

Proceedings
of
THE FIRST
INTERNATIONAL WORKSHOP
ON
DRUG INTERACTION KNOWLEDGE REPRESENTATION
(DIKR 2014)

THE SECOND
INTERNATIONAL WORKSHOP
ON
DEFINITIONS IN ONTOLOGIES
(IWOOD 2014)

THE WORKSHOP
ON
STARTING AN OBI-BASED BIOBANK ONTOLOGY
(OBIB 2014)

*Proceedings
edited by*

*Richard D. Boyce
Mathias Brochhausen
Philip E. Empey
Melissa Haendel
William R. Hogan
Daniel C. Malone
Patrick Ray
Alan Ruttenberg
Selja Seppälä
Christian J. Stoeckert
Jie Zheng*

*The workshops were held in conjunction with the 5th International Conference on Biomedical
Ontology (ICBO 2014)
Houston, TX, USA
October 6-7, 2014*

Preface

This volume collects the papers of three workshops held at the International Conference on Biomedical Ontologies (ICBO 2014) in Houston, TX. For more detailed information about the First International Workshop on Drug Interaction Knowledge Management (DIKR 2014), the Second International Workshop on Definitions in Ontologies (IWOOD 2014) and the Workshop on Starting an OBI-Based Biobank Ontology please refer to the individual workshop prefaces. The editors would like to thank the ICBO organizers, the members of all three scientific committees, the authors and, of course, the participants of the workshops. In addition, we would like to thank Mark Jensen for his advice regarding the creation of CEUR-WS submissions.

December 2014

The editors

TABLE OF CONTENTS

The First International Workshop on Drug Interaction Knowledge Representation (DIKR 2014) <i>R. D. Boyce, M. Brochhausen, P. E. Empey, W. R. Hogan, D. C. Malone</i>	1
Redesign of a clinical decision support system for a drug – drug interaction alert <i>D. Luna, C. Otero, F. G. Bernaldo de Quiros</i>	3
Towards a foundational representation of potential drug - drug interaction knowledge <i>M. Brochhausen, J. Schneider, D. C. Malone, P. E. Empey, W. R. Hogan, R. D. Boyce</i>	16
The Second International Workshop on Definitions in Ontologies (IWOOD 2014) <i>S. Seppälä, P. Ray, A. Ruttenberg</i>	32
Textual and logical definitions in ontologies <i>S. Seppälä, Y. Schreiber, A. Ruttenberg</i>	35
Defining ‘sign’ and ‘symptom’ <i>A. P. Cox, P. L. Ray, M. Jensen, A. D. Diehl</i>	42
An alternative terminology for pain assessment <i>W. Ceusters</i>	49
The Workshop on Starting an OBI-based Biobank Ontology (OBIB 2014) <i>M. Brochhausen, M. Haendel, C. J. Stoeckert, J. Zheng</i>	55
Penn Medicine Biobank Informatics <i>H. Williams, D. Birtwell</i>	57
Towards a Common Semantic Representation of Informed Consent for Biobank Specimens <i>F. J. Manion, Y. He, E. Eisenhauer, Y. Lin, A. Karnovsky, M. R. Harris</i>	61
A Specimen-based View of the World - Using the Biological Collections Ontology to Model Biodiversity Collections <i>R. Walls, R. Guralnick, J. Deck, A. Matsunaga</i>	64

Proceedings
of
THE FIRST
INTERNATIONAL WORKSHOP
ON
DRUG INTERACTION KNOWLEDGE
REPRESENTATION
(DIKR 2014)

*Proceedings
edited by*

*Richard D. Boyce
Mathias Brochhausen
Philip E. Empey
William R. Hogan
Daniel C. Malone*

*DIKR 2014 was held in conjunction with the 5th International
Conference on Biomedical Ontology (ICBO 2014)
Houston, TX, USA
October 6-7, 2014*

Preface

Optimizing drug interaction knowledge representation is a pressing need. Currently, the combination of poor quality evidence and a general lack of drug-drug interaction (DDI) knowledge by prescribers results in many thousands of preventable medication errors each year. While many sources of DDI evidence exist to help improve prescriber knowledge, no meta-data standard currently exists that is built based on the requirements derived from an analysis of the information synthesis workflows of pharmacists and drug compendium editors. Such a metadata standard could enable a more effective synthesis of DDI evidence during tasks such as consulting and guideline development.

The First International Workshop on Drug Interaction Knowledge Representation, held on October 6th 2014, brought clinical and ontology development experts together to discuss:

- a) potential DDI knowledge representation solutions that reflect the state-of-the-art of both the clinical understanding of DDIs and biomedical ontology development,
- b) how to best link DDI ontologies to pre-existing drug terminology efforts, and
- c) roadblocks to the adoption of ontology-driven solutions such as coverage, usability, and scalability. Our aim is to shape the workshop into an annual event that addresses issues of optimizing representation of drug interaction for meaningful use. The workshop's focus is on discussing solutions to bridging the gap between the representation of drug interaction in knowledge managements systems and the requirements by those using that information in clinical practice.

To ensure the clinical relevance of state-of-the-art knowledge representation for drug interactions doesn't get overlooked the workshop started with a keynote lecture by Daniel Malone (University of Arizona) entitled *Jumping the crevasse between assertions of drug interactions and clinical relevance*. (<http://www.slideshare.net/boycer/keynote-maloneclinicalrelevanceofddievidence>)

Acknowledgments

One of the aims of the workshop organizers was to start cooperation between groups developing ontologies related to drug interactions. We greatly appreciated the input of all of participants on the interactive panel discussion:

Panelists:

Moderator: William R. Hogan (University of Florida)

Olivier Bodenreider (National Library of Medicine)

Richard D. Boyce (University of Pittsburgh)
Mathias Brochhausen (University of Arkansas for Medical Sciences)
Daniel Malone (University of Arizona)
Maria Herrero-Zazo (University Carlos III de Madrid)

Organizing Committee:

Richard D. Boyce (University of Pittsburgh)
Mathias Brochhausen (University of Arkansas for Medical Sciences)
Philip Empey (University of Pittsburgh)
William R. Hogan (University of Florida)
Daniel Malone (University of Arizona)

Program Committee:

Richard D. Boyce (University of Pittsburgh)
Mathias Brochhausen (University of Arkansas for Medical Sciences)
Michel Dumontier (Stanford University)
Jon Duke (Regenstrief Institute)
Philip Empey (University of Pittsburgh)
William R. Hogan (University of Florida)
Daniel Malone (University of Arizona)
Alan Ruttenberg (University at Buffalo)
David Weinstein (Wolters Kluwer Health)

Funding Support:

This workshop was partially supported by a grant from the National Library of Medicine:
"Addressing gaps in clinically useful evidence on drug-drug interactions"
(1R01LM011838-01).

Redesign of a clinical decision support system for a drug - drug interaction alert.

Daniel Luna¹, Carlos Otero¹, Fernan Gonzalez Bernaldo de Quiros¹

¹ Health Informatics Department. Hospital Italiano de Buenos Aires. Argentina

Abstract. The clinical decision support systems (CDSS) in general and those related to prescription systems in particular have the potential to reduce and prevent morbidity and mortality associated with adverse events and improve the quality of patient care. Most electronic health records (EHR) has CDSS used as support for clinical decision making. In the context of electronic prescribing systems, knowledge databases necessary for the implementation of systems alerting for drug - drug interactions are not adapted to local contexts of use, generating a high rate of false positives and producing "alert fatigue". Redesigning of the notification system for drug interactions in areas like structuring a knowledge database of drugs, modification and validation of a knowledge database on drug - drug interactions, generating case studies and frequency of interactions at the local level and redesigning of the alert interface could be beneficial. Create a CDSS for making decisions about drug - drug interaction is a complex process that requires supportive evidence, structured databases, good interface design and trained staff to adapt the evidence to the healthcare context of a health institution.

Keywords. Clinical decision support system, Drug-drug interaction system, Electronic health record

1. Introduction

Decrease medical errors, improve health processes and ensure high quality health care for patients has been the focus of constant concern of all members of the health care team. In this context arise computerized support systems (CDSS) in response to the need to improve the process of clinical care (1). The CDSS in general and those related to prescription systems in particular have the potential to reduce and prevent morbidity and mortality associated with adverse events and improve the quality of patient care (2). Most electronic health records (EHR) has CDSS used as support for clinical decision making. In the context of electronic prescribing systems, knowledge databases necessary for the implementation of systems alerting for drug - drug interactions use commercial knowledge databases, usually created in the United States and in English language. These databases are not adapted to local contexts of use, generating a high rate of false positives and producing "alert fatigue" (3–5), situations in which the user, after receiving numerous warnings, with no real clinical impact, ignore and / or dismiss this advise, even though in some cases have medical relevance. To solve this problem, the databases of the CDSS should take into account the context of clinical use, the health system and the clinical evidence.

Although the benefits of CDSS are known, it is not uncommon to find reports indicating a high rate of omission of these alerts on grounds ranging, as previously discussed, since the inadequate content of the knowledge database of such systems and the lack of clinical significance of the recommendations to the poor design of human-computer interfaces (6,7) therefore, all references to improve these aspects result in improved safety for patients (8).

Interactions among drugs administered to a patient, occur when a drug causes changes in metabolism of the other, a phenomenon known as drug-drug interactions (DDI) (9). These DDI can cause unwanted adverse events in the patient and the severity of symptoms can vary from

negligible to potentially lethal. The occurrence of DDI is associated with increased morbidity and mortality (10,11), with prolonged hospitalization (12) and high health care costs (6). The increasing use of new pharmacological agents (7), the clinical context of the patient (8) increasingly complex and other factors such as prolonged hospitalization, makes identifying these DDI is beneficial and at the same time increasingly difficult (13).

El Hospital Italiano de Buenos Aires (HIBA) developed and implemented an electronic health record (EHR) with a notified drug interactions system (14), but in first instance the rate of cancellation of alerts presented was high. The decision was then to redesign the system components for notify drug interactions, consistent with the recommendations in the literature and using techniques of user centered design (UCD). First we worked on the first phase of debugging the knowledge database and the categorization of its recommendations (4) and an analysis of the cases was conducted to determine the local occurrence and to then, move forward in redesigning alerts with techniques based on user-centered design (UCD).

The redesigned of the notification system for drug interactions was organized in the following steps:

- Structuring a knowledge database of drugs
- Modification and validation of a knowledge database on drug - drug interactions
- Case studies and frequency of interactions at the local level
- Redesign of the alert interface

2. Setting

El Hospital Italiano de Buenos Aires (HIBA) is a university hospital of high complexity founded in 1853 belongs to a nonprofit health network including a second hospital, 25 outpatient centers and 150 private clinics distributed in the city of Buenos Aires. The infrastructure is complete with 750 inpatient beds, 200 of which are for critical care, a home care service and 41 operating rooms. A team of 2800 doctors, 3000 agents of the health team and 1900 persons for administrative tasks and management process work at the hospital. Approximately 45,000 discharges per year, 3 million annual visits and 45,000 surgical procedures were performed. Since 1998 has been gradually implemented a Health Information System (HIS) development "in house" that handles the medical and administrative information from capture to analysis. It includes a unique problem-oriented and patient-centered health record, known by the name of ITALICA (14). EHR allows documentation of care in areas including: outpatient, inpatient, emergency and home care. ITALICA allows the request of complementary studies, drug prescriptions and results display that includes a PACS (Picture archiving and communication system). Since the implementation of the EHR, a Terminology Server for vocabulary representation was created. The Terminology Server allows linking free text entered by the health team in the EHR and references it with SNOMED CT, plus the ability to associate them with different classifications, such as ICD-9-CM, ICD10, ICPC, LOINC, among others. In 2006 the HIBA start the Personal Health Portal Project (PHR). The PHR is linked to ITALICA, which provides services and unified access to multiple data applications, allowing the patients see their health data stored in the health network, and allow them to interact or consult their medical or administrative information.

3. Redesign process

3.1. Structuring a knowledge database of drugs

In order to integrate scientific knowledge with the clinical decision support systems the hospital decided to start its e-prescribing project, progress on both the development of its own CPOE, and will develop a structured knowledge database on drugs and serve substrate for the creation of support systems. From this decision created a group comprised of physicians, pharmacists, and students of medicine and pharmacology that dealt with several databases of national and international information and created a data structure that allow the storage of structured and coded information. The uploaded data are referenced to SNOMED CT, the same standard terminology that the health information contained in patient clinical data repository.

This process enables the creation of CDSS, and the use of international standards facilitates the decision system creation process, EHR integration and implementation.

The knowledge database of drugs, in addition to basic drug information contains also commercial information like products' names and presentations, and also, has information on the interaction of the drug with other active ingredients.

This database that fed in a first instance to the notification system of interactions was based on international knowledge database, so its lack of contextualization to local realities generated false positives causing alert fatigue. It was then decided to modify and validate this knowledge database on DDI and also make taxonomy of recommendations for actions to be taken by professionals, with the aim of providing the warning message to perform a specific action.

3.2. Modification and validation of a knowledge database on drug - drug interactions

In the electronic prescribing module the Notify System for Drug Interactions, worked at the beginning using a commercial database, Evaluations of Drug Interactions (EDI) (4,15) of the company First Data Bank. The database uses the terminology of Table 1 for the categorization of risk of each interaction.

Table 1 Terminology used in Evaluations of Drug Interactions (EDI) - First Data Bank

Code	Category	Description
I	<i>High clinical significance</i>	DDI that have great potential to harm the patient, are predictable or occur frequently and are well documented
II	<i>Moderate clinical significance</i>	DDI who have a moderate potential to harm the patient, are less predictable or occur infrequently, or lack of complete documentation
III	<i>Low clinical significance</i>	DDI who have a low potential to harm the patient, have a varied predictability or occur infrequently,

		or has poor documentation
IV	<i>No clinical significance</i>	Although these DDI can occur, the documentation is based on unproven theoretical foundations or effect resulting from the interaction is not clinically significant and / or any adverse event is not expected

To avoid false positives and alert fatigue evidenced in the literature we asked pharmacologists at the Clinical Pharmacology Section of HIBA to review and validate the DDI that are the basis of system.

Work were organized in two phases: In the first, we worked with DDI level I of EDI, which were evaluated between two pharmacologists physicians, determining each other, which ones, according to the literature, the health care setting and prevalence, were relevant. In the second phase, the remaining levels (II, III and IV) were distributed between two other pharmacologists. In the absence of agreement on an interaction, the level was assigned by the intervention of a third pharmacologist to arrive at a consensus.

We decided then to re-classify interactions using a standardized layering system of interactions with the knowledge database Lexicomp® (16). For the re-classification of the clinical significance of DDI by severity, likelihood of occurrence in our environment and level of care (outpatient, critical or noncritical hospitalization), two pharmacologists evaluated according to clinical criteria, scientific evidence in the literature and using standardized tools(Lexicomp®) (Table 2) each of the interactions after the first purification step and assigned the corresponding risk. If there is a disparity between the two, a third pharmacologist evaluated the same and defined.

Table 2. Lexicomp® terminology.

Severity	Action	Description
A	<i>Unknown interaction</i>	The data have not demonstrated pharmacodynamic or pharmacokinetic interactions between agents
B	<i>Not action required</i>	The data demonstrate that the specified agents may interact with each other, but there is little or no evidence of clinical interest resulting from concomitant use
C	<i>Monitoring</i>	The data demonstrate that the specified agents may interact with each other in a clinically meaningful way. The benefits of concomitant use of these two medications usually outweigh the risks. An appropriate monitoring plan should be implemented to identify potential negative effects. Adjustments may be necessary dose of one or both agents in

		a minority of patients
D	<i>Consider modification of treatment</i>	The data demonstrate that the two drugs can interact with others in a clinically significant way. A specific patient assessment should be conducted to determine whether the benefits of the combined therapy outweigh the risks. Specific actions to be taken in order to obtain the benefits and / or minimize the toxicity resulting from concurrent use of agents. These actions may include aggressive monitoring, empirical dose changes, the choice of alternative agents
X	<i>Avoid combination</i>	The data demonstrate that the specified agents may interact with others in a clinically significant way. The risks associated with the concomitant use of these agents usually outweigh the benefits. These agents are generally considered contraindicated.

So the categories I, II, III and IV of the first stage be re analyzed and reclassified using the nomenclature A, B, C, D and X of Lexicomp®.

As a first step before starting the analysis of interactions, DDI drugs not available in Argentina were eliminated. So from the total of 4148 DDI (768 level I, 1736 level II, 1591 level III and 53 level IV) were discarded in the first instance 381 cases, since this drugs are not available or authorized for use in the country.

For the remaining 3767 pairs of drugs, that generated some interaction, the evaluation task was performed by two clinical pharmacologists. Of the 3767 interactions evaluated the degree of agreement between the two observers was very good. However in cases where there was disagreement, a third pharmacologist evaluated all cases where there had been no match for defining risk. The final database included 3767 DDI re-categorized as level I (437), level II (1713), level III (1563), level IV (54).

In the second stage, the DDI classified as Level I and II (2150) were reclassified to level X or D 683 (548 X and 135 D). This process bases debug and modify the value of risk of several pairs of drugs, which meant a reduction of 69% of interactions with important clinical significance. For this stage the strength of agreement was also good, and discrepancies were defined by a third pharmacist.

Validation of pharmacological databases for use in a CDSS integrated into an EHR from commercial databases is a complex but necessary process. Adapting to health care settings and the local reality of the health system will impact the alerts that interrupt clinical workflow and acceptance of electronic prescribing alerts. The 69% reduction in the number of possible DDI would alert a possible interaction during the medical procedure not only impacts the quality of alerts, but on the quality of information for the physician and patient. At the same time reduces the possibility of false positives with consequent alert fatigue.

In this paper the concordance between observers was very good, which adds an extra value to the validation process of the knowledge databases. The way on observers classified the DDI was based on the pharmacological and pharmacokinetic potential changes of DDI, and accordingly the clinical consequences. The database Lexicomp® better fits this model, so we took the decision to use it for the award of levels D and X.

The lack of primary literature to evaluate commercial databases hinders its applicability in a particular scenario. The DDI databases vary by country and its laws, the users, the characteristics of the patient population, funding and technology infrastructure.

3.3. Modification and validation of a knowledge database on drug - drug interactions

In order to select cases of real and representative of the different levels of care in our network, was carried out a work to analyze the frequency of potential drug-drug interactions in pharmacological indications of the different levels of care and analyzed the characteristics of the most clinically significant. The study population in order to cover the widest possible spectrum of care processes implied two care models of drug prescriptions, episodic and longitudinal (outpatient). The episodic model looked pharmacological requirements of the care received by patients in both the central emergency and inpatient episodes; the instructions given in this care model are updated by the attending physician daily. The longitudinal model considered pharmacological prescriptions made by professionals in outpatient visits, these particulars are updated as needed in the patient's medication list in the EHR.

In both models of care, when a new drug indication was registered by a professional in the EHR, new drugs and preexisting indicated were combined and then contrasted with the knowledge base in the system in order to reach for potential DDI. The pair combinations of drugs that showed a positive result, after being processed with the System were considered for analysis. Both the longitudinal model and the episodic one was recorded only the first occurrence of the pair of drugs that generated a potential DDI, discarding subsequent repetitions during the same episode.

Data Analysis: A descriptive cross-sectional study in which two types of units of analysis were considered was performed: the only positive combinations, by patient or episode, in the year of study, in order to determine their impact; and on the other hand, patients who met the inclusion criteria in order to know how many of them could potentially suffer at least one DDI. Data are presented as mean and standard deviation (SD). In order to obtain the frequency of occurrence of DDI by the severity of the clinical cases seen in the episodic care model, the requirements were grouped into critical episodic (those made in critical care units in hospital and emergency center) and episodic not critical (the rest). The characteristics of the DDI were analyzed according to the amount of co-prescription drugs, their potential clinical significance (severity of the interaction) and taxonomy of recommendations (actions) contained in the knowledge database. Of these groups the serious severity selected and recommending an action to avoid the absolute torque joint use drug, and that this subset includes active and intrusive process alerts prescriptive. In this last subgroup of the prescribing professional specialties analyzed and the routes of administration.

Between March 2011 and February 2012 on the episodic model of care drug prescriptions, 1,587,167 drug combinations were generated, showing 19,162 potential DDI. From a total of

36,013 patients treated in this model of care, 6,338 had at least one DDI, with an annual incidence of 17.6 per 100 patients with DDI patients seen. While in ambulatory model drug combinations were 4,084,043 observed potential DDI in 29,251 opportunities. From a total of 55,119 patients treated in the longitudinal model, 15,267 had at least one DDI, with an annual incidence of 27.7 per 100 patients with DDI patients.

Once made in analysis of frequency of occurrence of potential DDI model episodic attention was divided into episodic critics (those made in the critical care units of hospital and emergency center) and episodic non-critical (the rest) in order to obtain different combinations grouped by positive indirect appearance in the complexity of the clinical cases. First the amount of drugs administered simultaneously (co-prescribed) when generating the DDI, divided by the different levels of care, showing a maximum of 5 to 10 concurrent drugs prescribed, mean were analyzed in the critical sectors of 11.76 (SD 4.70), in non-critical of 11.68 (SD 4.63), and outpatient 11.19 (SD 8.51).

In a previous study, the first phase of debugging the knowledge database for the Interactions System contemplated removing multiple interactions considered irrelevant, by classifying the remaining clinical significance (severity) in beneficial, mild, moderate and severe interactions. In this study, positive combinations tested were grouped according to this classification. Of the 48,413 total observed potential DDI, 3,180 were severe (6.6%), of which 167 were in the episodes grouped as critical, 1,694 in non critical and 1,319 in the outpatient setting.

In previous work, taxonomy of recommendations for actions to be taken by the practitioner, in order to deliver the message of the alert a specific action to perform will also drew up and that acceptance of that action is accounted for as a cancellation alert. In applying this taxonomy to the knowledge database, each DDI could have more than one recommendation. The recommendation "Avoid combination" was the only one considered in the EHR as an active and intrusive (regardless of severity) alert. As with severity, in this paper, the DDI potential observed were grouped according to the model of care and according to the above taxonomy, being 301 potential DDI avoid combination, not positive combinations in critical incidents, 23 in non-critical and 278 in the outpatient setting.

