniors' care setting is currently underway in GASHA. This project makes essential use of semantic web techniques, and constitutes a living expression of the efficacy of ontological knowledge bases and their employment in concrete everyday situations, and of the usefulness of standards such as SNOMED-CT. It is our belief that semantically-structured knowledge has yet to bear its fruit as a large-scale implementational venue for health informatics systems, and we view our present endeavor as a contribution towards streamlining efforts targeting massive adoption of such techniques.

The evolution of the interaction between the workflow system and the ACO includes a customization phase, which requires implementing a workflow mechanism that queries the patient/client on specific access control preferences during several predetermined phases of the workflow, plus a mechanism that builds new patient-specific access control ontologies that will be combined with the default ACO described above in order to customize the access control policy for each patient. Further, more sophisticated, access control scenarios may require implementing task-based access control policies; this, however, may require tools more expressive than DL fragments of OWL. Among others, we envisage im-

[2] 2.6 GHz Intel Core 2, 8 GB DDR3 RAM.

plementing an emergency “break the glass” mechanism, according to which certain system users can bypass security policies in case of emergency.

Development of the PCSO will progress under the guidelines detailed above: we will continue to inject more palliative and seniors’ care knowledge into it, particularly via enriching it with terms relating to the *care* aspect (as opposed to terms and knowledge pertaining to strictly *medical* issues), plus workflow-specific axioms and restrictions. The growth and fine-tuning of the PCSO will be essentially informed by the results and experience accumulated during the run of the first pilot. At the same time, our efforts will also be directed at coping with the tractability aspect, and possibly investigating new paths in this respect.

References

[1] B. Tsoumas, S. Dritsas, and D. Gritzalis, An Ontology-Based Approach to Information Systems Security Management in V. Gorodetsky, I. Kotenko, and V. Skormin, (eds.): MMM-ACNS 2005, LNCS 3685, Springer-Verlag Berlin Heidelberg, pp. 151 – 164, 2005.

[2] C. Cocos and W. MacCaull, An Ontological Implementation of a Role-Based Access Control Policy for Health Care Information, 2010, submitted.

[3] C. Kuziemsky, Environmental Scan of Hospice Palliative Minimum Datasets and Databases, workshop presentation at the “Towards a Hospice Palliative Care Minimum Data Set: Clinical, Research and Policy Implications” workshop, St. Francis Xavier University, Antigonish, Nova Scotia April 23-24, 2009.

[4] Canadian Council on Health Services Accreditation, ed., CCHSA’s Accreditation Program 2008: Standards of Hospice Palliative and End-of-Life Services, CCHSA, 2007.

[5] F. Ferris et al., A Model to Guide Hospice Palliative Care: Based on NationalPrinciples and Norms of Practice. Canadian Hospice PalliativeCare Association, 2002.

[6] H. Stuckenschmidt, C. Parent, S. Spaccapietra (eds.), Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization, LNCS 5445, Springer, 2009.

[7] http://code.google.com/p/information-artifact-ontology/

[8] http://obi-ontology.org/page/Main_Page

[9] http://obi-ontology.org/page/MIREOT

[10] http://ontofox.hegroup.org/

[11] http://www.icn.ch/pillarsprograms/international-classification-for-nursing-practicer/

[12] http://www.ifomis.org/wiki/ACGT_Master_Ontology_(MO)

[13] http://www.ihtsdo.org/snomed-ct/

[14] http://www.obofoundry.org/

[15] http://www.oracle.com/technology/tech/semantic_technologies/index.html

[16] J. Bock at al., A Reasoning Broker Framework for OWL, in R. Hoekstra and P.F. Patel-Schneider (eds.), OWL: Experiences and Directions, proceedings of OWLED 2009, http://www.webont.org/owled/2009/, 2009.

[17] J. Bock, Parallel Computation Techniques for Ontology Reasoning, in A Sheth et al. (eds.), ISWC 2008, LNCS 5318, Karlsruhe, Germany, 2008, pp. 901-906.

[18] K. Miller and W. MacCaull, Toward careflow management systems, Journal of Emerging Technologies in Web Intelligence (JETWI) Special Issue, E-health: Towards System Interoperability through Process Integration and Performance Management, 1(2):137–145, 2009.

[19] M.-M. Bouamrane, A. Rector and M. Hurrell, A Hybrid Architecture for a Preoperative Decision Support System Using a Rule Engine and a Reasoner on a Clinical Ontology, in A. Polleres and T. Swift (eds.): RR 2009, LNCS 5837, pp. 242–253, Springer-Verlag, Berlin Heidelberg 2009.

[20] P. Li et al., The Quest for Parallel Reasoning on the Semantic Web, in J. Liu et al. (eds.): AMT 2009, LNCS 5820, pp. 430–441, 2009.

[21] Q. Fang et al., Scalable Distributed Ontology Reasoning Using DHT-Based Partitioning, in J. Domingue and C. Anutariya (Eds.): ASWC 2008, LNCS 5367, pp. 91–105, 2008.

[22] T. Finin et al., ROWLBAC - Representing Role Based Access Control in OWL, in SACMAT’08, June 11–13, 2008, Estes Park, Colorado, USA. Copyright 2008 ACM 978-1-60558-129-3/08/06.

[23] W. Hogan, Aligning the top-level of SNOMED-CT with Basic Formal Ontology, poster presentation at KR-MED 2008, 01 June 2008, http://precedings.nature.com/documents/2373/version/1, 2008

[24] Y. Kazakov $\mathcal{SRIQ}$ and $\mathcal{SROIQ}$ are Harder than $\mathcal{SHOIQ}$, in F. Baader et al. (eds.), DL 2008. Vol. 353 of CEUR Workshop Proceedings, 2008.

[25] T. Bittner and B. Smith, Normalizing Medical Ontol-ogies using Basic Formal Ontology, in Kooperative Versorgung, Vernetzte Forschung, Ubiquitäre Information (Proceedings of GMDS Innsbruck, 26-30 September 2004), Niebüll: Videel OHG, pp. 199–201, 2004.

[26] P. Grenon, B. Smith, and L. Goldberg, Biodynamic Ontology: Applying BFO in the Biomedical Domain, in D. M. Pisanelli (ed.), Ontologies in Medicine, Amsterdam: IOS Press, 2004, pp. 20–38, 2004.

[27] World Health Organization (WHO), eds., Workflow Management Coalition Terminology and Glossary Technical Report (WFMC-TC- 1011) v3.0. Workflow Management Coalition, Winchester, UK, 1999.

[28] T. D. Gunter and N. P. Terry, The emergence of national electronic health record architectures in the united states and australia: Models, costs, and questions. J Med Internet Res, 7(1), 2005.

[29] http://www.obofoundry.org/ro/