In order to create representative clinical cases of our clinical reality for the user-centered redesign of alerts, and to focus the second treatment stage in our database for interactions, we select the most frequent potential DDI in our subgroup represented by the DDI of "high" severity (no matter what action or actions recommended in the taxonomy) and those recommending action "Avoid combination" (no matter how severe they had). Subsequently those with both attributes were considered only once (subtracted from the total) for the net number of potential DDI that would have been presented to professionals and active and intrusive alerts. In this subgroup analysis of a net total of 3,356 positive combinations, 167 episodes were classified as critical, 1,709 as non critic and 1,480 in the outpatient setting were found.

When the characteristics of this subgroup were analyzed by specialty doctor who performed the indication can see that cardiology and internal medicine are the specialties most exposed to positive combinations. Because the route of drug administration is a consideration to avoid false positive alerts we were analyzed which route was frequently prescribed in the potentially active alerts, resulting intravenous and oral routes the most frequently used.

3.4. Redesign of the alert interface

For the redesigned interfaces alert, a team consisting of 3 doctors and 2 computer usability specialists work in the process of designing user-centered interfaces for the development of a new support system for warning drug - drug interactions at the time of prescription. The project was raised in iterative steps.

In total, in the three stages, 24 doctors participated (Table 1).

Data	Value
Sex	male 58%
Age (mean in years)	33
Years of graduated (mean)	9
Years of experience using the EHR (mean)	7
Setting	outpatient 37%
	inpatient 25% (20% non intensive care - 5% intensive care)
	Both 38%

TABLA 1: Users profile

STEP 1: INDAGATION

This first stage consisted of observations and contextual interviews to physicians who perform both electronic prescriptions in the outpatient setting and inpatient areas. In these interviews, semi structured questionnaires dealing topics and scenarios concerning the processes of drug prescribing in situations where drug interactions are presented were performed. From the analysis of these interviews were generating low-fidelity prototypes, in order to confront the ideas that emerged from these interviews with the design of interfaces for interaction with medical drug – drug interaction alerts.

Profile of participants: at this stage physicians working in ambulatory care, inpatient in critical and not critical areas participated in the study. The minimum required experience was 4 years of use of our EHR, trying to select participants with experience in the prescription process that could convey their knowledge in the field (know how). At this stage a semi-structured questionnaire to guide the user towards dynamic test cases, basing these on real clinical cases were created using the case study as the incidence of possible interactions in relation to the history of prescriptions presented (17).

6 doctors participated in the first stage, the interviews in this stage were the basis for structuring the prototypes. Emerges from these interviews that the DDI is a common concern among physicians and the tools that support these issues would be useful in this instance. From this concept, several participants took pertinent comments that have agreement with those published in the literature, such as the difficulty of finding good resources for detecting interactions, poor interface design, the amount false positive, and alert fatigue that this causes. We took what was expressed by users referring to past experiences to the structuring of the prototypes that will be used in the following steps. They also expressed the need for convenience or alerts provide information and guidance concerning the conduct to follow.

STEP 2: PARTICIPATORY DESIGN:

From the previous stage and on the basis of low-fidelity prototypes stage participatory design was performed, where the participating physicians, through their opinions arising from the interaction with the first prototypes were guiding the development of a new prototype closest to the user feedback. Were performed in this instance 2 cycles of prototyping and testing. At this stage the focus was qualitative, seeking saturation domain. The opinions and thoughts of physicians were obtained and recorded from the Think Aloud technique (18). The experiments were recorded on a mobile usability lab, and tests were carried out mainly in the workplace participants. Participants in this stage different users of the above but with a similar profile. The script used in the test was also based on real cases. At this stage, two prototypes were used, the 1st made in Balsamiq (19) was printed on paper and showed to the users, the 2nd done in Balsamiq was exported and used in order to reproduce the prescription process as closely as possible.

At this stage, a different medical group as above, from participatory design techniques, generated, along with the team carrying out this study, a design of the interfaces to be used when facing a DDI alert.

The results at this stage were analyzed by non-stringent qualitative techniques, in order to analyze the results and semi structures the results. The opinions given at this stage were: alerts were generally well received, were highly valued the ability to take actions from the same alert, without interrupting the workflow and restart the prescription process. Recommended actions that are integrated as operations were received as a great advantage. On this workflow, this was the most complex to elaborate on the design of interfaces. Among the negatives, navigating the various options offered by an alert, were complex. These results were considered for the redesign and re testing of prototypes, until they were considered adequate by users. Comments from users about the relevance of an alert appears or were not assessed and taken into account for debugging knowledge databases and developing test cases. Perhaps the main result of this stage was to understand, from the views of users, the interface should be action-oriented, and not just information. In other words, the user wanted to see not only the data, but perform the recommended action from the same screen.

STEP 3: USABILITY TEST

High fidelity prototype was created in Axure (20) for testing usability, which was presented as a functional prototype to different users of the participants in previous stages. In turn, this instance is used for measuring EFFECTIVENESS (understood from two variables, the first referring to the course pursued, i.e., if they ignore the warning or take into account, being the effective design if the alert is taken into account and the recommendations are followed. The other endpoint is whether they can complete the process of prescribing) and SATISFACTION (which was assessed using a questionnaire SUS -. system usability scale (21)).

EFFECTIVENESS: the first variable, understood as the course pursued, i.e., if they ignore the warning or take into account, showed the following results: of the 24 participating physicians, 11 physicians ignored the warning (45.8%) and 13 (54.2%) agreed.

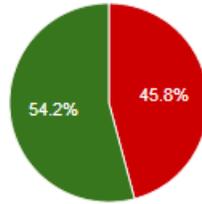


Figure 2: effectiveness of the alert acceptance

The other measure of effectiveness is whether it can finish the prescription process, in which the following results were evident: 13 (52.17%) physicians completed the task without difficulty, 10 (43.48%) completed the activity with questions, 1 (4.35%) with some errors and there were no serious errors that can't allow to complete the action.

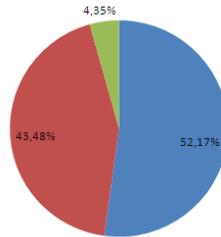


Figure 3: effectiveness of alert process

The number of errors was also measured during the process and are shown below: no error 82,6%, small mistakes 13%, severe mistake 4,4%

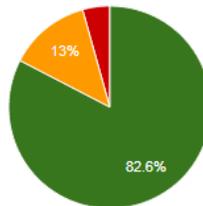


Figure 4: Errors

SATISFACTION, which was evaluated from a SUS questionnaire, yielded an average value per participant of 77.90 on a scale of 0 to 100 that indicates a value more than adequate, this equals a percentile of 83%, which means in other words, that the acceptance rate of the system is above 80%.

4. Discussion

Structure, validate and adapt the knowledge databases that served as a substrate for the development of this system resulted extremely difficult, in the same way, the creation of clinical cases also resulted complex, as they refer to rare problems that doctors do not handle usually,

what is presented as one of the challenges for the design framework, in this sense, the use of DDI alert systems is not easy for training because due to their low prevalence, it may be a long time from training until the appearance of the case, so the challenge is that these interfaces should be as intuitive as possible (having a short or fast learning curved), then there lies the benefit of investing time and resources in the participatory design of this type of support systems.

As for the interface design process, the parameters in the studies were more than satisfactory. From the detailed analysis of the metrics considered in the later stages, both the values of effectiveness, and satisfaction were improved to achieve the final prototype that will be used in the last stage. With respect to satisfaction, it is interesting to note that in comparison with the literature, where these types of systems have yielded no positive values, percentile values greater than 80% shows the importance of involving users in the design of the system. With regard to the acceptance of the recommendations, although a percentage above 50% is generally low impact, this performance is more than acceptable in this particular topic, where similar studies have thrown alerts omission rate much higher.

Getting doctors to participate in the test was not easy, however the number of participants was appropriate as many users as needed were interviewed to obtain saturation.

Future lines: the final stage will be held by the final prototype, with which laboratory test will be performed from cases generated from the actual incidence based on the frequency of cases recovered in pre analysis.

Limitations: This study was conducted at a single center, with doctors trained in the same institution, so that the processes for incorporating external validation must be done elsewhere. On the other hand, the correlation between the clinical case and physician specialty was not always desired by both, some physicians may face situations not common practice, so the behavior is not taken from the knowledge of field, but influence by the alert itself, which tends to select "suspend" rather than harm.

5. Conclusion

Create a CDSS for making decisions about drug - drug interaction is a complex process that requires supportive evidence, structured databases, good interface design and trained staff to adapt the evidence to the healthcare context of a health institution.

6. Bibliography

1. Greenes RA. Clinical Decision Support: The Road Ahead. Academic Press; 2011. 600 p.
2. Bates DW, Kuperman GJ, Wang S, Gandhi T, Kittler A, Volk L, et al. Ten Commandments for Effective Clinical Decision Support: Making the Practice of Evidence-based Medicine a Reality. *J Am Med Informatics Assoc JAMIA*. 2003;10(6):523–30.
3. Van der Sijs H, Aarts J, Vulto A, Berg M. Overriding of drug safety alerts in computerized physician order entry. *J Am Med Informatics Assoc JAMIA*. 2006 Apr;13(2):138–47.
4. Luna D, Otero V, Canosa D, Montenegro S, Otero P, de Quirós FGB. Analysis and redesign of a knowledge database for a drug-drug interactions alert system. *Stud Health Technol Inform*. 2007;129(Pt 2):885–9.

5. Luna, D., Gonzalez Bernaldo de Quiros, F., Martinez, M., Gambarte, M. L., Morchón, A. ¿Porque se ignoran los alertas sobre interacciones farmacológicas? Buenos Aires; 2004.
6. Bates DW, Spell N, Cullen DJ, Burdick E, Laird N, Petersen LA, et al. The costs of adverse drug events in hospitalized patients. Adverse Drug Events Prevention Study Group. *JAMA J Am Med Assoc.* 1997 Jan 22;277(4):307–11.
7. Seymour RM, Routledge PA. Important drug-drug interactions in the elderly. *Drugs Aging.* 1998 Jun;12(6):485–94.
8. Hartholt KA, van der Velde N, Looman CWN, Panneman MJM, van Beeck EF, Patka P, et al. Adverse Drug Reactions Related Hospital Admissions in Persons Aged 60 Years and over, The Netherlands, 1981–2007: Less Rapid Increase, Different Drugs. *PLoS ONE.* 2010 Nov 12;5(11):e13977.
9. Kuperman GJ, Bobb A, Payne TH, Avery AJ, Gandhi TK, Burns G, et al. Medication-related clinical decision support in computerized provider order entry systems: a review. *J Am Med Informatics Assoc JAMIA.* 2007 Feb;14(1):29–40.
10. Phillips DP, Christenfeld N, Glynn LM. Increase in US medication-error deaths between 1983 and 1993. *Lancet.* 1998 Feb 28;351(9103):643–4.
11. Classen DC, Pestotnik SL, Evans RS, Lloyd JF, Burke JP. Adverse drug events in hospitalized patients. Excess length of stay, extra costs, and attributable mortality. *JAMA J Am Med Assoc.* 1997 Jan 22;277(4):301–6.
12. Fanikos J, Cina JL, Baroletti S, Fiumara K, Matta L, Goldhaber SZ. Adverse Drug Events in Hospitalized Cardiac Patients. *Am J Cardiol.* 2007 Nov 1;100(9):1465–9.
13. Glassman, P. A., Simon, B., Belperio, P., & Lanto, A. Improving recognition of drug interactions: benefits and barriers to using automated drug alerts. *Med Care.* 2002;40(12):1161–71.
14. Fernán González Bernaldo de Quirós, Daniel Luna, Analía Baum, Fernando Plazzotta, Carlos Otero, Sonia Benítez. Incorporación de tecnologías de la información y de las comunicaciones en el Hospital Italiano de Buenos Aires. CEPAL;
15. First Data Bank - Evaluation of Drug Interactions | EDI [Internet]. Available from: <http://www.fdbhealth.com/fdb-medknowledge-clinical-modules/drug-drug-interaction/>
16. Lexicomp [Internet]. Available from: <http://www.lexi.com/>
17. Daniel Luna MD. MSc cand. PhD cand, Martín Waldhorn MD, , Lic. Adriana Ianicelli, Marcelo Risk Ing. MBA MSc PhD3, y Fernán González Bernaldo de Quirós MD MSc FACMI. Drug-drug interaction in a healthcare network of Argentina: frequency and characteristics. Brasil; 2012.
18. Jaspers MWM, Steen T, van den Bos C, Geenen M. The think aloud method: a guide to user interface design. *Int J Med Inf.* 2004 Nov;73(11-12):781–95.
19. Balsamic. Available from: <https://balsamiq.com/>
20. Axure. Available from: <http://www.axure.com/>
21. P. W. Jordan, B. Thomas, B. A. Weerdmeester, & A. L. McClelland. SUS: a “quick and dirty” usability scale. SUS: a “quick and dirty” usability scale.

Contact

Dr Carlos Otero
 Hospital Italiano de Buenos Aires, Argentina
carlos.otero@hospitalitaliano.org.ar

Towards a foundational representation of potential drug-drug interaction knowledge

Mathias Brochhausen¹, Jodi Schneider², Daniel Malone³, Philip E. Empey⁴, William R. Hogan⁵ and Richard D. Boyce⁶

¹ Division of Biomedical Informatics, University of Arkansas for Medical Sciences, 4301 W. Markham St., #782, Little Rock, AR, 72205-7199

²Web-instrumented man-machine interactions, communities and semantics group, INRIA Sophia Antipolis – Méditerranée, France

³College of Pharmacy, University of Arizona, Tucson, AZ, USA

⁴Department of Pharmacy and Therapeutics, School of Pharmacy, University of Pittsburgh, Pittsburgh, PA, USA

⁵Department of Health Outcomes and Policy, College of Medicine, University of Florida, Gainesville, FL, USA

⁶Department of Biomedical Informatics, University of Pittsburgh, Pittsburgh, USA

mbrochhausen@uams.edu, jschneider@pobox.com,
malone@pharmacy.arizona.edu, pempey@pitt.edu, hoganwr@ufl.edu,
rdb20@pitt.edu

Abstract. Inadequate representation of evidence and knowledge about potential drug-drug interactions is a major factor underlying disagreements among sources of drug information that are used by clinicians. In this paper we describe the initial steps toward developing a foundational domain representation that allows tracing the evidence underlying potential drug-drug interaction knowledge. The new representation includes biological and biomedical entities represented in existing ontologies and terminologies to foster integration of data from relevant fields such as physiology, anatomy, and laboratory sciences.

Keywords: Potential drug-drug interactions, ontologies, knowledge management

1 Introduction

Every year, many thousands of people are harmed by exposure to two or more drugs for which there exists a known interaction potential. Exposure to such “potential drug-drug interactions” (PDDIs), are a significant source of *preventable* drug-related harm, leading to clinically important events in 5.3% - 14.3% of inpatients, and accounting for 0.02% to 0.17% of the 129 million emergency department visits that occur in the U.S. each year[1][2]. Multiple defenses exist in the healthcare system to prevent patient harm from PDDIs including clinician knowledge, computer screening,

and monitoring. Each defense depends on complete, accurate, and current knowledge of what drugs have the potential to interact, and the most appropriate methods for managing patients when exposure to a PDDI is unavoidable [3]. However, most sources of clinically-oriented PDDI knowledge disagree substantially in their content, including about which drug combinations should never be never co-administered. For example, only one quarter of 59 contraindicated drug pairs were listed in three PDDI information sources[4], only 18 (28%) of 64 pharmacy information and clinical decisions support systems correctly identified 13 PDDIs considered clinically significant by a team of drug interaction experts[5], and four clinically oriented drug information compendia agreed on only 2.2% of 406 PDDIs considered to be “major” by at least one source[6].

A key factor underlying the existing disagreements among sources of drug information that are used by clinicians is the inadequate representation of PDDI evidence and knowledge. In practice, organizations that provide PDDI information as part of their information services employ an expert or panels of experts (editorial boards) to search, evaluate, synthesize, and stay current with evidence. The process involves applying some criteria to judge whether a drug combination could lead to an interaction, what impact it might have on exposed patients, and how to best manage patient exposure. In the current paradigm, these individuals or groups must search across multiple information sources, including the scientific literature, drug product labeling, and documents submitted to regulatory groups during the drug development/approval process. There is significant variation across drug knowledge bases with respect to ratings of specific drug pairs and currency. Moreover, the available sources rarely include first-hand clinical experience, information that can help contextualize management recommendations[7]. In addition, those multiple sources are currently not created in a way that fosters semantic integration of their data at a later stage. This leads to inefficient and discordant approaches to the acquisition of PDDI evidence and synthesis of that evidence into knowledge. The result is that there is general disagreement among drug information systems about what PDDI exist and their clinical importance.

A goal of the “Addressing gaps in clinically useful evidence on drug-drug interactions” project is to identify the core components of a new PDDI knowledge representation paradigm that addresses these issues. As we describe below, the project makes a fundamental distinction between assertions of PDDI knowledge and the evidence that supports or refutes such assertions. The central thesis of the project is that a framework for representing PDDI assertions and evidence as interoperable Linked Data[8] will enable a more integrated approach to the acquisition and synthesis of PDDI evidence into knowledge. Linked Data methodologies should be used to semantically integrate the various relevant sources of PDDI evidence so that experts can more easily retrieve all relevant evidence items. This will lead to more complete, accurate, and current PDDI information provision to any single evidence board than is possible with current resources.

The proposed framework requires a new foundational representation of PDDIs that covers the material entities and processes in the domain of discourse for PDDI evidence and knowledge claims. The representation will enable the integration of drug

interaction mechanisms, effects, risk factors, severity, and management options with the chemical and pharmacological properties (e.g., chemical structure, function, pharmacokinetic and pharmacodynamic properties) of the interacting drugs. This paper specifies the design requirements for such a foundational representation that we are calling the Drug-drug Interaction and Drug-drug Interaction Evidence Ontology (DIDEO). Section 2 provides clinical background. Section 3 discusses the basic design principles and decisions for the new ontology. Finally in section 4, we show that the classes in DIDEO are sufficient to represent a concrete example of PDDI evidence selected from the Drug Interaction Knowledge Base.

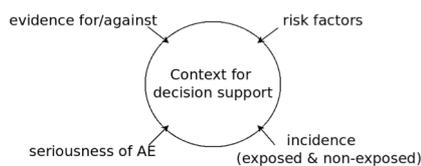


Fig. 1. The four types of information used by van Roon et al. to determine if a PDDI warrants clinical action.

AE – Adverse Event

2 Background

2.1 A conceptual framework for clinically useful PDDI Knowledge and Evidence

There is a rather complex relationship between the evidence that establishes a PDDI, and information that can help clinicians accurately assess the risk of exposure within a given patient[7]. The foundational model we envision would benefit from an explicit conceptual model of that relationship. Eric van Roon *et al.* proposed a conceptual model of PDDI information using the definition that clinically-useful PDDI information is that which helps discern whether some action should be taken with respect to a PDDI (Figure 1)[9]. Evidence for, or against, the existence of a PDDI is an important component in that model, along with consideration of patient risk factors, the potential severity of an adverse event that could be caused by exposure, and prior experience with exposure in relevant patient populations. While the van Roon model is not considered a standard for representing PDDI knowledge, it captures the essence of recommendations by other PDDI experts[10][11], including developers of PDDI databases in the United States and Europe[12][13].

The van Roon model helps to conceptually outline the principal information domains for clinically-useful PDDI knowledge. We think that it is also important to consider the relationship between PDDI evidence and claims of PDDI knowledge established by evidence. The evidence for, or against, PDDI assertions is dynamic and of varying robustness to various forms of bias. For example, in prior work on the Drug Interaction

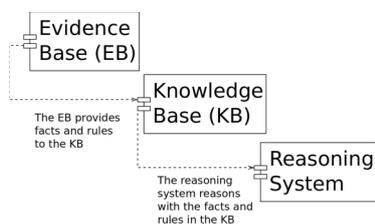


Fig. 2. Distinctions between evidence, knowledge, and reasoning with respect to PDDIs

Knowledge Base (DIKB) [14][15][16][17], the editorial board considered certain pharmacologic assertions written in a FDA guidance to industry[18][19][20] useful as PDDI evidence. The assertions reflected the state of science at the time the documents were published. The guidance has been updated, from 1999 to 2006 to 2012, each update leading to changes in the DIKB evidence base.

Based on these observations, it's possible to conceive of PDDI evidence board as sociotechnical *reasoning system* that manages both an *evidence base* and a *knowledge base* (Figure 2)[21]. This implies that the foundational PDDI knowledge representation we envision would find application within three specific contexts:

- 1) within the *evidence base*, by defining the four types of information from van Roon's model (Figure 1): evidence that can be used to establish the existence of a PDDI, patient risk factors, the potential severity of an adverse event that could be caused by exposure, and prior experience with exposure in relevant patient populations;
- 2) within the *knowledge base*, by representing the entities explicit within PDDI assertions (e.g., "*drug X* interacts with *drug Y*") and pharmacologic assertions that can be used to infer PDDIs (e.g., "*drug X* inhibits *enzyme Q* which is important for the clearance of *drug Y* from the body"); and
- 3) within the *reasoning system*, by constraining the inference activities of the evidence board so that inferred knowledge is logically consistent with all of the other assertions in the knowledge base.

These distinctions can be illustrated using the artifacts used to support the DIKB. Underlying the DIKB's *evidence base* are specific examples of evidence types; these evidence types are outlined in a draft online document[22]. Meanwhile the DIKB's *knowledge base* contains PDDI and pharmacologic assertions; these assertions address competency questions that were identified during prior work on the DIKB, which can be found in a different draft online document[23].

2.2 Related work

Currently, there are two ontologies built specifically for the domain of drug-drug interactions: the Drug Interaction Ontology (DIO)[24] and the Drug-drug Interaction Ontology (DINTO)[25]. Both provide insights that are valuable for representing the domain. However, neither was designed with the perspective outlined above. Nor do they allow for a consistent and scalable representation of the ontological distinctions relevant to representing clinically useful drug-drug interaction assertions, the drugs involved, and the supporting or refuting evidence. We introduce these ontological distinctions in the course of discussing the existing ontologies.

The DIO is an ontology of drug interactions developed with the goal of predicting drug interactions[24]. While DIO[24] is inspired by both Basic Formal Ontology (BFO) and the NCI Thesaurus (via UMLS), it is not aligned with either one. Although the DIO [24] specifically refers to BFO's distinction between continuants and processes (occurrences)[26], the BFO's representation of process (its definition and entity URI), is not reused in the OWL implementation of DIO accompanying the aforementioned paper[27]. Rather, within the DIO OWL file, process is defined as: "A se-

quence of events which produces some outcome” [27], reusing the CUI and definition from the National Cancer Institute Thesaurus (NCIT)[28]. But the DIO OWL file includes several axioms that are inconsistent with both the NCIT definition of process and the BFO’s representation of process. In particular, DIO specifies necessary conditions for processes, such as

- `hasEnableTriggerParticipant min 1 Thing`
- `hasResultantPopulationChange min 1 Thing`
- `hasResultantPopulationChange only Increased`

A possible explanation for this inconsistency is that the developers of DIO intended to map their domain representation to those found in the UMLS (NCIT) and other terminologies and ontologies rather than actually reuse them. However, the reuse of terms should always be accompanied by ensuring that the intended meaning and the ontological commitments of source and target resource match. (We will elaborate on the role of ontological commitments in Section 3.)

Another shortcoming of DIO is that it does not represent roles. Each instance of a chemical is a drug, regardless of whether its dosage or formulation allows it to act as a drug. We find, for instance, that:

```
Capecitabine
  rdfs:subClassOf Drugs
  rdfs:subClassOf DrugOrMetabolite
  rdfs:subClassOf Chemicals
```

But active ingredients can only bear a role as a drug in a specific dose and in conjunction with excipients[10]. However, it is by now standard accepted practice in numerous drug terminologies and ontologies to carefully distinguish among drug products, their ingredients, and the molecules that constitute those ingredients[29][30][31][32]. Most recently, Hogan et al. in this regard showed that assigning therapeutic properties to active ingredients disregards the effect of dose form and therefore leads to mistakes that contradict scientific knowledge (e.g., oral vancomycin treating bacterial endocarditis)[30].

The second ontology that we took into consideration is the Drug Interaction Ontology (DINTO). DINTO is intended “to represent all possible mechanisms that can lead to a drug-drug interaction. The ontology provides the general pharmacological principles of the domain”[25]. The developers have provided a version of DINTO that is an extension of BFO[33]. The key limitation of DINTO with respect to our goals is that DINTO does not represent potential drug-drug interactions at all, but only drug-drug interactions (DDIs). Representing PDDIs, as we aim to do, is quite different from representing a DDI. For each individual instance of a drug-drug interaction it is possible to specify the individual patient who suffered from its effects. However, this is not possible for all instances of PDDIs, because some of them are not actualized. DIDEO will be based on a novel definition of PDDI (Section 3).

While DINTO does not represent PDDIs at all, the way it represents the actual occurrences of DDIs and information about those occurrences is problematic. DINTO specifies a subclass of DDIs named *DDI described in a database*. The members of

this class are intended to be DDIs¹ that are linked to a database by the *is described in*-relation. According to the DINTO OWL file[33] the class ‘DDI in database’ is intended to “represent those DDIs imported in DINTO from the DrugBank database[34][35] with the purpose of distinguishing them from those inferred from the ontology”. Notably, the information loaded from DrugBank will not be about drug-drug interactions, but about PDDIs, as [35] clearly indicates. This demonstrates that DINTO does not provide ways to distinguish information about actual DDIs from evidence pointing at drug co-medications that are suspected to lead to unwanted effects.

3 Methods

Building from the pioneering work of the DIO and DINTO we propose to develop a new ontology, DIDEO. The DIDEO will address the aforementioned limitations of the two ontologies while being in alignment with the van Roon conceptual model. Moreover, the ontology will comply with principles of good practice in ontology development, such as formulated by the OBO Foundry[36], for instance:

Reuse of pre-existing resources – Because integration of data is among the key rationales for using ontologies, it is crucial to use Unique Resource Identifiers to refer to the same entities even across domains. We strive to reuse entities from pre-existing ontologies wherever reasonable. Reuse of entities may be limited by the fact that the basic ontological commitments of the source and the target ontology need to be the same. For instance, an ontology that defines ‘drug’ as ‘a chemical entity that bears a drug role that is realized by its use in a pharmacotherapy’ cannot import an individual drug, for instance ‘acetaminophen’ from an ontology that defines drugs as chemicals that are used in pharmacotherapy². In our example both ontologies represent ‘drug’, but each representation comes with a different ontological commitment. One way to assure consistency of ontological commitment is to select entities from ontologies using the same upper ontology.

Use of an upper ontology relevant in biomedical informatics – The DIDEO should support the integration of drug-drug interaction data with data on other biomedically relevant phenomena, for example proteins, protein interactions, laboratory methods and clinical studies. We reuse entities from the Drug Ontology (DRON) [29][30][37] the Ontology of Biomedical Investigation (OBI)[38][39], the Gene Ontology (GO)[40][41] and the Information Artifact Ontology (IAO)[42] These ontologies are all listed on the OBO Foundry[43] webpage. The reuse of DRON, OBI and IAO commits the DIDEO to use the Basic Formal Ontology (BFO) [26][44] as the upper ontology. The Gene Ontology is not using any upper ontology, but multiple classes from GO have been subsumed under BFO classes in the aforementioned ontologies.

¹ The axiomatization of the class actually falls short of specifying that, since being a DDI is not part of its necessary and sufficient condition.

² The aspirin in my medicine cabinet is not a drug according to the second definition given above since it is not participating in any pharmacotherapy, yet it is most certainly a drug according to the first definition.

Representation of biological and biomedical entities was one of the use cases in the development of BFO[26]. In addition, BFO provides well-documented representations for roles, functions and dispositions[45], which are also relevant for biological and biomedical phenomena. Hence we plan to use BFO 2.0 as our upper level and we will import existing OWL entities for reuse using the MIREOT methodology[46] implemented in the MIREOT Protégé plugin[47]. One open question when using MIREOT to import terms from pre-existing ontologies is how to track changes in the source ontologies. MIREOT [46] relies on OBO Foundry’s internal (non-automatic) monitoring process, which might not be an optimal solution, especially when the methodology is applied outside the Foundry.

Compliance with relevant standards of drug representation – The initial development of DIDEO re-uses the drug representation of the Drug Ontology[29] (DRON) which is based on RxNorm[32]. For a PDDI ontology, active ingredients must not be assigned the status of drugs, because the excipients, route of administration, and dose impact the potential for, likelihood of, and severity of interactions. DRON provides ontologically sound representations of Clinical Drugs and Branded Drugs. These representations contain information about dosage and intended route of administration. In addition, DRON provides information about drug ingredients that is linked to the Chemical Entities of Biological Interest ontology [29][30][48].

Community-driven development – Once the initial OWL version of DIDEO is created it will be made publicly available and the project will be continued as an open source project. In addition, we aim to build a community of consumers/contributors to help us create, expand and maintain the ontology. We have already reached out to the developers of DINTO. Since DINTO is using BFO as the upper ontology just as DIDEO we want to investigate possibilities to align our efforts.

4 Results

To appropriately represent PDDI knowledge and its evidence, a definition of PDDIs is crucial. We start our definition by making a basic ontological categorization based on BFO 2.0[49]. The most basic ontological categories of BFO are independent continuant (such as material entities), generically dependent continuants (such as information content entities), specifically dependent continuants (such as qualities and dispositions) and occurrents (such as processes)[26]. Consider an individual PDDI. We might be tempted to categorize it as an occurrent, since the term “interaction” points to a process. However, a PDDI is not an actual process. It is also not a potential or disposition inhering in a substance that may or may not be realized. Representing PDDI in that way, would neglect the fact, that in PDDI research we are collecting information about occurrences that could be drug-drug interactions. Rather, a PDDI is a piece of information about the possible effects of a certain event, for instance the co-administration of azithromycin and ergot alkaloids. We propose the following definition for PDDI:

“A potential drug-drug interaction (PDDI) is an information content entity that specifies the possibility of a drug-drug interaction based on either reasonable extrapolation about drug-drug interaction mechanisms or a data item

created by clinical studies, clinical observation or physiological experiments.”

From this starting point it is crucial to represent the informational bases of PDDIs, namely a) reasonable extrapolation, b) physiological observations from clinical studies and c) drug-drug interaction observational data and d) mechanistic assertions that are useful for inferring drug-drug interactions (derived either from clinical studies or from various experiments, such as inhibition and transport protein experiments)

The aim of DIDEO is to adequately represent all types of PDDI evidence, as well as their differing bases. Not all instances of PDDI evidence data are based on actually observing a drug-drug interaction. In many instances, it is unclear from the evidence whether any actual interaction between the object and precipitant drug has occurred. Hence, it would not be appropriate for us to code the ontology in a way that implies, from the existence of the PDDI evidence, the existence of at least one instance of the specific drug-drug interaction. These bases imply the existence of specific physiological processes, drugs, drug components, and in some cases even DDIs. The following example describes a case of PDDI without evidence for an instance of the actual DDI:

Assume the class ‘potential azithromycin-ergot alkaloid interaction evidence data’ exists in an OWL ontology. To link the data item to actual physiological processes we could axiomatize that each element of this class is about some element in the class ‘azithromycin-ergot alkaloid interaction’. But our axiom will then imply that at least one element of the latter class exists.

This *existential import* can be avoided by using a feature novel in OWL2 called ‘punning’. ‘Punning’ enables users to assign the same name to an OWL class and an OWL individual, allowing the use of the individual when referring to the type and the use of the class when referring to individuals or aggregate of individuals. Despite the two entities bearing the same name, no cross-inferences are made when reasoning[50]. Thus, punning would allow stating that a PDDI is about at least two types of drugs, without affirming that each individual entity is about one individual portion of that drug.

Figure 3 shows a DIDEO representation (without punning) of a clinical study potentially useful as evidence for a mechanistic assertion that could be used for inferring drug-drug interactions. The figure shows how information about the drug and enzyme involved in the study can be traced from the study data item. It would also be possible to track the type of clinical study. Some ontological commitments and design decision depicted in Figure 3 warrant more detail:

- A particular simvastatin metabolism process is the proper occurrent part of a drug metabolism assay. In natural language we might say that the metabolism *participates in* the assay. But one of the ontological commitments of BFO is that the ‘participates in’-relation only holds between a continuant and an occurrent[51]. To simply say that the metabolism and the assay temporarily overlap would not be sufficient here. Many processes overlap in time, without being interrelated in any other way.
- When we represent substances that are referred to by mass nouns, we talk about a portion of that substance (e.g. the simvastatin metabolism has a portion of simvastatin as a participant). This is inline with the practice used in numerous OBO Foundry ontologies, to distinguish between a specific instance of a portion of the entity and the term denoting its type[52].

- Both the portion of simvastatin and the portion of CYP3A4 are affirmed to participate in the process at some times. Claiming their participation at all times we would exclude the possibility that there is no CYP3A4 available to participate in the metabolism, while the final part of the process are still occurring. Moreover, ‘participate at some time’ is not a negation of ‘participate at all time’, but, entails the latter [53].
- Assays of a simvastatin metabolism establish that CYP3A4 is the bearer of a disposition called ‘drug metabolism enabler disposition’ that enables the metabolism, and is realized by that metabolism.
- The outcome of the assay is a simvastatin metabolism data item. Data items are defined as “an information content entity that is intended to be a truthful statement about something (modulo, e.g., measurement precision or other systematic errors) and is constructed/acquired by a method which reliably tends to produce (approximately) truthful statements”[42] . The simvastatin metabolism data item is a member of the class of data items that are specified output of some drug metabolism assays.

5 Conclusion

This paper offers justification for a new foundational domain representation for PDDI knowledge and describes the initial steps toward its development. This new semantic model, the Drug-drug Interaction and Drug-drug Interaction Evidence Ontology (DIDEO), is motivated by the needs of experts who must search, evaluate, and synthesize PDDI evidence into knowledge claims. The results reported in this paper form a foundation for the further development of DIDEO. We will now start to implement DIDEO in OWL in order to test its applicability with respect to competency questions [23] specified in as part of the “Addressing gaps in clinically useful evidence on drug-drug interactions” project. During implementation we will seek to coordinate our efforts with the developers of DINTO and DRON.

Acknowledgments: For all authors: This project is supported by a grant from the National Library of Medicine: “Addressing gaps in clinically useful evidence on drug-drug interactions” (R01LM011838-01) and the National Institute of Aging “Improving medication safety for nursing home residents prescribed psychotropic drugs” (K01 AG044433-01). The authors thank Michel Dumontier and Alan Ruttenberg for their valuable comments, which have significantly improved the paper.

For PE: This work is supported by the National Center For Advancing Translational Sciences of the National Institutes of Health under Award Number KL2TR000146.

For DM: This work is partially supported by the Agency for Healthcare Research and Quality (AHRQ) Grant No. 1R13HS021826-01 (Malone DC-PI)

For JS: This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no 246016.

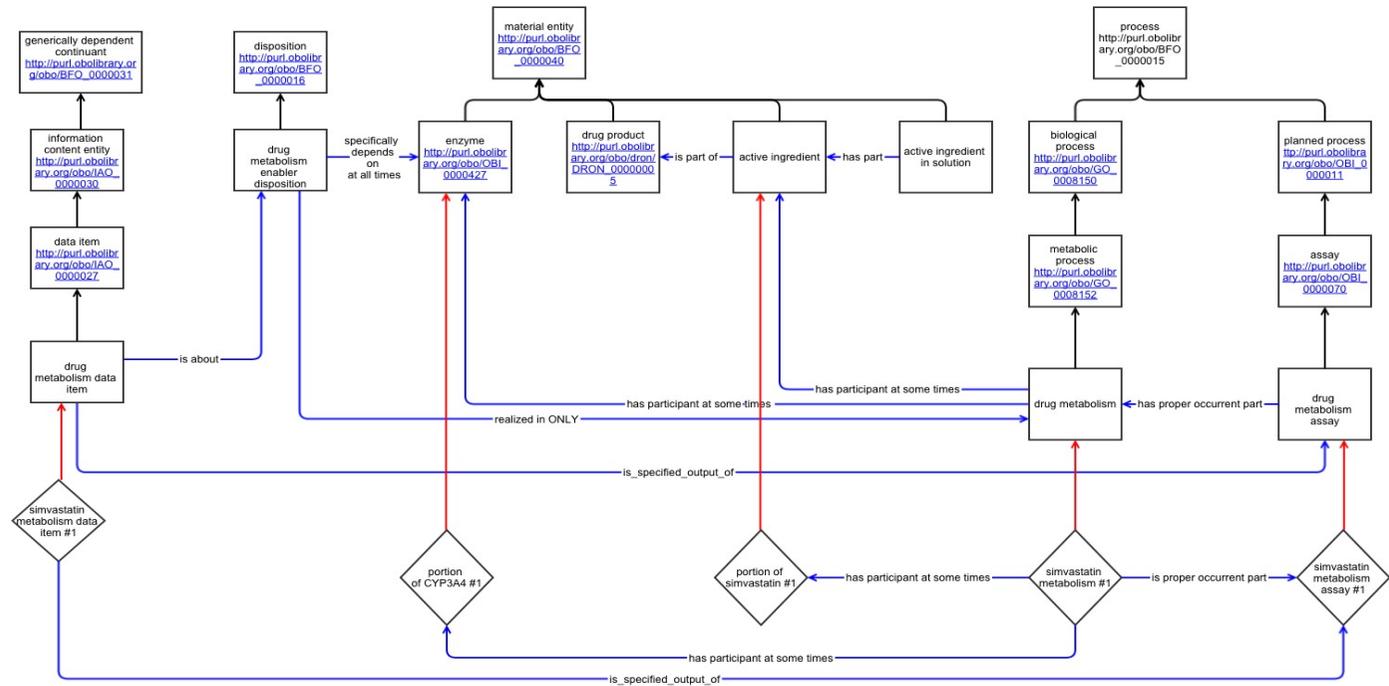


Fig. 3: Representation of instances derived from DIKB in the new semantic model, DIDEO. (Boxes represent OWL named classes; Text in boxes gives the label for the class and for imported terms the class URI, diamonds represent individuals, arrows represent object properties or `rdfs:subClassOf` (i.e., the “is a” relation)).

References

1. Magro, L., Moretti, U., Leone, R.: Epidemiology and characteristics of adverse drug reactions caused by drug-drug interactions. *Expert Opin. Drug Saf.* 11, 83–94 (2012).
2. FastStats - Emergency Department Visits, <http://www.cdc.gov/nchs/fastats/emergency-department.htm>. Last accessed: 07/29/2014.
3. Horn, J.R., Hansten, P.D.: Sources of error in drug interactions: the Swiss cheese model. *Pharm. Times.* 70, 53–55 (2004).
4. Wang, L.M., Wong, M., Lightwood, J.M., Cheng, C.M.: Black box warning contraindicated comedications: concordance among three major drug interaction screening programs. *Ann. Pharmacother.* 44, 28–34 (2010).
5. Saverno, K.R., Hines, L.E., Warholak, T.L., Grizzle, A.J., Babits, L., Clark, C., Taylor, A.M., Malone, D.C.: Ability of pharmacy clinical decision-support software to alert users about clinically important drug-drug interactions. *J. Am. Med. Inform. Assoc. JAMIA.* 18, 32–37 (2011).
6. Abarca, J., Malone, D.C., Armstrong, E.P., Grizzle, A.J., Hansten, P.D., Van Bergen, R.C., Lipton, R.B.: Concordance of severity ratings provided in four drug interaction compendia. *J. Am. Pharm. Assoc. JAPhA.* 44, 136–141 (2004).
7. Hines, L.E., Murphy, J.E., Grizzle, A.J., Malone, D.C.: Critical issues associated with drug-drug interactions: highlights of a multi-stakeholder conference. *Am. J. Health-Syst. Pharm. AJHP Off. J. Am. Soc. Health-Syst. Pharm.* 68, 941–946 (2011).
8. Marshall, M.S., Boyce, R., Deus, H.F., Zhao, J., Willighagen, E.L., Samwald, M., Pichler, E., Hajagos, J., Prud'Hommeaux, E., Stephens, S.: Ontology Paper: Emerging Practices for Mapping and Linking Life Sciences Data Using RDF - A Case Series. *Web Semant.* 14, 2–13 (2012).
9. Van Roon, E.N., Flikweert, S., le Comte, M., Langendijk, P.N.J., Kwee-Zuiderwijk, W.J.M., Smits, P., Brouwers, J.R.B.J.: Clinical relevance of drug-drug interactions : a structured assessment procedure. *Drug Saf. Int. J. Med. Toxicol. Drug Exp.* 28, 1131–1139 (2005).

10. Preskorn, S.H.: Drug-drug interactions: proof of relevance (part I). *J. Psychiatr. Pract.* 11, 116–122 (2005).
11. Preskorn, S.H.: Drug-drug interactions: Proof of relevance (Part II): cause of tolerability problems or noncompliance. *J. Psychiatr. Pract.* 11, 397–401 (2005).
12. Hansten, P.D., Horn, J.R., Hazlet, T.K.: ORCA: OpeRational ClassificAtion of drug interactions. *J. Am. Pharm. Assoc. WashingtonDC* 1996. 41, 161–165 (2001).
13. Ylva Böttiger, K.L.: SFINX-a drug-drug interaction database designed for clinical decision support systems. *Eur. J. Clin. Pharmacol.* 65, 627–33 (2009).
14. Boyce, R.D.: Drug Interaction Knowledge Base, <http://purl.org/net/drug-interaction-knowledge-base/>. Last accessed: 08/09/2014
15. Boyce, R.D., Collins, C., Horn, J., Kalet, I.: Modeling Drug Mechanism Knowledge Using Evidence and Truth Maintenance. *IEEE Trans. Inf. Technol. Biomed.* 11, 386–397 (2007).
16. Boyce, R., Collins, C., Horn, J., Kalet, I.: Computing with evidence Part I: A drug-mechanism evidence taxonomy oriented toward confidence assignment. *J. Biomed. Inform.* 42, 979–989 (2009).
17. Boyce, R., Collins, C., Horn, J., Kalet, I.: Computing with evidence Part II: An evidential approach to predicting metabolic drug-drug interactions. *J. Biomed. Inform.* 42, 990–1003 (2009).
18. US Food and Drug Administration: Guidance for Industry Drug Interaction Studies — Study Design, Data Analysis, Implications for Dosing, and Labeling Recommendations, (1999).
19. US Food and Drug Administration: Guidance for Industry Drug Interaction Studies — Study Design, Data Analysis, Implications for Dosing, and Labeling Recommendations, (2006).
20. US Food and Drug Administration: Guidance for Industry Drug Interaction Studies — Study Design, Data Analysis, Implications for Dosing, and Labeling Recommendations, (2012).
21. Boyce, R.D., Collins, C., Horn, J., Kalet, I.: Modeling drug mechanism knowledge using evidence and truth maintenance. *IEEE Trans. Inf. Technol. Biomed. Publ. IEEE Eng. Med. Biol. Soc.* 11, 386–397 (2007).
22. Boyce, R.D.: A Draft Evidence Taxonomy and Inclusion Criteria for the Drug Interaction Knowledge Base (DIKB),

- <http://purl.net/net/drug-interaction-knowledge-base/evidence-types-and-inclusion-criteria>. Last accessed: 10/01/2014.
23. Boyce, R.D.: Draft Competency Questions for a Drug Interaction Knowledge Base Derived from Discussions with Clinical Experts, <http://purl.net/net/drug-interaction-knowledge-base/competency-questions>. Last accessed: 10/01/2014.
 24. Arikuma, T., Yoshikawa, S., Azuma, R., Watanabe, K., Matsu-mura, K., Konagaya, A.: Drug interaction prediction using ontology-driven hypothetical assertion framework for pathway generation fol-lowed by numerical simulation. *BMC Bioinformatics*. 9, S11 (2008).
 25. Herrero-Zazo, M., Hastings, J., Segura-Bedmar, I., Croset, S., Martinez, P., Steinbeck, C.: An ontology for drug-drug interactions. http://ceur-ws.org/Vol-1114/Session3_Herrero-Zazo.pdf. Last ac-cessed: 10/05/2014
 26. Grenon, P., Smith, B., Goldberg, L.: *Biodynamic Ontology: Applying BFO in the Biomedical Domain*. Stud. Health Technol. In-form. pp. 20–38. IOS Press (2004).
 27. The Drug Interaction Ontology, <http://www.biomedcentral.com/content/supplementary/1471-2105-9-S6-S11-S1.owl>. Last accessed: 08/09/2014.
 28. NCI Thesaurus, <http://ncit.nci.nih.gov>. Last accessed: 08/09/2014.
 29. Hanna, J., Joseph, E., Brochhausen, M., Hogan, W.R.: Building a drug ontology based on RxNorm and other sources. *J. Biomed. Semant.* 4, 44 (2013).
 30. Hogan, W.R., Hanna, J., Joseph, E., Brochhausen, M.: Towards a Consistent and Scientifically Accurate Drug Ontology. http://www2.unb.ca/csas/data/ws/icbo2013/papers/research/icbo2013_submission_40.pdf. Last accessed: 10/05/2014.
 31. Luciano, J.S., Andersson, B., Batchelor, C., Bodenreider, O., Clark, T., Denney, C.K., Domarew, C., Gambet, T., Harland, L., Jentzsch, A., Kashyap, V., Kos, P., Kozlovsky, J., Lebo, T., Mar-shall, S.M., McCusker, J.P., McGuinness, D.L., Ogbuji, C., Pichler, E., Powers, R.L., Prud'hommeaux, E., Samwald, M., Schriml, L., Tonellato, P.J., Whetzel, P.L., Zhao, J., Stephens, S., Dumontier, M.: The Translational Medicine Ontology and Knowledge Base: driving personalized medicine by bridging the gap between bench and bed-side. *J. Biomed. Semant.* 2 Suppl 2, S1 (2011).

32. Nelson, S.J., Zeng, K., Kilbourne, J., Powell, T., Moore, R.: Normalized names for clinical drugs: RxNorm at 6 years. *J. Am. Med. Inform. Assoc.* 18, 441–448 (2011).
33. The Drug-Drug Interaction Ontology, https://4625527d7e0f0589865f115fb0c87fc18bef216f.googleusercontent.com/host/0B-7Po9tR1KLUNkpNdmFEcG44RjA/DINTO_1BFO.owl. Last accessed: 10/01/2014
34. Knox, C., Law, V., Jewison, T., Liu, P., Ly, S., Frolkis, A., Pon, A., Banco, K., Mak, C., Neveu, V., Djoumbou, Y., Eisner, R., Guo, A.C., Wishart, D.S.: DrugBank 3.0: a comprehensive resource for “omics” research on drugs. *Nucleic Acids Res.* 39, D1035–1041 (2011).
35. Law, V., Knox, C., Djoumbou, Y., Jewison, T., Guo, A.C., Liu, Y., Maciejewski, A., Arndt, D., Wilson, M., Neveu, V., Tang, A., Gabriel, G., Ly, C., Adamjee, S., Dame, Z.T., Han, B., Zhou, Y., Wishart, D.S.: DrugBank 4.0: shedding new light on drug metabolism. *Nucleic Acids Res.* 42, D1091–1097 (2014).
36. The OBO Foundry: OBO Foundry Principles, <http://www.obofoundry.org/crit.shtml>. Last accessed: 09/15/2014.
37. The Drug Ontology, <http://purl.obolibrary.org/obo/dron.owl>. Last accessed: 10/05/2014.
38. The Ontology of Biomedical Investigation, <http://purl.obolibrary.org/obo/obi.owl>. Last accessed: 10/05/2014
39. Brinkman, R.R., Courtot, M., Derom, D., Fostel, J.M., He, Y., Lord, P., Malone, J., Parkinson, H., Peters, B., Rocca-Serra, P., Rutenber, A., Sansone, S.-A., Soldatova, L.N., Stoeckert, C.J., Turner, J.A., Zheng, J., OBI consortium: Modeling biomedical experimental processes with OBI. *J. Biomed. Semant.* 1 Suppl 1, S7 (2010).
40. The Gene Ontology, <http://purl.obolibrary.org/obo/go.owl>. Last accessed: 10/05/2014
41. Ashburner, M., Ball, C.A., Blake, J.A., Botstein, D., Butler, H., Cherry, J.M., Davis, A.P., Dolinski, K., Dwight, S.S., Eppig, J.T., Harris, M.A., Hill, D.P., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J.C., Richardson, J.E., Ringwald, M., Rubin, G.M., Sherlock, G.: Gene ontology: tool for the unification of biology. *The Gene Ontology Consortium. Nat. Genet.* 25, 25–29 (2000).
42. The Information Artifact Ontology, <http://purl.obolibrary.org/obo/iao.owl>. Last accessed: 10/05/2014

43. Smith, B., Ashburner, M., Rosse, C., Bard, J., Bug, W., Ceusters, W., Goldberg, L.J., Eilbeck, K., Ireland, A., Mungall, C.J., OBI Consortium, Leontis, N., Rocca-Serra, P., Ruttenberg, A., Sansone, S.-A., Scheuermann, R.H., Shah, N., Whetzel, P.L., Lewis, S.: The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nat. Biotechnol.* 25, 1251–1255 (2007).
44. Grenon, P., Smith, B.: SNAP and SPAN: Towards Dynamic Spatial Ontology. *Spat. Cogn. Comput.* 4, 69–103 (2004).
45. Arp, R., Smith, B.: Function, Role, and Disposition in Basic Formal Ontology. *Nat. Preced.* (2008).
46. Courtot, M., Gibson, F., Lister, A.L., Malone, J., Schober, D., Brinkman, R.R., Ruttenberg, A.: MIREOT: The Minimum Information to Reference an External Ontology Term. *Appl Ontol.* 6, 23–33 (2011).
47. Hanna, J., Chen, C., Crow, A., Hall, R., Liu, J., Pendurthi, T., Schmidt, T., Jennings, S.F., Brochhausen, M., Hogan, W.R.: Simplifying MIREOT: a MIREOT Protégé Plugin. *International Semantic Web Conference Posters & Demos*, volume 914 of CEUR Workshop Proceedings. CEUR-WS.org, Boston (2012).
48. Hastings, J., de Matos, P., Dekker, A., Ennis, M., Harsha, B., Kale, N., Muthukrishnan, V., Owen, G., Turner, S., Williams, M., Steinbeck, C.: The ChEBI reference database and ontology for biologically relevant chemistry: enhancements for 2013. *Nucleic Acids Res.* 41, D456–D463 (2013).
49. The Basic Formal Ontology, Release 2014-05-03, <http://purl.obolibrary.org/obo/bfo/2014-05-03/bfo.owl>. Last accessed: 10/05/2014
50. Golbreich, C., Wallace, E.K.: OWL 2 Web Ontology Language New Features and Rationale (Second Edition), <http://www.w3.org/TR/owl2-new-features/>. Last accessed: 09/25/2014.
51. Smith, B., Almeida, M., Bona, J., Brochhausen, M., Ceusters, W., Courtot, M., Dipert, R., Goldfain, A., Grenon, P., Hastings, J., Hogan, W.R., Jacuzzo, L., Johansson, I., Mungall, C.J., Natale, D., Neuhaus, F., Petosa, A., Rovetto, R., Ruttenberg, A., Ressler, M., Schulz, S.: Basic Formal Ontology 2.0 Draft Specification and User's Guide, <https://bfo.googlecode.com/svn/trunk/docs/bfo2-reference/BFO2-Reference.docx>, (2014). Last accessed: 10/05/2014

52. Kumar, A., Smith, B., Novotny, D.D.: Biomedical informatics and granularity. *Comp. Funct. Genomics*. 5, 501–508 (2004).
53. Schulz, S., Ruttenberg, A., Courtot, M., Hastings, J.: Release notes for the BFO 2 2012-07-20 “Graz” OWL Release, <http://purl.obolibrary.org/obo/bfo/2012-07-20/ReleaseNotes>, (2012). Last accessed: 10/05/2014

Proceedings
of
THE SECOND
INTERNATIONAL WORKSHOP
ON
DEFINITIONS IN ONTOLOGIES
(IWOOD 2014)

*Proceedings
edited by*

*Selja Seppälä
Patrick Ray
Alan Ruttenberg*

*IWOOD 2014 was held in conjunction with
the
5th International Conference on Biomedical Ontology (ICBO 2014)
Houston, TX, USA
October 6-7, 2014*

Preface

The IWOOD 2014* workshop was the second workshop on definitions in ontologies in as many years. The first workshop (DO 2013) was held last year in Montreal in conjunction with ICBO 2013. The focus of this second workshop was on definition practices in either human or machine-assisted ontology development.

Explicit definitions of terms in ontologies serve a number of purposes. Logical definitions allow reasoners to create inferred hierarchies, lessening the burden of asserting and checking the validity of subsumptions. Natural language definitions help to ameliorate the pervasive problem of low inter-annotator agreement. In specialized domains, experts will know their own field well, but may only have limited knowledge of adjacent disciplines. Good definitions make it possible for non-experts to understand unfamiliar terms and thereby make it possible for more confident reuse of terms by external ontologies, which in turn facilitates data integration.

The goal of this workshop was to bring together researchers and developers in the biomedical domain to discuss difficulties that arise in definition construction with a view to sharing strategies. Even within the narrow domain of definition construction, cross-fertilization of ideas from related disciplines should yield benefits in quality and help reinforce common approaches and identify novel ones.

The communications published in these proceedings address the theoretical, methodological, and pragmatic criteria one should consider before engaging in the activity of defining terms in ontologies. The main goals of the first article are to shed light on the nature of logical and textual definitions, to explore the relationship that exists between them, and to make recommendations for an improvement between the two. The second paper addresses issues related to the general activity of defining terms as illustrated by the definition of 'sign' and 'symptom': asking preliminary questions; looking up and comparing existing definitions; considering one's ontological framework (BFO, other...); taking or not into account actual usages of the terms; discussing the relevance of alternative definitions; etc. The last communication describes how definitions can be used to assess the coherence and overall quality of an ontology according to some prerequisites, and proposes a systematic methodology for revising and updating existing definitions to conform to the set quality criteria.

Selja Seppälä, Yonatan Schreiber and Alan Ruttenberg
Textual and logical definitions in ontologies

Alexander P. Cox, Patrick Ray, Mark Jensen and Alexander D. Diehl
Defining 'sign' and 'symptom'

Werner Ceusters
An alternative terminology for pain assessment

1 We changed the acronym of the workshop from DO 2014 to IWOOD 2014 to avoid any confusion with the Disease Ontology acronym, also DO.

Acknowledgements

This workshop would not have been possible without the thorough reviews of the international scholars who accepted to be part of the program committee. We would like to thank them for their participation and support to this workshop. We would also like to thank the participants for helpful insights and thoughtful discussions. Finally, we thank the ICBO organizers for hosting the workshop and providing assistance with logistics, as well as the Swiss National Science Foundation and the State University of New York at Buffalo for their support.

October 2014
IWOOD organizers

More information is available on the workshop's website:
<https://sites.google.com/site/definitionsinontologies/>.

Organizing Committee

Selja Seppälä (University at Buffalo, USA)
Patrick Ray (University at Buffalo, USA)
Alan Ruttenberg (University at Buffalo, USA)

Program Committee

Nathalie Aussenac-Gilles (National Center for Scientific Research (CNRS), France)
Mélanie Courtot (MBB Department Simon Fraser University and BC Public Health Microbiology & Reference Laboratory, Canada)
Natalia Grabar (Université de Lille 3, France)
Janna Hastings (European Bioinformatics Institute, Cambridge, UK)
James Malone (European Bioinformatics Institute, Cambridge, UK)
Alexis Nasr (Aix Marseille Université, France)
Richard Power (The Open University, UK)
Allan Third (The Open University, UK)

Textual and logical definitions in ontologies

Selja Seppälä
Department of Philosophy
University at Buffalo
Buffalo, NY, 14260, USA
Email: seljamar@buffalo.edu

Yonatan Schreiber
CUBRC
4455 Genesee St.
Buffalo, NY, 14225, USA
Email: yonatan.schreiber@cubrc.org

Alan Ruttenberg
School of Dental Medicine
University at Buffalo
Buffalo, NY, 14215, USA
Email: alanruttenberg@gmail.com

Abstract—We discuss the structure and functions of definitions and axioms in ontologies from the perspective of a terminologist and logician respectively. By working through a few examples of the correspondence between parts of the textual definitions and the axioms, we show how to compare and contrast each and how each perspective reveals areas for improvement. Having established a correspondence between the textual and logical parts of ontology term definitions, we discuss the possibility of developing tools that help developers improve their ontologies. Such tools could be used to check both the textual definitions against the asserted axioms and vice versa. In addition, we propose a few other ways of checking the contents of textual definitions.

Keywords—*textual definitions, natural language definitions, logical definitions, OWL axioms, checking definition contents, problems in definitions, functions of definitions in ontologies, recommendations for definitions in ontologies, ontologies, terminology*

I. INTRODUCTION

Ontologies have on the one hand axioms that form parts of the logical definition of terms, and on the other hand natural language definitions and other documentation of those terms.

However, the ontological world does not seem to have a theory of what the functions of textual as opposed to logical definitions are. The result of that is authoring practices that vary widely. There are nevertheless correspondences (to a certain extent) between phrases in the textual parts and the logical parts. We can use an expectation of correspondences between the textual and logical parts to build tools that help developers improve their ontologies and provide guidelines for identifying issues in axioms and definitions. Aspects we can exploit are:

- Leverage logic to help establish correspondences between the textual definition and the axioms.
- Leverage principles of organizing terminological entities (definitions, notes, . . .) to characterize the logical parts.
- Measure some part of the quality of an ontology in terms of these correspondences.

Thus, it may be feasible to bring automated methods used in the terminological world to bear on both establishing the correspondences and identifying quality issues in the textual part that could be mapped to quality issues in the logical form.

In this communication, we show examples of varying definition practices in ontologies to support our first thesis and describe issues in definition practices. We discuss the structure and functions of definitions and logical parts from

the perspective of a terminologist and a logician respectively. By working through a few examples of the correspondence between parts of the textual definitions and logical parts, we show how to compare and contrast each and how each perspective reveals areas for improvement.

We suggest that it is possible to write tools that analyze textual definitions with the goal of offering places for improvement. We discuss how such tools could be leveraged to check the contents of both textual and logical definitions for terms in ontologies. Our recommendations could also contribute to supplementing the specifications of the OBO Foundry principles on textual definitions.¹

II. TEXTUAL DEFINITIONS

In an ontology, a textual definition is, ideally, a short sentence found as the object of an annotation property designated for that purpose. This kind of natural language definition is also found in specialized terminological dictionaries. The account we give in the present communication is thus based on the more developed account of terminological definitions in [1], [2].

A good definition conveys the intended meaning of an ontology term — we will come back to this later — by describing the type of thing to which the term refers. For example, the Cell Type Ontology (CL) contains the following definition for the term *leukocyte*:

- (a) *An achromatic cell of the myeloid or lymphoid lineages capable of ameboid movement, found in blood or other tissue.*

This example shows that the term *leukocyte* refers to those things that are of the type achromatic cell and that are distinguished from other achromatic cells in virtue of being: of the myeloid or lymphoid lineages; capable of ameboid movement; found in blood or other tissue.

As we can see, a definition normally states the type of thing to which the instances of the defined term belong, and distinguishes these instances from the type and from other things falling under the same type by listing one or more of the characteristics of the instances of the term.

The first part, the head of the definition is called the *genus*; a distinguishing part, *differentia*. Thus, a definition has a structure where each part is related to the defined term's instances by some type of relation:

¹http://obofoundry.org/wiki/index.php/FP_006_textual_definitions

- In the classical Aristotelian form, the genus (implicitly) expresses an *is_a* relation, as in example (a) above, which we read as: *a leukocyte is an achromatic cell*.
- The differentia may express any kind of relation relevant for describing and distinguishing the kinds of things to which the defined term refers. In example (a) above, the relations expressed in the definition of *leukocyte* are respectively *develops_from* (*of the myeloid or lymphoid lineages*), *capable_of* (*capable of ameoboid movement*), and *located_in* (*found in blood or other tissue*).

A textual definition also has a logical form that derives from the relationship between its *intension* (that which is said about the referent) and its *extension* (the set of instances that fall under the intension). We can distinguish three main logical forms:²

Classical definition A definition where the intension holds for *all* instances of the type that is defined, as in *Every instance of X is a Y and all instances of X Z...* In this case, the characteristics expressed by Y and Z are necessary and, in the ideal case, they are jointly sufficient for including all instances of X and distinguishing them from other instances of Y. The ideal case corresponds to the Aristotelian definition by necessary and sufficient conditions. A standard example of classical definition is that of triangle: *A rectilinear figure that has three sides*. (All triangles are rectilinear and have three sides.)

Typical or prototypical definition A definition where the intension holds for *most* of the instances of the type that is defined, as in *Every instance of X is a Y and most instances of X Z...* An example of prototypical definition for a swan would be *An aquatic bird with a long neck, usually having white plumage*. (Most swans are white.)

Instance definition A definition where the intension holds for *only a single* instance, as in *X is the only Y that Z...* These correspond to proper definite descriptions. This kind of definition would apply, for example, to ontologies that include what may be considered as proper names, such as the *Large Hadron Collider* (LHC) in an ontology of nuclear physics. In this case, the relevant kind of differentiae would probably inform us about the geographical location of the LHC and specify that it is (or was until some point in time) “the world’s largest and most powerful particle accelerator.”³ The definition could be even more specific and tell us about the length of the ring and the number of magnets that compose it.

Normally, ontologies contain classical definitions because their function is to disambiguate terms. This is not to say that the other forms cannot appear in the textual definitions, but this would not be ideal with respect to the function they are meant to fulfill in this context; without necessary and sufficient conditions it becomes possible to interpret terms in a manner that deviates from their intended use.

Indeed, the main function of textual definitions in ontologies is to specify the intended meaning of the ontology terms in order to avoid ambiguities and errors when, for example,

annotating biomedical research texts or importing terms into other ontologies. Of course, this is also the function of the axioms, as we will see in the next section. However, the latter can be somewhat obscure to non-ontologists who may need more detailed and explicit information about the term and its referent.

Therefore, there is a cognitive advantage in including textual definitions in ontologies. As argued in [1, section 1.3], dictionary-type definitions are meant to adjust users’ lexical competence [3] by modifying (or confirming) their knowledge about the use of terms. In ontologies, definitions allow users to make their use of a term converge toward that of the rest of the users of the ontology. Both the genus and the differentia contribute to the cognitive adjustment: the genus is meant to provide a sort of cognitive anchor by stating a term that should be familiar to the user of the definition; the differentiae are meant to tell the user how the defined thing differs from the thing that is expected to be already known.

III. AXIOMS IN ONTOLOGIES

Axioms in ontologies restrict the intended meaning of a term by asserting necessary conditions for its use. They thus function in a manner analogous to the necessary conditions previously discussed under **Classical definition** in section II. In OWL, it is rarely possible to provide sufficient conditions, so axioms do not on their own constitute full definitions. We distinguish three primary functions of ontology axioms: disambiguation, taxonomic schematization, and fact-modeling.

The function of axioms in the disambiguation of terms is analogous to the function that textual definitions play in disambiguation. Every axiom represents a necessary condition for entities in the terms extension. Axioms thus help to determine the extension of a term by restricting it to those entities meeting the asserted condition. Each additional axiom restricts the extension further, though it is usually not possible to restrict the term to only its intended extension by providing conditions that are jointly sufficient. The most common type of axiom asserts an *is_a* relation that relates the defined term to a parent class by means of the *subclassOf* relation. For the most part, the relatum of such an axiom should correspond directly to the genus in the textual definition.

We call the second function we identify ‘taxonomic schematization’. When employed in this capacity, an axiom asserted for a class provides a schema or template for the axioms of any subclasses. This provides, in our view, robust, principled taxonomic relations between parent, child, and sibling classes. A class’s axioms are inherited by all of its subclasses. This makes it possible to use axioms to suggest differentiae for its child classes, in other words to use these axioms as templates for the axioms of the subclasses. This can be done by asserting a relational axiom for the parent class relating it to some other kind of entity (e.g. by writing an axiom for a class X asserting that any X is ‘part_of some Y’). For every subclass of this related kind, a subclass of the parent can then be distinguished. For example, the axioms specifying the term *infection* in the Infectious Disease Ontology (IDO) can be used to generate the subclass axioms of its child terms, such as *amebiasis* (see the axiom under *SubClass Of (Anonymous Ancestor)* in Figure 1; see also the discussion of this example in section IV-C below).

²X, traditionally called the *definiendum*, stands for the defined term’s referent; Y for the genus; Z for a differentia; Y and Z together for the definition itself, traditionally called the *definiens*.

³<http://home.web.cern.ch/topics/large-hadron-collider>

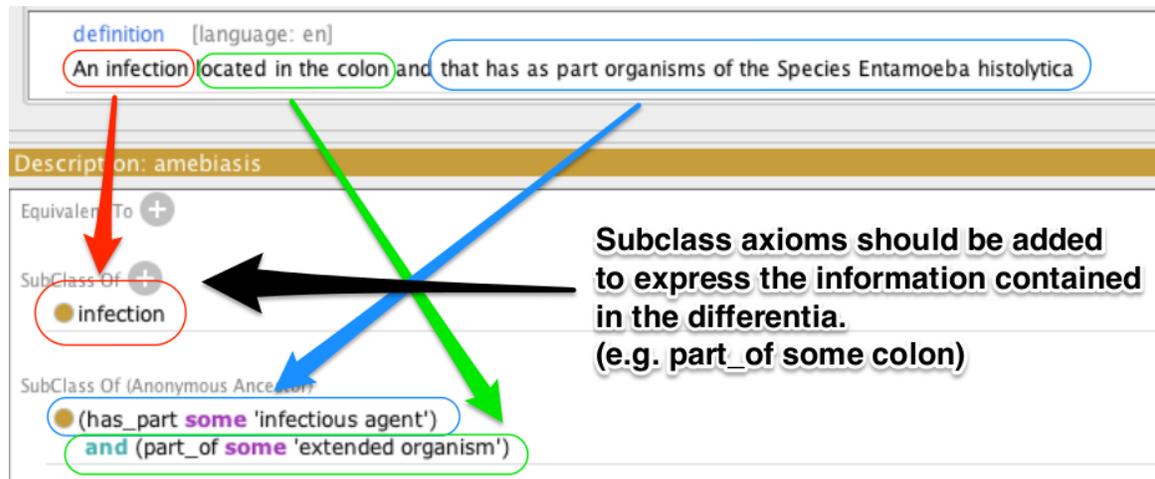


Fig. 1. Correspondences in the parts of the textual definition and the axioms of the IDO term *amebiasis*.

Lastly, we distinguish a fact-modeling function of axioms. An ontology can be considered a specification of a controlled vocabulary for expressing facts in a given domain. Such a vocabulary is much sparser than the vocabulary that would be used to express these facts in natural language, that is, there is a one-many correspondence between ontology terms and words in domain-relevant portions of natural language. This means that the syntax for expressing facts (i.e., assertions between instances) using ontology terms necessarily diverges from the syntax used for expressing the same facts in natural language. The RDF-schema regularizes this syntax substantially, but it is still generally the case that RDF syntax plus the list of terms in the ontology underdetermine how any given fact should be translated from natural language into an expression using the ontologies controlled vocabulary. An important function of axioms in ontologies is to provide a schematic suggestion of how this should be done. Thus, axioms complement textual definitions in contributing cognitively towards regularizing users' employment of terms. For example, the axiom 'is_about some document' in one of the axioms specifying the term *abstract* in the Information Artifact Ontology (IAO) tells us that the relation expressed by the verb *to summarize* in natural language is expressed at the logical level by the *is_about* relation that is part of the controlled vocabulary of the ontology (see annotations in blue in Figure 2).

IV. CORRESPONDENCES BETWEEN TEXTUAL AND LOGICAL DEFINITIONS

As we have seen, axioms and textual definitions have overlapping and complementary functions. Hereafter, we examine how they contribute to conveying the intended meaning of terms. We compare and discuss some examples in the biomedical domain to show how these different forms relate. The examples will show what kinds of issues or inconsistencies can be identified by these comparisons; they reveal at least five types of correspondences. We also give some recommendations as to how to improve both the textual definitions and the related axioms. For sake of readability, we will illustrate the cases with screenshots of the ontology editor Protégé.

A. General recommendation

Based on the identified functions for textual definitions and axioms, we make the following general recommendation: textual definitions should contain content analogous to what is expressed in the axioms, i.e., descriptive content that motivates the logical axioms. The expressions used in natural language may however be more idiomatic than the ontology vocabulary (e.g., the expression *inheres_in* is not very natural). Any complementary information that is deemed useful for understanding the intended meaning of the term but which cannot be included in the axioms should be systematically asserted using other annotation properties.

B. Exact correspondence

Figure 3 shows that the parts of the textual definition of *dead-end host* in IDO correspond exactly to the logical definition by necessary and sufficient conditions. The only difference is in the natural language expression (*bearing*) used for the *has_role* ontological relation — perhaps to avoid the seemingly redundant use of 'role twice. Here, the logical part is useful to fix the intended meaning of the natural language expression.

C. Structural correspondence but more specific content in textual definitions than in axioms

Figure 1 shows that both differentia of the textual definition of the IDO term *amebiasis* contain information of the type expressed in the subclass axioms inherited from the parent class *infection* (see annotations in blue). However, the content conveyed by the parts of the textual definition of *amebiasis* are more specific than the properties and classes expressed in the axiom; they are subproperties of the relations and subclasses of the relata in the axiom.

If these inherited parts are relevant for distinguishing all the subclasses, then all textual definitions at that subclass level should include that kind of information with the specific content that actually distinguishes each entity at that level. If the comparison reveals a match of logical and textual parts at the level of inherited logical parts, this might be a sign that

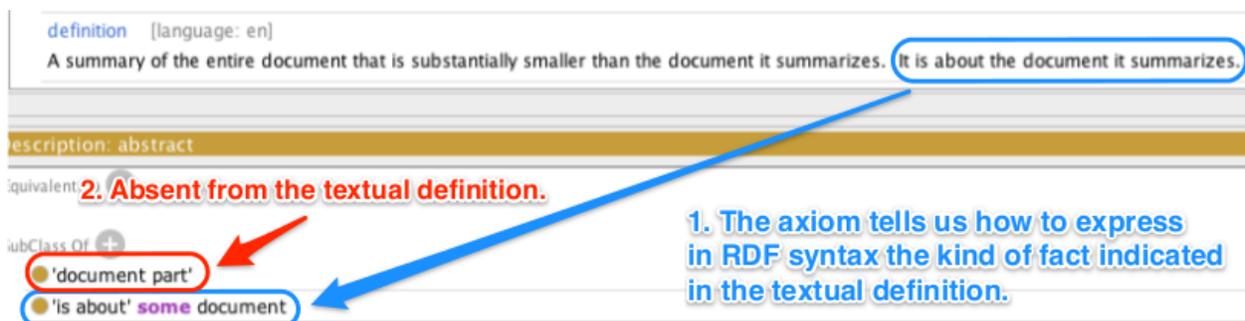


Fig. 2. Correspondences in the parts of the textual definition and the axioms of the IAO term *abstract*.

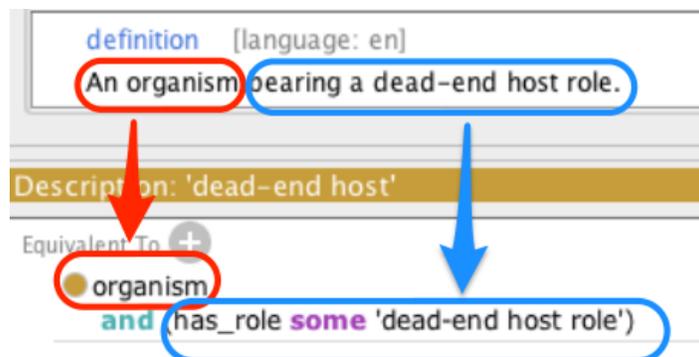


Fig. 3. Correspondences in the parts of the textual definition and the axioms of the IDO term *dead-end host*.

the entity lacks an available subclass axiom. If this is the case, the textual definition can be used as a basis for creating the missing axioms.

We thus recommend that more specific axioms be added whenever the ontology has the resources to include them, i.e., if the terms are defined elsewhere in the ontology. For example, the axioms specifying the IDO term *antiseptic role* in Figure 4 could be completed as follows:

```
subClassOf
  realized_by only has_participant
    some (anatomical entity
      and part_of some organism)
```

D. Incomplete textual definitions

Figure 2 shows that the axiom specifying the term *abstract* in the IAO contains the information ‘document part’ which is absent from the textual definition.

We recommend that the textual definition be completed with this information.

E. Missing axioms

Figure 4 shows that the last part of the textual definition of the IDO term *antiseptic role* does not correspond to any logical part (see annotations in green). However, this more specific differentia serves to distinguish the defined term from (1) *antimicrobial disposition*, which has the same subclass axiom (in blue), and (2) the sibling term *disinfectant role* which is specified by exactly the same axioms. It would therefore be

useful to have an axiom that allows these three terms to be logically distinguished.

Here again, we recommend that the axiom be added whenever the ontology has the resources to include the missing axiom.

F. Redundant parts of axioms or definitions

Logical parts may contain axioms specifying other terms.

Figure 4 shows that part of the axioms specified for *antiseptic role* in IDO correspond to:

- the subclass axioms specifying the term ‘antimicrobial’ — the ‘material entity’ (see annotations in red);
- the subclass axioms specifying the term ‘antimicrobial disposition’ (see annotations in blue).

This should not be a problem at the logical level, since the inferences that are made based on the logical expressions end up being the same.

We recommend nevertheless that the axiom be simplified by using the terms that are specified by those axioms. For example, in this example, the first part of the axiom

```
(inheres_in some
  ('material entity'
    and (has_disposition
      some 'antimicrobial disposition')))
```

can be replaced by the following simpler expression:

```
inheres_in some 'antimicrobial'
```

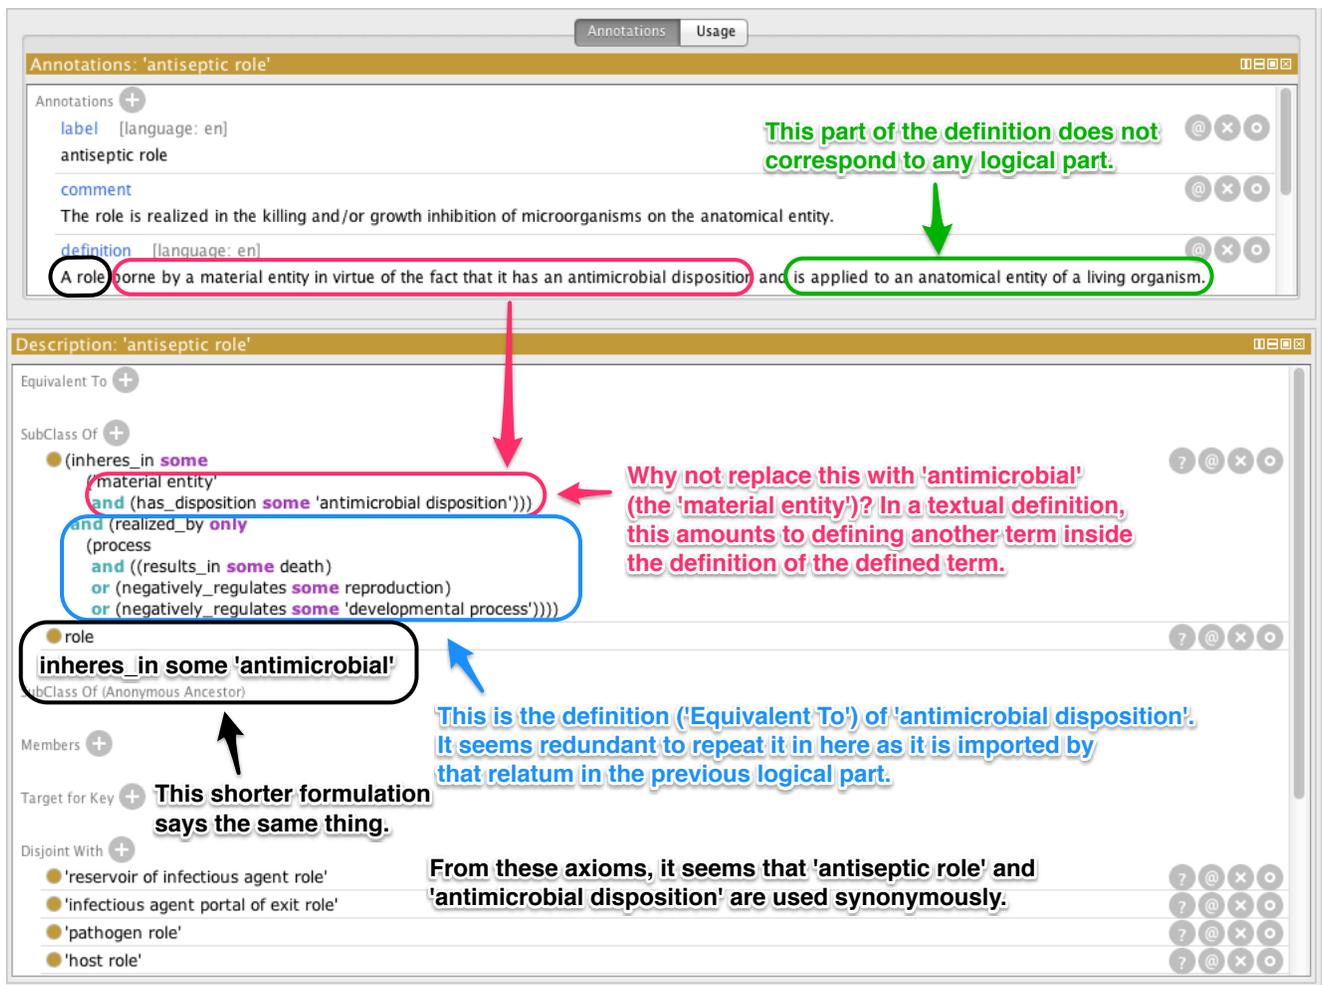


Fig. 4. Correspondences in the parts of the textual definition and the axioms of the IDO term *antiseptic role*.

In a textual definition, this amounts to defining another term within the definition of the defined term, as can be seen in the first differentia of the example (in red), which contains the definition of *antimicrobial*. This lacks conciseness and is generally considered bad practice (see for example [4, 28]). It unnecessarily overloads the contents of the definition — imagine if each term of a definition was replaced by its definition. More importantly, the reader might not recognize that it is the definition of another term and fail to link the defined term with that other one.

We thus recommend that whenever a textual definition contains the definition of another term from the same ontology or an imported ontology, this sub-definition be replaced by the corresponding term. In this example, the differentia *borne by a material entity in virtue of the fact that it has an antimicrobial disposition* should be replaced by *borne by an antimicrobial*. If the reader does not know the term used in the definition, she can (in principle) look it up in the ontology. A system of hyperlinks should also be provided for easier access, as it is done in electronic dictionaries and in the axioms.

V. USING THE CORRESPONDENCES TO HELP IN DEFINITION CHECKING

In ontologies that use semi-automated systems to create the logical and the corresponding textual definitions, such as *TermGenie*⁴, both definition forms are expected to be reasonably consistent. However, when definitions are hand-crafted or imported from other sources, such as other ontologies or, for example, from Wikipedia, various kinds of errors or inconsistencies can creep in, as discussed above. Identifying these problems manually is less rigorous if no guidelines are provided.

To increase reliability of definition-content checking, we propose a method that could be implemented in a computer program to assist ontology editors/curators in carrying out this task in a systematic way. This method can also be used as a guide to manual identification of issues in definitions.

The method consists in the following steps:

- 1) Determine whether any of the terms from either the ontology that is being checked or the imported ontologies appear in the textual definitions.

⁴*TermGenie* is used for creating definitions in the Gene Ontology (GO), (<http://go.termgenie.org>).

- 2) Get the taxonomic hierarchy of the matched terms to the top level.
- 3) Determine whether any of the terms in this hierarchy corresponds to one of the relations in the axioms.
- 4) If no correspondence is found between terms in the textual definition and terms in the axioms, look for a correspondence between the relations expressed in the definition of the textual definition and the object properties in the axioms. This can also be done by taking into account the hierarchy of object properties (if available).
- 5) If matches are found, tag the corresponding part of the textual definition with the corresponding relation–relatum pair (the tagging could supplement the textual definitions with hyperlinks to the entries of the terms and relations used in the definition).
- 6) If mismatches of this kind are identified, manually correct, modify or complete either the textual definition or the axioms, or both according to the recommendations put forward in this paper.

The proposed method may raise some implementation challenges. For example, the first and fourth steps require natural language processing (NLP) methods to correctly identify existing terms and relations in the textual definition. This involves using methods to find inexact matches, for example, plural forms of terms and partial matches, as when only the head of a complex term is used. Matching ontology relations to natural language expressions can also be challenging, as there can be several ways to express a single ontological relation. A solution for relation identification that also involves NLP methods would be analyzing large amounts of definitions in which each part is matched to the corresponding ontological relation to identify the different corresponding expressions. This solution might reveal domain-specific expressions for the more general ontological relations.

VI. OTHER USEFUL WAYS OF CHECKING THE CONTENTS OF TEXTUAL DEFINITIONS

In ontologies, definitions should include only necessary conditions that have the classical all-some form. Thus, they should avoid:

- Particularizing expressions such as *for example, especially, in particular, i.e., such as, ...*, and punctuation signs such as parentheses and colons. Sometimes, differentia may contain hidden examples that should also be avoided, as in the definition of *leukocyte* above which states *found in blood or other tissue*. Here, the specification *blood* is superfluous since it is embedded in a conjunction of which the other conjunct is its superclass.
- Overly generalizing expressions such as *etc., in general, normally, ...*, and disjunctions, as these are linguistic markers of conditions that are not necessary.

Although particularizing and generalizing expressions can be useful for a better understanding of the term (as in example (a) above). These kinds of information should be asserted using other annotation properties.

Furthermore, textual definitions should not contain definitions of other terms, as in the definition of *antiseptic role* examined above (Figure 4). Thus, they should avoid:

- Punctuation signs such as parentheses and colons which are also a sign of new definitions.
- Expressions introducing new information such as *i.e., that is, ...*

The content-related issues presented in this section can be automatically checked with a simple rule-based program that uses, for example, lexico-syntactic patterns. This kind of program can also be used for checking the conformity of the surface form of the definitions to the editorial line of the ontology (if any) [5].

In addition to these ontology-specific recommendations, terminological manuals and guidelines state a number of other general principles and recommendations relating to definition writing [4], [6]–[8].⁵

VII. CONCLUSION

In this communication, we showed through examples that the defining practices in the ontology world lack systematic principles and theory. To fill this gap, we presented some background on textual definitions and axioms in ontologies from the terminologist’s and logician’s viewpoint, emphasizing their overlapping and complementary functions.

Based on a discussion of various kinds of correspondences between the parts of textual definitions and axioms, we put forward two primary recommendations to improve the contents of both textual definitions and axioms:

- Textual definitions and axioms should, whenever possible, represent the same content. As we hope our examples have indicated, it is frequently possible to do this with the resources of the ontology.
- Neither textual definitions nor axioms should include content that defines another term in the same ontology.

Finally, we proposed an implementable procedure to help systematize content-checking of textual and logical definitions in ontologies.

VIII. ACKNOWLEDGMENT

Work on this paper was supported by the Swiss National Science Foundation (SNSF).

REFERENCES

- [1] S. Seppälä, “Contraintes sur la sélection des informations dans les définitions terminographiques: vers des modèles relationnels génériques pertinents,” Ph.D. dissertation, Département de traitement informatique multilingue (TIM), Faculté de traduction et d’interprétation, Université de Genève, 2012. [Online]. Available: <http://archive-ouverte.unige.ch/unige:21874>
- [2] —, “An ontological framework for modeling the contents of definitions,” *Terminology*, forthcoming.
- [3] D. Marconi, *Lexical Competence*. The MIT Press, 1997.
- [4] ISO 704:2009(E), *Terminology work – Principles and methods (ISO 704:2009)*, 3rd ed. Geneva: ISO, 2009.
- [5] S. Seppälä, “Semi-automatic checking of terminographic definitions,” in *International Workshop on Terminology design: quality criteria and evaluation methods (TermEval) — LREC 2006*, Genoa, Italy, 2006, pp. 22–27.

⁵For a list of (terminology) manuals that contain definition writing principles and recommendations, as well as other writings on definitions see <https://sites.google.com/site/definitionsportal/literature>.

- [6] S. Pavel and D. Nolet, *Handbook of Terminology*. Canada: Public Works and Government Services - Translation Bureau, 2001.
- [7] (2012) The pavel terminology tutorial. [Online]. Available: <http://www.bt-tb.tpsgc-pwgsc.gc.ca/btb.php?lang=eng&cont=308>
- [8] R. Vézina, X. Darras, J. Bédard, and M. Lapointe-Giguère, *La rédaction de définitions terminologiques*, ser. Version abrégée et adaptée par Jean Bédard et Xavier Darras. Montréal: Office québécois de la langue française, 2009.

Defining ‘sign’ and ‘symptom’

Alexander P. Cox^{*}, Patrick L. Ray, Mark Jensen, Alexander D. Diehl
State University of New York at Buffalo, Buffalo, NY, USA
^{*}apcox@buffalo.edu

Abstract—The terms ‘sign’ and ‘symptom’ have proven difficult to define and represent in a biomedical ontology. Medical professionals use ‘sign’ and ‘symptom’ to refer to medically relevant information about patients; however, they do not agree on the definitions. In particular, while medical professionals agree that there is an important distinction between signs and symptoms, they do not agree on the precise nature of this distinction. It is unsurprising then that attempts to provide ontological representations of these entities have repeatedly fallen short. As an added complication, a variety of entities—including material entities, qualities, and processes—may reasonably be understood as signs or symptoms. Thus, the ontological nature of a sign or symptom raises many questions about the meanings and proper use of these terms. We discuss specific challenges to defining ‘sign’ and ‘symptom’, identify essential features of these entities, explore the ontological implications of existing definitions, and propose our own definitions. We evaluate several competing ontological representations and present our proposed representation within the framework of the Ontology for General Medical Science. The proposed representation of sign and symptom is ontologically sound, provides precise definitions of each term, and enables users to easily create customized groups of signs and symptoms. Our experience highlights general issues about developing definitions in ontologies.

Keywords—*sign; symptom; definition; clinical finding; OGMS; ontology*

I. INTRODUCTION

Clinicians and other medical professionals regularly use the terms ‘sign’ and ‘symptom’ to refer to medically relevant information about patients. Yet, the use of these terms is often imprecise, inconsistent, or both. This is due, in part, to the tendency to use these terms loosely. For example, by broadly referring to both signs and symptoms as symptoms [1]. As a further complication, many medical texts—including those dedicated to the study of signs and symptoms—fail to provide even preliminary definitions for these terms [2, 3]. When definitions are provided, they are not always consistent with one another. See TABLE I for a list of definitions.

Comparison of lists of signs and symptoms that are presented in the absence of definitions reveals numerous potentially inconsistent applications of ‘sign’ and ‘symptom’. According to [2], examples of symptoms include: fatigue, dizziness, fever, headache, insomnia, lymphadenopathy, night sweats, muscle weakness, weight gain, weight loss, pain, nausea, bloating, itching, sore throat, hearing loss, diarrhea, constipation, confusion, memory loss, tremor, anxiety, cough, and jaundice. According to [1], examples of signs include:

jaundice, swollen joints, and cardiac murmurs. According to [4], examples of vital signs include: temperature, respirations, pulse, and blood pressure. Notice that jaundice appears on both a list of symptoms and on a list of signs. While some definitions of ‘sign’ and ‘symptom’ allow certain features of the patient to be both a sign and a symptom, others do not. Additionally, which features can be both a sign and a symptom can change based on which definition is used.

Representing sign and symptom in an ontology is an ideal means by which to enforce their precise definitions and encourage their consistent application. At the same time, it emphasizes the importance of these terms to the medical community. Our goal is to precisely define the terms ‘sign’ and ‘symptom’ and to provide sound ontological representations of these entities. In doing so, we hope that our experience will serve as a primer on some of the challenges involved in developing rigorous definitions in ontologies.

II. METHODS

There are theoretical concerns regarding definition formation that must be considered prior to an attempt to define a term or set of terms. Definition formation is goal-driven and, as such, there are certain desiderata for what typically constitutes a “good” definition. These desiderata often depend on the type of definition one is seeking to provide as well as on the field one is working in [5, 6]. Nonetheless, we can identify certain desiderata that should hold irrespective of these concerns. In general, definitions ought to be: a) sufficiently inclusive so as to include or capture all of the actual instances of their definiens, b) sufficiently exclusive so as to exclude or discount all of the instances that are not their definiens, and c) informative enough to impart information to the audience [7]. We acknowledge that many groups may require additional desiderata. The considerations listed here are minimal desiderata for definitions.

There is also an issue of conceptual priority underpinning our process. Since we acknowledge that there are general desiderata for definitions before we engage in analysis of the current literature, these concerns are conceptually prior to any considerations discovered in the process of evaluating existing efforts. We also acknowledge that there may emerge more desiderata for specific definitions or types of definitions as a result of the evaluation of a set of attempted definitions. These should also be considered when determining whether a definition is adequate. For example, if a definition meets the three initial desiderata listed above but is criticized for obscurity or inconsistency with dominant views expressed in the literature, then one should seek to find a consistent and

TABLE I. DEFINITIONS OF ‘SIGN’ AND ‘SYMPTOM’ FROM THE MEDICAL LITERATURE

Source	Definition of ‘sign’	Definition of ‘symptom’
DeGowin’s Diagnostic Examination [4]	Abnormalities detected by the examiner.	Abnormalities the patient perceives.
MedlinePlus www.merriam-webster.com	An objective evidence of disease especially as observed and interpreted by the physician rather than by the patient or lay observer.	Subjective evidence of disease or physical disturbance observed by the patient; broadly: something that indicates the presence of a physical disorder.
The Free Dictionary www.thefreedictionary.com	A body manifestation, usually detected on physical examination or through laboratory tests or x-rays, that indicates the presence of abnormality or disease.	A sign or an indication of disorder or disease, especially when experienced by an individual as a change from normal function, sensation, or appearance.
MediLexicon www.medilexicon.com	Any abnormality indicative of disease, discoverable on examination of the patient; an objective indication of disease.	Any morbid phenomenon or departure from the normal in structure, function, or sensation, experienced by the patient and indicative of disease; a subjective indication of disease.
MedicineNet www.medterms.com	Any objective evidence of disease. It is evidence that can be recognized by the patient, physician, nurse, or someone else.	Any subjective evidence of disease; only the patient can perceive them.
Towards an Ontology of Pain [20]		A restricted family of phenomena, which are of their nature experienced in the first person. Symptoms can be reported to, and associated behaviors and bodily qualities can be observed by, the clinician; but the symptoms themselves cannot be observed or objectively measured.
Signs and Symptoms: Applied Pathologic Physiology and Clinical Interpretation [1]	Signs are detectable by another person and sometimes by the patient himself.	(a) As broadly and generally employed, the word symptom is used to name any manifestation of disease. (b) Strictly speaking, symptoms are subjective, apparent only to the affected person. (c) In ordinary clinical usage, the term symptom refers to what the patient experiences and reports as manifestations of illness. Thus, symptoms are subjective (psychological) in the sense that the patient can report only that of which he is aware.
Textbook of Diagnostic Medicine [24]	As opposed to revealing symptoms, physical examination reveals information that is comparatively more objective, measurable, and reproducible.	Symptoms are clinical manifestations of the disorder of organs or systems as experienced by patients. Symptoms are subjective and often difficult to quantify.
Rational Diagnosis and Treatment: Evidence-based Clinical Decision-making [25]	(a) Physical signs comprise all those observations which are made by the doctor during the physical examination. (b) Some of the recorded ‘signs’ fall into a special group: provoked symptoms. They are subjective symptoms which are only noticed by the patient during the physical examination.	(a) Subjective symptoms are the sensations noted by the patient and the patient’s mood. (b) Objective symptoms are observations made by the patient or the relatives concerning the patient’s body and its products.
The Mosby Medical Encyclopedia [26]	Something seen by an examiner. Many signs go along with symptoms, as bumps and rashes are often seen when a patient complains of itching.	Something felt or noticed by the patient that can help to detect a disease or disorder.

non-obscure definition; thus adding to our initial set of desiderata. Considerations such as these are not universal for definitions as they are relative to a community or sub-community. In contrast to general desiderata, let us call these subject-specific desiderata. Both general and subject-specific desiderata should be considered equally when determining the success of a definition or set of definitions [8].

Sources for the definitions of ‘sign’ and ‘symptom’ were gathered by performing a literature review. The literature review drew primarily from medical dictionaries and medical texts—especially those texts whose asserted focus is on signs or symptoms. These texts typically discussed the diagnostic process, clinical encounters, identification of diseases, or lists of signs and symptoms based on their relative importance and possible etiology. It is notable that many texts failed to provide definitions of either ‘sign’ or ‘symptom’ and thereby implicitly presumed familiarity on the part of their readers with the meanings of these terms. We compiled a list of available definitions and present a representative selection in TABLE I.

Biomedical ontologies that represent signs or symptoms were identified by performing queries in BioPortal for the terms ‘sign’ and ‘symptom’ [9]. Each search result was screened to identify and eliminate ontologies that returned inappropriate matches. The remaining results were reviewed to

identify and set aside ontologies that reused the relevant term from another ontology. Finally, we recorded the representations and definitions of sign and symptom in each remaining ontology. TABLE II displays the pertinent information for each ontology that provides a unique definition for at least one of these terms.

At the time of our research, querying the term ‘symptom’ returned 30 results in BioPortal. 9 results were screened out as irrelevant to our project. Of the remaining 21 results, 8 projects reused the symptom class from another ontology. Of the remaining 13 ontologies, only 6 provide a definition of ‘symptom’. 2 projects, the Translational Medicine Ontology (TMO) and the Ontology for General Medical Science (OGMS), use the same source and therefore give identical definitions [10]. This leaves 5 ontologies that uniquely define ‘symptom’. Querying the term ‘sign’ returned 21 results in BioPortal. 8 results were screened out as irrelevant to our project. Of the remaining 13 results, 5 projects reused the sign class from another ontology. Of the remaining 8 ontologies, only 4 provide a definition of ‘sign’. Again, TMO and OGMS give identical definitions. This leaves 3 ontologies that uniquely define ‘sign’.

Of the 8 reuses of ‘symptom’ and 5 reuses of ‘sign’, OGMS: ‘symptom’ is reused by 5 ontologies and OGMS: ‘sign’

TABLE II. DEFINITIONS OF ‘SIGN’ AND ‘SYMPTOM’ FROM ONTOLOGIES IN BIOPORTAL

Ontology	Term	Definition	Parent Class
Ontology for General Medical Science (OGMS)	symptom	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.	entity
	sign	A 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.	entity
National Cancer Institute Thesaurus (NCIT)	Symptom	Subjective evidence of disease perceived by the patient.	Sign or Symptom
	Sign	Objective evidence of disease perceptible to the examining healthcare provider.	Sign or Symptom
International Classification for Nursing Practice (ICNP)	Symptom	Phenomenon: Change in the body, subjective experience of change in bodily sensation, function or appearance.	Phenomenon
	Sign	N/A	Phenomenon
Symptom Ontology (SYMP)	symptom	A symptom is a perceived change in function, sensation, loss, disturbance or appearance reported by a patient indicative of a disease.	(Root Term)
Computer Retrieval of Information on Scientific Projects Thesaurus (CRISP)	sign/symptom	Clinical manifestations that can be either objective when observed by a physician, or subjective when perceived by the patient.	pathology

is reused by 3 ontologies, which makes OGMS the most widely reused source of both classes. The Systematized Nomenclature of Medicine – Clinical Terms (SNOMED CT) has the second most reuses. SNOMED CT boasts a massive medical terminology with over 300,000 classes and is designed for the primary purpose of improving Electronic Health Records (EHRs) [11]. In contrast, OGMS is a small mid-level ontology that is compliant with the Basic Formal Ontology (BFO) and is designed to be easily imported and used by other biomedical ontologies [12, 13]. While SNOMED CT currently has more end users, there are reasons to doubt that it has the logical capacity to meaningfully assist in automated reasoning over its classes [14]. Thus, OGMS’s versatility and compatibility with other biomedical ontologies makes it better suited to enable term reuse and is the best candidate ontology for hosting the representations of sign and symptom. For these reasons, we focus on the representation of these entities within the OGMS framework.

Following OGMS and BFO, we employ the methodology of ontological realism in developing our representations of sign and symptom [15]. According to ontological realism, when developing an ontology, the goal is to identify the sorts of entities that exist in reality and then represent them according to the best current scientific understanding. We are committed to upholding the OBO (Open Biological and Biomedical Ontologies) Foundry principles for best practices in ontology development [8]. In particular, we adhere to the principles of avoiding redundancy, exploiting compositionality, and using common architecture [16, 17]. The existence of at least 13 distinct representations of symptom and 8 distinct representations of sign in ontologies available through BioPortal creates redundancy and multiple architectures. Making OGMS the sole host of sign and symptom respects these OBO Foundry principles. Our proposed representations exploit compositionality by using existing terms from multiple ontologies to define ‘sign’ and ‘symptom’.

III. RESULTS

Examination of particular signs and symptoms reveals that, taken as a whole, they are not instances of a single universal. That is, sign and symptom are not natural kinds. Rather, instances of each group are comprised of a variety of types of entities including material entities, processual entities, and

qualities. Adherence to ontological realism therefore requires that sign and symptom not be asserted as named universal classes in an ontology.

Our solution is to introduce relations to connect entities that can be a sign, symptom, or, in some cases, both to the diseases, disorders, or syndromes that they are a sign or symptom of. Given the frequent use of the terms ‘sign’ and ‘symptom’ in non-clinical settings, we chose to use the terms ‘clinical sign’ and ‘clinical symptom’. In addition to reducing confusion, the use of specialized terms emphasizes the need for specialized definitions and can reduce objections to the definitions’ potentially counter-intuitive entailments. Hence, we propose the relations ‘is clinical sign of’ and ‘is clinical symptom of’ as subtypes of the Information Artifact Ontology’s ‘is about’ object property, which relates an information artifact to an entity. We define these relations as follows:

is clinical symptom of =_{df} X is a symptom of Y if and only if: (i) X is a clinical finding about a patient that is reported by a patient, family member, caretaker, or other non-medical professional; (ii) Y is a disease, disorder, or syndrome; and (iii) X is hypothesized by a clinician to be of clinical significance to Y.

is clinical sign of =_{df} X is a sign of Y if and only if: (i) X is a clinical finding about a patient that is observed by a clinician or reported by another medical professional; (ii) Y is a disease, disorder, or syndrome; and (iii) X is hypothesized by a clinician to be of clinical significance to Y.

While we contend that these relations most accurately represent the meanings of ‘sign’ and ‘symptom’, users may find it desirable to have named classes. Named classes make it easier to annotate terms and to identify and compose lists of entities of interest. Adoption of our relational approach does not necessitate a loss of functionality. Anonymous defined classes (ADCs) can be created for this purpose [18, 19]. Unlike a named class, an ADC need not represent a natural kind. Thus, ADCs can be constructed to represent just those entities that ontology users are interested in. For example, if a user wants to query her ontology for a list of all clinical signs, she can create an anonymous class defined as (‘clinical finding’ and (‘is clinical sign of’ some (disease or disorder or syndrome))).

This approach can be used to generate lists of signs, symptoms, or both that are hypothesized to be of significance to specific diseases, disorders, or syndromes. For example, a user who is only interested in symptoms of cardiovascular disease can create an anonymous class defined as (‘clinical finding’ and (‘is clinical symptom of’ some ‘cardiovascular disease’)). If an ADC is of particular value to the user, it can be given a name—such as ‘clinical sign’ or ‘clinical symptom of cardiovascular disease’. Naming an ADC produces a named defined class. Although a named defined class need not be a universal type, users can interact with it in much the same way that they interact with named universal classes. In this way, our representations of sign and symptom can accommodate the diversity of users’ needs.

IV. DISCUSSION

A. Defining ‘sign’ and ‘symptom’

The definitions in TABLE I suggest the adoption of one of the following criteria for distinguishing signs from symptoms:

1. Who reported or observed the phenomenon.
2. Whether the patient or the clinician reported or observed the phenomenon.
3. Who is capable, at least in theory, of observing or experiencing the phenomenon.

The first distinction can, but need not, allow persons other than the patient to observe and report the patient’s symptoms. The second distinction limits symptoms to only those things the patient observes and reports. Both distinctions allow certain features of patients to be both signs and symptoms.

“The distinction between symptoms and signs is frequently unclear. For instance, jaundice may be a symptom that brings the patient to the physician, but it is also a sign visible to the clinician. [...] Vomiting, although it can be witnessed, is more often a symptom, while tenderness, although it may be noted by the patient, is a sign that can be elicited by the examiner.” [4]

The third distinction makes a stronger claim. According to this distinction, signs can, at least in theory, be observed by more than one person, but symptoms can only ever be observed by the patient [20]. Thus, nothing can be both a sign and a symptom. What is essential is who *could* have observed the feature, not who *actually* observed or reported it.

Yet, an historically compelling reason for creating and continuing to use the sign/symptom distinction is that observations made by medical professionals are, as a whole, typically considered to be more reliable than reports made by the patient, a family member, or someone who is not trained in medicine [4]. Thus, while the third distinction is *prima facie* ontologically superior because it does not allow the same feature of the patient to be both a sign and a symptom, it fails to account for the primary motivation for making the distinction. More significantly, the third account relies on a distinction between objectivity and subjectivity that may be metaphysically untenable.

Consider pain. Pain is arguably the archetypical symptom because, while people can observe behavioral cues and then infer that another individual is in pain, only that individual can definitively say whether he or she is experiencing pain. Yet, neuroscientists have made incredible progress both in imaging

the human brain and in mapping specific functions to specific areas of the brain [21]. In some cases, such as neurons in the hippocampus called place cells, precise locations of specific memories have been identified [22, 23]. Thus, it is becoming increasingly plausible that neuroscientists will eventually be able to objectively observe pain and other features of the patient. If this is possible, then, according to the third distinction, pain and other archetypical symptoms are—and always have been—signs. For this reason, we reject the third distinction in favor of an account of sign and symptom based on who reported the feature.

This leaves either the first or the second proposal. The second distinction is more restrictive since only the patient can report a symptom. If signs are similarly restricted to reports made by clinicians, then observations reported by a family member, caretaker, or other non-clinician fall outside the range of signs and symptoms. One implication of this is that, while a parent can report observations about his or her child and a doctor can use these reports to aid in diagnosing the child, a parent cannot report his or her child’s symptoms. Rejecting the second distinction and allowing non-clinicians to report symptoms avoids this oddity while preserving the initial motivation for the sign/symptom distinction. On the resulting view, symptoms are reports about the patient’s health made by a non-clinician; signs are reports about the patient’s health made by a clinician. This can be refined to allow reports made by certain non-clinicians, namely those persons who play related medical roles, to report signs. Identifying what these roles are, who has them, and in what settings they are realized are important issues that we set aside for the purposes of this paper.

Having distinguished signs from symptoms, it remains to distinguish them from other entities. We contend that an essential criterion of both signs and symptoms is that they be hypothesized to be clinically significant. A competing view is that signs and symptoms are clinically significant regardless whether anyone ever hypothesizes them to be so. We reject this view because we take signs and symptoms to have an important epistemic component. That is, something cannot be a sign or symptom unless it is known by someone. For example, a genetic mutation may be the material basis of a particular genetic disease, but it is not a sign of that disease until a test has detected the presence of the mutation and a qualified professional has interpreted the test results. Prior to that, the genetic mutation is simply a disorder. The epistemic component of signs and symptoms is due to the social construction of clinical settings. Determining clinical significance requires interpretation by clinicians. Furthermore, clinicians use signs and symptoms as part of the diagnostic process—the goal of which is to arrive at a diagnosis, which is a hypothesis about the patient’s health. Hence, it would be a mistake to divorce signs and symptoms from their clinical interpretation.

It is a further question whether the role of the person who formulates the hypothesis of clinical significance matters. There are three plausible answers:

- (i) It does not matter who hypothesizes the feature to be of clinical significance as long as someone does.
- (ii) It only matters whether the person who reported the feature hypothesizes that it is clinically significant.

- (iii) It only matters whether the clinician hypothesizes that the feature is clinically significant.

We reject (ii) because it entails that, if a patient reports something but fails to postulate that it is important, it is not a symptom. This is true even if the clinician correctly identifies the reported feature as important. We reject (i) because it permits too many things to be signs or symptoms. For example, any observation a clinician makes about a patient, regardless of its relevance to the patient's health, can become a sign simply because another person hypothesizes that it is clinically significant.

Option (iii) has its own potentially counter-intuitive consequences because it ignores patients' hypotheses. This entails that only reports made within a clinical setting can be signs or symptoms. Nonetheless, we endorse (iii) for several reasons. First, clinicians are in a privileged position to identify which features of a patient are clinically significant. Their knowledge and experience prevents a lot of irrelevant information from being misidentified as significant and limits the likelihood that something significant will be overlooked.

Second, the social nature of signs and symptoms is important. The clinician role is a special social entity that endows its bearer with the power to medically diagnose patients within a clinical setting. This is similar to how only a judge has the authority to sentence a defendant within an appropriate legal setting. Furthermore, since signs and symptoms are used to diagnose patients and determine treatment plans, they are only needed within a clinical setting. This does not, however, prevent people from using the terms 'sign' and 'symptom' in a very broad manner to refer to any number of things; however, the general application of these terms is technically incorrect and any meaning that is conveyed is derivative of their proper clinical usage. The prevalence of non-clinical applications of 'sign' and 'symptom' is ample reason to prefer the use of 'clinical sign' and 'clinical symptom' in order to avoid confusions of this sort. Once the terminological confusion is eliminated and the importance of the clinical setting is emphasized, we contend that the initial counter-intuitiveness of (iii) becomes negligible. Thus, we conclude that a health feature of a patient is only a sign or symptom if it is hypothesized by a clinician to be of clinical significance.

B. Representing Sign and Symptom

Recall from TABLE II that the Ontology for General Medical Science (OGMS) defines 'sign' as "A 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." OGMS defines 'symptom' 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."

These definitions raise several issues. First, they allow material entities to be signs but not symptoms. If this is due to acceptance of the subjective/objective distinction, it has not been fully implemented because these definitions are consistent with a quality or processual entity being both a sign and a symptom. Yet, if the subjective/objective distinction is not

being employed, it is unclear why material entities, such as a rash or abnormal lump, cannot be symptoms. Second, it is not explicit whether being "deemed... to be of clinical significance" is the same as being "hypothesized... to be the realization of a disease". Third, OGMS is built using the Basic Formal Ontology (BFO), which states that qualities are not realizable entities. So OGMS's definition of 'symptom' is incorrect. Finally, and most significantly, these definitions combine fundamentally different types of entities. Qualities are dependent continuants, material entities are independent continuants, and processes are occurrents. As a result, these classes do not fit within BFO's representational structure. Hence, they are defined classes and are represented as direct subtypes of 'entity'.

The current representations of sign and symptom in OGMS limits what can be axiomatically asserted of these classes because anything that is asserted must hold for qualities, material entities, and processes. This means that not even the most fundamental relations, for example 'inheres in', 'bearer of', or 'realizes', can be asserted of either class. While this does not prevent simple annotation using these terms and these relations can still be asserted at the instance level, it severely limits the automatic reasoning power of any system that uses these OGMS terms. This undermines one of the major advantages of using an ontology. The problem is compounded because the meaning of many other OGMS terms depends on a clear account of sign and symptom. These include: syndrome, treatment, acute disease course, clinical picture, and clinical history.

Before presenting the reasoning for our representations of sign and symptom, we present four alternative representations and briefly discuss why each one should be rejected. First, eliminate 'sign' and 'symptom' from OGMS. Everything that is currently a sign or symptom could instead be represented as a clinical finding. This would require the redefinition of other OGMS terms that explicitly refer to signs and symptoms, which might lead to further difficulties. More importantly, it is highly unlikely that the medical community would accept the elimination of 'sign' and 'symptom'. So, even if the distinguishing characteristics of signs and symptoms were incorporated in OGMS using logical definitions to preserve important information about these clinical findings, this representation would fail to satisfy the desires of the ontology's intended user base. Nonetheless, of the four alternatives discussed here, this solution is ontologically superior because, unlike the others, it is ontologically self-consistent. Readers who are ultimately left with the sense that 'sign' and 'symptom' are overly confused or possibly indefinable, may be inclined to endorse this solution.

Second, make 'sign' and 'symptom' roles. These roles may be played either by clinical findings or by qualities, processes, or material entities. Both representations fail because BFO does not permit qualities, processes, or dependent continuants to be the bearers of roles. Even if these entities were permitted to bear roles, this solution would create the logistical challenge of constructing a particular role for each disease, syndrome, and disorder. It is not sufficient to simply create the roles 'sign of' and 'symptom of' because each role is specific to the particular disease, disorder, or syndrome it is a sign or symptom of. Thus, the ontology would have to include

thousands of roles (e.g., ‘sign of Alzheimers disease’, ‘sign of heart attack’, ‘sign of influenza’, etc.), which is not an ontologically parsimonious solution.

Third, make ‘sign’ and ‘symptom’ subtypes of ‘clinical finding’. Yet, a clinical finding becomes a sign or symptom once it has been hypothesized to be of clinical significance to a particular disease, disorder, or syndrome. Thus, this solution permits clinical findings to shift their type simply because a clinician makes a hypothesis about it. This sort of type shifting is especially ontologically vicious because the “change” that occurs involves no change in the clinical finding itself. While a role can be acquired or lost without a corresponding change in its bearer, gaining or losing a role does not change the type of entity that its bearer is. Thus, this solution should be rejected.

Fourth, make ‘sign’ and ‘symptom’ relations between qualities, processes, or material entities and the diseases, syndromes, or disorders they are hypothesized to be of clinical relevance to. It is unclear that these relations are needed since more explicit relations already exist to connect these entities to their respective diseases, disorders, or syndromes. Pathological processes, such as tremors, are part of the disease course that realizes the disease. Pathological qualities, such as an elevated temperature, inhere in the patient as a result of certain pathological processes. Pathological material entities are part of the patient and can be a manifestation of the disease, such as a rash, or part of its material basis, such as neurofibrillary tangles. Furthermore, this solution is incompatible with the absence of a feature being a sign or symptom. For example, hyporeflexia, the lack of a deep tendon reflex, can be a sign of neuromuscular disease. Thus, material entities, qualities, and processes do not exhaust the domain of signs and symptoms.

Our solution is to represent sign and symptom as relations between clinical findings and the illnesses they are hypothesized to be of clinical relevance to. The result is X ‘is sign of’ Y and X ‘is symptom of’ Y where the domain X is a clinical finding and the range Y is a disease, disorder, or syndrome. These relations specify the nature of aboutness that holds between certain clinical findings and certain diseases, disorders, and syndromes. Which relationship is used depends on the role played by the person who reported the finding. Clinical findings reported by the patient, the patient’s family, or another non-clinician are potential symptoms. Clinical findings reported by a clinician are potential signs. In both cases, only findings hypothesized by a clinician to be of clinical significance to a disease, disorder, or syndrome will have one of these relations.

While laboratory tests, imaging techniques, and other medical procedures can provide diagnostically valuable clinical findings, they often are not performed by a clinician. Thus, it is necessary to allow the medical professionals who perform these procedures to report findings that may be hypothesized by a clinician to be signs. Additionally, while patients and non-clinician medical professionals must report their observations in order for them to be symptoms or signs, observations made by a clinician do not need to be reported in order to be signs. This is because a clinician must be informed about observations made by others in order to hypothesize that they are clinically significant, but does not need to report his own findings in order to hypothesize about them. If the clinician does not report his finding, the clinical finding is the clinician’s

mental representation. Thus, both features of the patient observed by a clinician and clinical findings about the patient that are reported by a medical professional can be signs.

Note that our representation is capable of handling cases where nonexistent entities are signs or symptoms. While there are no nonexistent entities, there can be a clinical finding about a feature that is not present. This clinical finding can then be hypothesized to be of clinical significance. In the case of hyporeflexia, the clinical finding would be the observation or report that no reflex occurred.

One might object that, unlike the subjective/objective distinction for sign and symptom, our representation fails because it permits a single feature of a patient to be both a sign and a symptom. If this were the case, it would mean that our definitions are too inclusive. This could lead to confusion and violate ontology best practices. On our account, the same clinical finding cannot have both the ‘is clinical sign of’ and the ‘is clinical symptom of’ relations. This is because only those clinical findings that are reported by a patient or non-medical professional can have the ‘is clinical symptom of’ relation. Similarly, only those clinical findings that are observed by a clinician or reported by an appropriate medical professional can have the ‘is clinical sign of’ relation. Thus, while there can be two findings about the same feature of a particular patient, no single finding can be both a sign and a symptom.

What happens if the patient or family member who reports a clinical finding is a clinician? Can such reports be both a sign and a symptom? The answer depends on which conditions one accepts for the realization of a clinician role. It is reasonable to assert that a clinician role can only be realized in the context of a clinical encounter. It is a further question whether an individual can play both a patient role and a clinician role in a single clinical encounter. Since the clinician role is a social construct, limitations—such as prohibiting a doctor from diagnosing or treating himself—can easily be asserted to resolve this dilemma. Another solution would be to allow clinicians to self-diagnose, but assert that the clinician role takes priority over the patient role with regard to clinical findings. Thus, clinical findings made during these encounters would always be either a clinical sign or just a clinical finding. The precise explication of this scenario is left open for further debate.

Finally, one might object that our proposed definitions are overly strict because they exclude prognostic signs from being clinical signs. Prognostic signs are signs that are indicative of the patient’s health outcome. These signs assist clinicians in determining a patient’s likelihood of survival, recovery time, or possible loss of physical ability or mental functioning. This is opposed to diagnostic signs, which are indicative of the nature of the patient’s illness. If—as our definition of ‘is clinical sign of’ requires—some prognostic signs are not about a disease, disorder, or syndrome, then not all prognostic signs are clinical signs and our definition is too restrictive.

There are several things to consider here. First, even if prognostic signs cannot always be understood as clinical signs, this may be due to prognostic signs and diagnostic signs being distinct types of signs. If this is the case, then the mistake may lie in grouping two distinct kinds of clinical findings together

as a single thing. Second, our definition permits clinical signs to be signs of disorders or syndromes as well as of diseases. It is plausible that most, if not all, prognostic signs are signs of disorders. For example, a death rattle is a prognostic sign of imminent death, but it is also a clinical sign of the buildup of fluid in the throat and upper chest, which can be understood as a syndrome and is the result of a disorder. Similarly, a clinician may determine that a gunshot victim will make a full recovery based on observing that the bullet missed all major organs and arteries. The wound is a disorder and it is based on this clinical finding that the clinician is able to make a prognosis. Thus, according to our definition, all prognostic signs are clinical signs.

V. CONCLUSIONS

We have presented an ontologically sound representation of sign and symptom and developed precise definitions for relations that capture the meaning of each term. This is important for the biomedical community because it unifies the representation of two commonly used terms while providing a clear delineation of their instances that does not allow for confusing overlap between their members. Furthermore, our representation enables the easy formulaic creation of customized groups of signs and symptoms in order to identify information relevant to each user's needs. This is, perhaps, the most significant contribution our work provides to the biomedical ontologies community.

Our experience in developing this account of signs and symptoms is indicative of general issues that can arise when developing definitions in ontologies. For example, when the ontological representation requires a more restrictive definition than the colloquial definition, it is advisable to create a special label for the entity (e.g., 'clinical sign' instead of 'sign' and 'clinical symptom' instead of 'symptom'). Changing the label reduces the risk of confusion as well as the risk that the specialized definition will elicit resistance from users familiar with the old term. Ontology development is typically a descriptive exercise in representing entities such that the ontology is made to conform to our understanding of the world; however, our experience here has shown that the direction of fit can operate in reverse. This occurs when the only ontologically sound means of representing the entities in question requires changing our everyday understanding of the meaning of those terms. In these cases, the ontological definition should be used to prescriptively enforce a new, more precise, use of the term.

ACKNOWLEDGMENTS

The authors thank Travis Allen and Isaac Berger for discussion of the manuscript; the members of the OGMS community, especially Albert Goldfain and Barry Smith for general discussion; and the participants of the OGMS workshop at the 2013 International Conference on Biomedical Ontology for feedback on an earlier version of this proposal.

REFERENCES

[1] C. M. MacBryde, R. S. Blacklow, and R. D. Aach, Signs and symptoms: applied pathologic physiology and clinical interpretation. Philadelphia: Lippincott, 1970.

[2] M. C. Henderson, L. M. Tierney, Jr, and G. W. Smetana, The patient history: an evidence-based approach to differential diagnosis. New York; London: McGraw-Hill Medical, 2012.

[3] S. D. C. Stern, A. S. Cifu, and D. Altkorn, Symptom to diagnosis: an evidence-based guide. New York: McGraw-Hill Medical, 2010.

[4] R. F. LeBlond, R. L. DeGowin, and D. D. Brown, DeGowin's diagnostic examination. New York: McGraw-Hill Medical, 2009.

[5] J. J. Cimino, "Desiderata for controlled medical vocabularies in the twenty-first century," *Methods Inf Med.* 37(4-5), Nov 1998: pp. 394-403.

[6] J. Michael, J. L. Mejino, Jr, and C. Rosse, "The role of definitions in biomedical concept representation," in *Proceedings of the AMIA Symp.* 2001: pp. 463-467.

[7] B. Smith, "Introduction to the logic of definitions," in *Proceedings of the 2013 International Conference on Biomedical Ontologies Workshop on Definitions in Ontologies.* 7 July 2013.

[8] B. Smith, et al., "The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration," *Nature Biotechnology*, vol. 25, 2007: pp. 1251-1255.

[9] P. L. Whetzel, et al., "BioPortal: enhanced functionality via new web services from the National Center for Biomedical Ontology to access and use ontologies in software applications," *Nucleic Acids Research*, July 1, 2011; 39(Web Server issue): W541-W545.

[10] R. H. Scheuermann, W. Ceusters, and B. Smith, "Toward an ontological treatment of disease and diagnosis," in *Proceedings of the 2009 AMIA Summit on Translational Bioinformatics*, San Francisco, California; 2009: pp. 116-120.

[11] Systematized Nomenclature of Medicine - Clinical Terms (SNOMED CT). [<http://www.ihtsdo.org/snomed-ct/>]

[12] Ontology for General Medical Science (OGMS). [<https://code.google.com/p/ogms/>]

[13] Basic Formal Ontology (BFO). [<http://www.ifomis.uni-saarland.de/bfo/>]

[14] G. Héja, G. Surján, and P. Varga, "Ontological analysis of SNOMED CT," *MC Medical Informatics and Decision Making*, 2008, 8(Suppl 1): S8.

[15] B. Smith and W. Ceusters, "Ontological realism as a methodology for coordinated evolution of scientific ontologies," *Applied Ontology* 5: 2010, pp. 139-188.

[16] B. Smith, "Ontology (Science)," in *Formal Ontology in Information Systems - Proceedings of the Fifth International Conference (FOIS 2008)*. Edited by: C. Eschenbach and M. Grüninger, Amsterdam: IOS Press, *Frontiers in Artificial Intelligence and Applications*; 2008: pp. 21-35.

[17] B. Smith, et al., "Relations in biomedical ontologies," *Genome Biology*, 6, 2005: R46.

[18] Protege 4.x Anonymous Classes. [<http://protegewiki.stanford.edu/wiki/P4AnonymousClasses>]

[19] OWL Web Ontology Language Reference. [<http://www.w3.org/TR/owl-ref/>]

[20] B. Smith, W. Ceusters, L. J. Goldberg, and R. Ohrbach, "Towards an ontology of pain," in *Proceedings of the Conference on Ontology and Analytical Metaphysics*; Keio University Press, 2011.

[21] K. Deisseroth and M. J. Schnitzer, "Engineering approaches to illuminating brain structure and dynamics," *Neuron*. Vol. 80, Issue 3, 30 October 2013, pp. 568-577.

[22] Xu Liu, et al., "Optogenetic stimulation of a hippocampal engram activates fear memory recall," *Nature*. 484: 19 April 2012, pp. 381-385.

[23] J. O'Keefe and J. Dostrovsky, "The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat," *Brain Research*. Vol. 34, Issue 1, November 1971, pp. 171-175.

[24] A. H. Samiy, R. G. Douglas, and J. A. Barondess, *Textbook of diagnostic medicine*. Philadelphia: Lea & Febiger, 1987.

[25] H. R. Wulff and P. C. Götzsche, *Rational diagnosis and treatment: evidence-based clinical decision-making*. Oxford; Malden, MA, USA: Blackwell Science, 2000.

[26] W. D. Glanze, K. Anderson, and L. E. Anderson, *The Mosby medical encyclopedia*. New York: Plume, 1992.

An alternative terminology for pain assessment

Werner Ceusters
Department of Biomedical Informatics
University at Buffalo
Buffalo, NY 14203, USA
Email: ceusters@buffalo.edu

Abstract — Background: the International Association for the Study of Pain (IASP) publishes since 1986 a relatively frequently updated list of pain terms with corresponding definitions and clarificatory notes currently known as the ‘IASP Taxonomy’. The last update, i.e. the May 2012 version of this taxonomy, was subjected to an analysis with the goal to assess whether the definitions of the IASP terms that are used to describe findings of somatosensory testing and pain assessment satisfy the conditions for these terms to become part of a realism-based ontology. **Results:** the taxonomy was found to be built on definitions that are not in every case based on necessary and sufficient conditions, nor satisfy the single inheritance principle for realism-based ontologies. Furthermore, although the documentation about introduced changes provided by the IASP makes it clear that the terminology authors tried to solve ambiguities and unclarity present in previous versions, they did not succeed completely and introduced even some inconsistencies. The analysis demonstrates that the main cause for this is not the choice of differentiating characteristics, but rather insufficient attention to the wide variability in stimulus/response combinations that these characteristics reveal. **Conclusions:** the IASP taxonomy is not fit to form the basis for a realism-based ontology. A new representation framework for describing pain assessment findings more accurately using the same set of differentiae is proposed and its correspondence with the traditional terminology explained.

Keywords—*pain terminology, ontological realism*

I. BACKGROUND

The Ontology for Pain-Related Mental Health and Quality of Life (OPMQoL) is being developed as part of the NIDCR-funded project R01DE021917 with the goal to integrate five datasets gathered in four different countries from patients suffering from one or other form of orofacial pain [1, 2]. Part of the data in these datasets describe findings that are based on the various kinds of responses that patients may report when subjected to stimuli to test their somatosensory status and that are typically described using terms such as ‘allodynia’, ‘hyperesthesia’, and so forth. Although these terms were already in practice since at least the early 19th century [3], standard definitions for these terms were first proposed in 1979 [4] and are since then regularly updated by the International Association for the Study of Pain (IASP), in print for the last time in 1994 [5], with more regular electronic updates on the IASP webpage [6] the last one in May 2012 (Table 1). These definitions are further based on the IASP definition for ‘pain’ as ‘an

unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage’.

For terms to be eligible as representational units in a realism-based ontology such as OPMQoL, they must not only (1) denote entities that can be classified following the principles of Ontological Realism [7], but also (2) be defined using Aristotelian definitions which specify the necessary and sufficient conditions for class membership, and further lead to a taxonomy based on single inheritance [8]. The goal of the work reported on here was to assess the adherence of the IASP pain assessment definitions to this second condition and to find ways for remediation if non-compliance was found.

Allodynia: pain due to a stimulus that does not normally provoke pain. Note: The stimulus leads to an unexpectedly painful response.
Analgesia: absence of pain in response to stimulation which would normally be painful.
Dysesthesia: an unpleasant abnormal sensation, whether spontaneous or evoked. Note: Special cases of dysesthesia include hyperalgesia and allodynia.
Hyperalgesia: increased pain from a stimulus that normally provokes pain.
Hyperesthesia: increased sensitivity to stimulation, excluding the special senses.
Hyperpathia: a painful syndrome characterized by an abnormally painful reaction to a stimulus.
Hypoalgesia: diminished pain in response to a normally painful stimulus.
Hypoesthesia: decreased sensitivity to stimulation, excluding the special senses.
Paresthesia: an abnormal sensation, whether spontaneous or evoked. Note: paresthesia is to be used to describe an abnormal sensation that is not unpleasant.

Table 1 - Pain terms analyzed

II. METHODS

Based on the definitions of the terms studied – note that table 1 contains only part of the relevant notes and that the reader should for complete understanding of the analysis method consult reference [6] - an analysis framework was designed by introducing nine hierarchically organized variables reflecting the type of stimulus, the presence or absence of a response, and the type of response when present, when a patient is subjected to a pain assessment investigation. The allowed values for these variables were defined, depending on what the variable stands for, either on a nominal or ordinal scale (Table 2).

Acronym	Term (plus meaning)	N
CONC	Normal case	9
ALLO-D	allodynia (definition) : unexpected evoked pain	10
ALLO-N	allodynia (note) : unexpected more intense evoked pain	30
ANAL	analgesia : unexpected absence of evoked pain	40
DYS-E	evoked dysesthesia	80
DYS-EP	painful evoked dysesthesia	50
DYS-EU	non-painful evoked dysesthesia	30
HYPERA	hyperalgesia : unexpected more intense evoked pain	20
HYPERE	hyperesthesia = increased sensitivity to stimulation	81
HYPERE-I	unexpected more intense evoked sensation	42
HYPERE-IP	unexpected more intense evoked pain	20
HYPERE-IM	unexpected more intense evoked modus M	26
HYPERE-P	unexpected presence of evoked sensation	49
HYPERE-PU	unexpected evoked unpleasant sensation other than pain	30
HYPERE-PP	unexpected pain	10
HYPERE-PM	unexpected modus M	13
HYPERP	hyperpathia	30
HYPOALG	hypoalgesia	20
HYPOE	hypoesthesia = decreased sensitivity to stimulation	58
HYPOE-P	decreased sensitivity to pain stimulation	40
HYPOE-PL	less pain to pain stimulation	20
HYPOE-PA	non painful unpleasant response to pain stimulation	20
HYPOE-M	decreased sensitivity to modus M stimulation	26
HYPOE-BI	decreased sensitivity to both kinds of stimulation	8
PAR-D-E	evoked paresthesia (definition)	81
PAR-D-EP	painful evoked paresthesia	30
PAR-D-EU	non-painful unpleasant evoked paresthesia	30
PAR-D-EN	non-painful not unpleasant evoked paresthesia	39
PAR-N-E	evoked paresthesia (note)	19

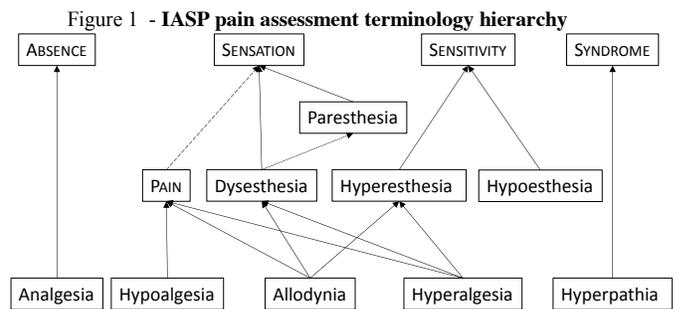
Table 5 - Terms and ontological subgroups for the IASP pain assessment terminology. Legend: N = number of stimulus/response combinations applicable (max = 130).

This step answers thus for each term pair ‘A B’ the question which and how many of the possible stimulus/response combinations can occur in the pair combinations A+/B+, A+/B-, A-/B+, A-/B- where ‘+’ and ‘-’ indicate that the stimulus/response combination can, resp. cannot occur under the definition of the term. As it became clear at this point that overlap was considerable, we designed a new terminology based on definitions that minimize the potential overlap using categories that are mutually exclusive. We then compared this new terminology with the traditional one, again using the stimulus/response combinations as benchmark.

III. RESULTS

A. The IASP terms do not satisfy the criteria for direct integration in a realism-based ontology.

Figure 1 - in which terms displayed in SMALL CAPS are the immediate superordinate terms found in the definitions and the arrows stand for the classical subsumption relation [10] – demonstrates that although the individual definitions follow the Aristotelian form ‘an A is a B which C’, the defined terms do not lead all together to a complete directed graph with an overarching top, not even if all 29 IASP terms would be included. Furthermore, the terms ‘allodynia’ and ‘hyperalgesia’ have superordinate terms which under their standard meanings should represent disjointed classes: although sensation and sensitivity are certainly related,



nothing which is a kind of one can also be a kind of the other. In addition, already a superficial reading of these terms and accompanying notes reveals ambiguities and inconsistencies. The definition of ‘allodynia’, for instance, indicates that the term should be used for pain evoked after applying a stimulus which is below the normal pain threshold. The corresponding note however suggests that also a response on an above-threshold stimulus may count as such when the stimulus leads to more pain than expected. The note for ‘dysesthesia’, as many similar notes for other terms which for space reasons are not reproduced in Table 1 but can be found in reference [6], indicate that there is considerable overlap between the terms.

B. Traditional pain assessment terminology shows considerable overlap

All terms of Table 1 could be mapped to the stimulus/response combinations. Table 6 illustrates how the parent terms relate to each other in function of the stimulus/response combinations. The individual cells contain the counts for the overlap, if any. For example, the overlap cells between hyperesthesia and hypoalgesia show - surprisingly - that these two conditions do not exclude each other: 6 of the 130 combinations fall under both definitions, 14 are such that hypoalgesia is present without hyperesthesia, 75 have hyperesthesia without hypoalgesia, and 35 don’t exhibit either. An additional color coding is used to highlight the type of overlap: white indicates a symmetric overlap for all 4 types of co-occurrence as exemplified by the hyperesthesia/hypoalgesia pair; green indicates mutual exclusion of the positive occurrences, the other three colors indicate an asymmetric overlap. An ideal terminology would be such that the classes defined are mutually disjoint. For 12 (n) classes as is the case here, there are 66 possible overlaps ($n*(n-1)/2$) between any pair of these classes, not counting overlap of a class with itself. As displayed in Table 6, there is no overlap in only 2 cases of these 66: (1) for hyperpathia versus allodynia (taking the note into account), and (2) for hyperesthesia and paresthesia (when the note is not taken into account).

C. Novel terminology with less overlap

Table 7 provides an overview of the proposed terminology which uses 6 variables (*Response expectation*, *Main finding*, *Sensation expectation*, *Sensation intensity*, *Sensation mode*, and *Stimulation type*) that can take a number of values and which are strongly related to the variables and values used to design the analysis framework of the 130 stimulus/response combinations.

It leaves however still many interpretations open, such as whether the resulting sensation must be alien to the given stimulus - would an erotic feeling induced by providing a pressure stimulus to the hand count as such a non-unpleasant abnormal sensation? - or whether it may be special cases of hypo- and hyperesthesia.

These reflections provide at the same time explanations for the very high degree of overlap between the majority of the traditional terms (Table 6). There is of course a symmetric non-overlap for each category with each negation, but the only non-overlap between distinct categories is found for the pairs allodynia (taking the note into account) -hyperpathia and hyperesthesia-paresthesia (as defined, without the limiting note).

The proposed terminology shows a much more limited degree of overlap. This lesser degree of overlap is because the parameters have been chosen in such a way that a specific combination of values cannot count for a specific class in more than one way, a feature which is not exhibited by the traditional terminology.

A disadvantage of the terminology is that it is more verbose, but this is compensated by the ease by which it can be implemented in systems for structured electronic reporting and automatic assigning of the categories using single select choice lists for each variable.

V. CONCLUSION

It is demonstrated that the IASP terms do not satisfy the criteria for direct integration in a realism-based ontology. A new terminology for stimulus based pain and somatosensory status assessment is proposed which exhibits less shortcomings in terms of overlap than the traditional terminology. This is because in contrast to the traditional approach, this proposal does not underestimate the various stimulus/response combinations that may occur.

VI. COMPETING INTERESTS

None

VII. AUTHORS' CONTRIBUTIONS

All analyses as well as paper writing were done by the author.

VIII. ACKNOWLEDGEMENTS

The work described is funded in part by grant 1R01-DE021917-01A1 from the National Institute of Dental and Craniofacial Research. The content of the paper is solely the responsibility of the authors and does not necessarily represent the official views of the NIDCR or the NIH.

REFERENCES

- [1] Smith B, Ceusters W, Goldberg LJ, Ohrbach R. Towards an Ontology of Pain. In: Okada M, editor. Proceedings of the Conference on Logic and Ontology. Tokyo: Keio University Press; 2011. p. 23-32.
- [2] Ceusters W. An information artifact ontology perspective on data collections and associated representational artifacts. *Stud Health Technol Inform.* 2012;180:68-72.
- [3] Welch WH. Papers and addresses. Baltimore: The John Hopkins Press; 1822.
- [4] IASP Subcommittee on Taxonomy. Pain terms: a list with definitions and notes on usage. *Pain.* 1979;6(3):249-52.
- [5] Merskey H, Bogduk N. Classifications of chronic pain: Description of chronic pain syndromes and definition of pain terms. Report by the International Association for the Study of Pain Task Force on Taxonomy. Seattle: IASP Press; 1994.
- [6] International Association for the Study of Pain. IASP Pain Taxonomy. 2012; Available from: <http://www.iasp-pain.org/Content/NavigationMenu/GeneralResourceLinks/PainDefinitions/default.htm>.
- [7] Smith B, Ceusters W. Ontological Realism as a Methodology for Coordinated Evolution of Scientific Ontologies. *Applied Ontology.* 2010;5(3-4):139-88.
- [8] Smith B. Introduction to the Logic of Definitions. International Workshop on Definitions in Ontologies, organized in conjunction with the Fourth International Conference on Biomedical Ontology (ICBO). Montreal: CEUR; 2013. p. 1-2.
- [9] Turk DC, Melzack R. Handbook of pain assessment. 3rd ed. New York: Guilford Press; 2011.
- [10] Ceusters W. Pain assessment terminology in the NCBO BioPortal: evaluation and recommendations. Proceedings of the International Conference on Biomedical Ontology 2014. Houston, TX2014. p. (accepted).

Proceedings
of

THE
WORKSHOP
ON
STARTING AN OBI-BASED BIOBANK
ONTOLOGY
(OBIB 2014)

*Proceedings
edited by*

*Mathias Brochhausen
Melissa Haendel
Christian J. Stoeckert
Jie Zheng*

*OBIB 2014 was held in conjunction with the 5th International
Conference on Biomedical Ontology (ICBO 2014)
Houston, TX, USA
October 6-7, 2014*

Preface

Biobanks are a critical resource for translational science. However, they often suffer from a lack of semantics in their ability to disseminate data and make it readily queryable. The OBO Foundry provides ontologies that are relevant to representing the structure and function of biobanks. The workshop organizers developed an ontology for biobanking based on the Ontology for Biomedical Investigations (OBI), an OBO Foundry ontology, aiming to support consistent annotation of biobank resources and provide semantics to biobank resources. The workshop focuses on problems and requirements of an ontology for biobanking, as well as potential solutions. Following topics are covered:

- Information models describing biobanks, their strengths, weaknesses, and competency criteria
- Reports of biobank related ontologies
- Case studies of applying ontologies to biobanks

The biobank workshop provides a platform to people from the biobank world who are interested in bringing ontologies into this effort. Over 30 people attended the workshop and initiated the potential collaboration following the workshop.

Acknowledgements

We would like to thank the keynote speaker, Dave Parrish, who gave an excellent talk on The Bio-Specimen Knowledge Bank. We would also like to thank the ICBO organizers for hosting the workshop.

Organizers of Biobank workshop ((in alphabetical order):

Mathias Brochhausen, University of Arkansas for Medical Sciences, Little Rock, AR

Melissa Haendel, Oregon Health & Science University, Portland, OR

Christian Stoeckert, Perelman School of Medicine University of Pennsylvania, Philadelphia, PA

Jie Zheng, Perelman School of Medicine University of Pennsylvania, Philadelphia, PA

Penn Medicine Biobank Informatics

OBI Influenced Software Design

Heather Williams, David Birtwell

Penn Medicine BioBank
University of Pennsylvania
Philadelphia, PA USA

hwilli@upenn.edu, birtwell@upenn.edu

Abstract— We present a use case of the **Ontology for Biomedical Investigations [1] (OBI)** informing the software design of a suite of biobanking applications. We describe how OBI has influenced the design of the Penn Medicine BioBank applications that support the collection, processing, and storage of biobank specimens and our work in creating a robust search system over data produced by BioBank applications and other sources. We show that applications that have been designed with the tenets of OBI in mind, particularly those of being reality based and modeling events as OBI style processes, have proven to effectively express richly interconnected data and be easily extendable.

Keywords— *BFO; Biobanking; OBI; Ontology; Process; Search; Software Design*

I. INTRODUCTION

Bio-specimens and the data gained by their analysis are valuable resources for bio-medical investigators. Biobanks, collections of bio-specimens (specimens) made available for research, are of extreme importance to investigators, because they can provide a large enough sample size to perform robust statistical analysis and can be used to find specimens with rare genotypes or phenotypes of interest. Information associated with specimens in biobanks and the subjects from whom the specimens were collected is frequently as important to research as the information gleaned from specimen analysis. Information technology such as databases and web application frameworks provide basic support for the storage and retrieval of biobank information. However, these technologies do not provide models for complex bio-medical data. Modeling such rich interconnected data remains a challenge for bio-medical investigators and informaticians, one that must be overcome for specimen based research to reach its full potential.

The Penn Medicine BioBank (PMBB) enables biomedical research by providing centralized access to a large number of annotated blood and tissue specimens. The Penn Medicine BioBank Informatics Team has been tasked with supporting this initiative by creating the informatics infrastructure to enable the collection, processing, and storage of specimens and associated subject data, and making the biomedical and demographic information associated with its subjects and specimens readily available to the research community. The

information must be easily accessible, discoverable, and query-able, and data provenance must be maintained.

Since 2013, the PMBB informatics infrastructure has been implemented by a suite of biobanking applications collectively called Squash that are founded on OBI concepts with the aim of presenting and interacting with biobank information in a semantically rich ontology adherent manner. We designed our data model to follow patterns and conventions established by OBI, its higher order ontology Basic Formal Ontology [8] (BFO), and the OBO Relation Ontology [9]. BFO is a theory of the basic structures of reality currently being developed at the Institute for Formal Ontology and Medical Information Science (IFOMIS) at the University of Leipzig [11]. The OBO Relation Ontology provides guidelines for creating ontologies with consistent relational assertions.

To date, we have implemented a web based specimen collection and processing application named Pumpkin that makes heavy use of the concept of a process [2] and have prototyped a query system that searches over OBI annotated data. We have found that modeling events such as pre-storage specimen processes like aliquoting, centrifugation, and freezing as processes with specific end-points, inputs, and outputs, has led to a powerful application with an expressive data model that reflects reality and is easily transformable to an ontology friendly format. We also found that keeping our data model reality based following the example set by OBI has resulted in a data model over which it is easy to reason and that facilitates organic extension.

II. METHODS

A. OBI Driven Software Design

Pumpkin was developed using the web application framework Grails [3] and is written in Groovy [4] and Java [5] using MySQL [6] as the relational database backend. Pumpkin supports the specimen collection process from the initial creation of specimen collection packets, through the processing and ultimate storage of specimens.

Grails incorporates an object-relational mapping (ORM) powered by Hibernate [7] that provides an abstraction layer over relational databases. Instead of creating tables with fields and foreign keys, one creates inter-related domain classes that specify the database schema. Data are written and read from the database at the Groovy object level rather than via SQL. Using the rich Grails ORM, we were able to model our persistent data in a manner very similar to the way classes are defined in an ontology like OBI -- reality based with class inheritance.

When designing our data model we considered the concepts represented in OBI and the relationships between them and theorized how new concepts would be represented as a guide to designing persistent domain objects. In this way, OBI informs both the software architecture and the structure of data that is created by the application. Some examples of OBI terms that were modeled as domain objects are the concepts of protocol, specimen collection, containers, and specimens.

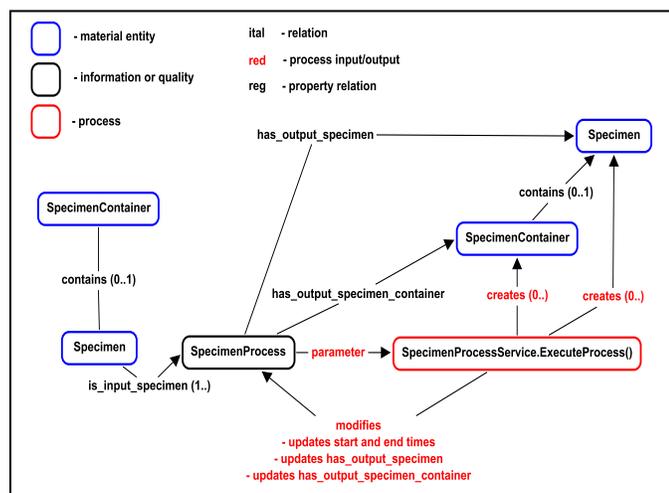


Fig. 1. The specimen process architecture expressed here in an informal graph exemplifies how OBI concepts influenced the design of Pumpkin. Specimens, SpecimenContainers, and SpecimenProcesses are all modeled as persistent domain objects.

The concept of a process heavily influenced our design. From the BFO concept of a process, we included both start and end times in our process domain classes. The OBI relationships `has_specified_input` and `has_specified_output` are implemented as well. For example, we modeled a domain super-class `SpecimenProcess` that includes input specimens, output specimens, start and end times, and a user (participant). Subclasses of `SpecimenProcess` include common specimen processes like aliquot, spin (centrifugation), dilute, and trash. Given specimen processes modeled in this way, each specimen is part of a directed specimen process graph that starts with a specimen extraction process and terminates with processes that have output specimens bound for storage.

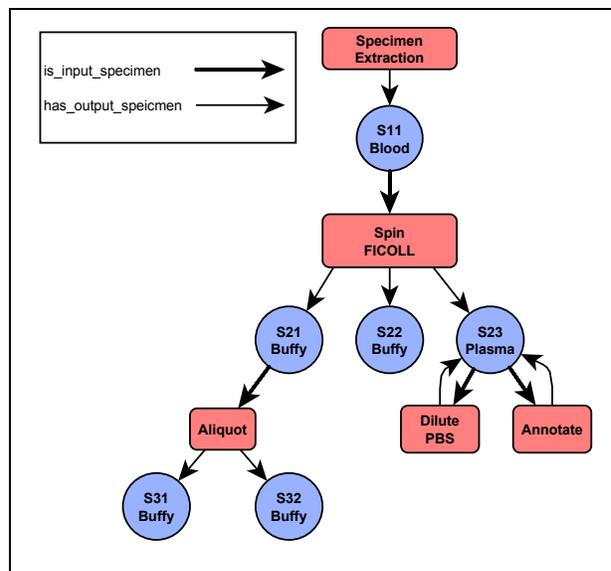


Fig. 2. An example of a typical specimen workflow showing specimen processes and their input and output specimens as modeled in Pumpkin following OBI guidelines. S11 is the primogenitor specimen. Because specimen processes are explicitly modeled, information can be directly associated with processes and specimens or inferred via the graph. For example, information pertaining to the specimen extraction of S11, like the study subject, is directly associated only with S11 and discoverable for derivative specimens by graph traversal.

B. OBI Annotated Data Search

We have developed a prototype search system that implements a natural language query interface over OBI annotated data. Ontology experts analyzed several small existing biomedical data sets and created a mapping between the data and concepts in the OBI ontology. D2RQ [10] was used to present these annotated data as a SPARQL endpoint.

To enable natural-language-like queries (NLQ), a pipeline following the standard programming language compilation process was created. An NLQ is first parsed as per a fully specified context free grammar. The resulting parse tree is fed to an interpreter that creates a logical query representation. A query generator takes as input this logical query representation and generates a SPARQL query that is run against the NLQ query endpoint.

III. RESULTS

Pumpkin has been in production since June 2013 and to date has stored over 90,000 specimens from over 8,000 collections. Its design has proven to be adequate to handle our initial collection specifications and be easily extendable to additional processes and concepts, such as new specimen and collection attributes. Since the data model is reality-based and expressive as it is in graph form, it provides a common representation for all biobank related data, independent of individual lab nomenclature and idiosyncrasies. Because the

data are stored in a harmonized data model, no transformation is required to query across these data.

IV. DISCUSSION

Early in the requirements gathering and design process, it became clear that one of the primary difficulties of biobanking informatics is the heterogeneity and interconnectedness of the information involved. Application developers are mostly unaccustomed to modeling entities and processes as diverse and complex as those found in biology and biobanking. While the volume of data is small in modern terms, the complexity and fragility is great. In order to remain useful, each bit of information concerning a biological process must be richly explained, which often means complex links to other bits of information and semantic definitions. Our development team, staffed with computer science and math majors, found itself ill equipped to meet the challenge of modeling the information of a robust biobanking informatics landscape. Traditional data modeling techniques as they apply to relational and document databases fall short. It was only after several months acquainting ourselves with OBI and ontology concepts in general that we were able to see a path to an informatics system that would provide data expressivity equal to the task. What ensued was the implementation of a web based biobanking application designed from the ground-up to be OBI compliant.

Two tenets of OBI stand out as particularly significant. The first is the dedication to remaining reality based. It is often more convenient to model data for a given requirement in a way that satisfies that requirement only, usually following the path of least resistance of the implementation technology, than it is to adhere to a reality based model. From the outset, we committed ourselves to a reality based data model following the example set by OBI. While this commitment did prove difficult and seemed unnecessarily so at times, inevitably it led to an understandable and often surprisingly easily extendable data model. The second is our choice to model events as BFO style processes, occurs with temporal boundaries, following the OBI convention of including process inputs and outputs. It was unclear at the outset that this approach would lead to an improved data model. We found however that much like our commitment to remaining reality based, modeling our processes in this way resulted in an understandable and easily extendable data model.

This approach has not been without challenges. One notable difficulty arose around efficient information retrieval from the database. To get the full data for a particular specimen, the specimen process graph must be generated, which in our initial implementation required recursive domain class traversals resulting in an explosion of computationally expensive database calls. We addressed this issue via shortcut pointers in the database. In most instances the data needed for a particular specimen are associated with either the specimen itself or its primogenitor specimen. To alleviate the

computational load of traversing the process graph for common tasks, each specimen was assigned a direct pointer to its primogenitor specimen allowing single database queries rather than recursive searches.

In the hopes of finding a more general solution to efficient data retrieval, we are experimenting with mirroring our data in a graph database. Graph databases are designed to store and operate efficiently over data in graph format and may provide a mechanism to perform efficient reads of our data.

In addition to specimen processes, we have loosely modeled the concept of a 'task' to follow the OBI methodology as a time-based process with inputs and outputs. Currently, we model specimen intake as a task. In the future, other tasks will be included. As with specimen processes, task data will be expressed in graph form and so the same efficiency considerations will exist and methods used.

Still to be developed is Carnival, the system that will tie together the subject and specimen data generated by Squash applications with data from other sources and present them in a discoverable and query-able format. This will be an expansion of the prototype natural language query tool. The data generated by and stored in Squash applications exist at rest in a form that is compatible with OBI. We plan to annotate any additional data from sources outside Squash with OBI terms and provenance information in order to create a unified search endpoint.

Through our experiences attempting to create ontology adherent database applications, we have gained an appreciation for the valuable work that has been and continues to be done in ontology development. We suspect that the perceived value of ontologies within the biomedical research community will increase over time as those outside the immediate ontology community learn the contributions that ontologies like OBI and BFO can make towards their efforts. It remains to be seen whether ontology influenced software design will be adopted by the broader software development community, but if there is continued success of the Penn Medicine Biobank, it will be due in large part to the influence ontologies have had on our software development team.

REFERENCES

- [1] Brinkman RR, Courtot M, Derom D, Fostel JM, He Y, Lord P, Malone J, Parkinson H, Peters B, Rocca-Serra P, Ruttenberg A, Sansone SA, Soldatova LN, Stoeckert CJ Jr, Turner JA, Zheng J; OBI consortium. (2010) Modeling biomedical experimental processes with OBI. *J Biomed Semantics*. 2010 Jun 22;1 Suppl 1:S7.PMID: 20626927
- [2] Whetzel PL, Noy NF, Shah NH, Alexander PR, Nyulas C, Tudorache T, Musen MA. BioPortal: enhanced functionality via new Web services from the National Center for Biomedical Ontology to access and use ontologies in software applications. *Nucleic Acids Res*. 2011 Jul;39(Web Server issue):W541-5. Epub 2011 Jun 14.
- [3] Grails. [<https://grails.org>].
- [4] Groovy. [<http://groovy.codehaus.org>].
- [5] Java. [<http://www.java.com>].
- [6] MySQL. [<http://www.mysql.com>].
- [7] Hibernate. [<http://hibernate.org>].
- [8] Grenon P, Smith B, Goldberg L. *Ontologies in Medicine*. IOS Press; 2004. *Biodynamic Ontology: Applying BFO in the Biomedical Domain*. pp.20–32.
- [9] Smith B, Ceusters W, Klagges B. Relations in Biomedical Ontologies. *Genome Biology*. 2005;6:R46. doi: 10.1186/gb-2005-6-5-r46.
- [10] D2RQ [<http://d2rq.org/>].
- [11] Grenon, P. and Smith, B. (2004) “SNAP and SPAN: Towards Dynamic Spatial Ontology”, *Spatial Cognition and Computation*, 4:1, 69-103.

Towards a Common Semantic Representation of Informed Consent for Biobank Specimens

Frank J. Manion^{1*}, Yongqun He², Elizabeth Eisenhauer³, Yu Lin², Alla Karnovsky⁴,
Marcelline R. Harris³

¹ Comprehensive Cancer Center, University of Michigan Medical School, Ann Arbor, MI 48109

² Unit for Laboratory Animal Medicine, Department of Microbiology and Immunology, and Center for Computational Medicine and Bioinformatics, University of Michigan Medical School, Ann Arbor, MI 48109, USA

³ Division of Systems Leadership and Effectiveness Science, University of Michigan School of Nursing, Ann Arbor, MI 48109

⁴ Department of Computational Medicine and Bioinformatics, University of Michigan Medical School, Ann Arbor, MI 48109

Abstract — Biospecimen-based research is rapidly growing in the post genomic era, and includes the need to retrieve specimens from distributed biobanks of various size and complexity in a fashion that ethically preserves the expressed wishes of specimen donors as represented by the informed consent process and its artifacts. This paper briefly describes existing work along these lines, presents some challenges unique to biobanks, and presents our own work on an ontology of informed consent.

Keywords— *BioBank, Informed consent; ontology; ICO; OBO Foundry, Basic Formal Ontology (BFO), OBI ontology*

I. INTRODUCTION

Research in the post-genomic era requires access to high quality biospecimens, often annotated with or linked to clinical data. Many groups at varying levels of institutional complexity, ranging from small scale individual laboratories to distributed international collaboratives, have established and operate biorepositories (also referred to by various names such as *biobanks*, *biolibraries*, and even *collections*). Often, there are needs to share data and specimens among multiple biobanks [1-3]. The act of requesting specimens from a biorepository may demand a complex series of transactions, each of which in turn may convey a series of rights, obligations, and permissions for access to specimens and data. Despite over a decade of experience incorporating biospecimens in the research process, formal models that describe the use of biorepositories in human research are a relatively recent development. Without a common formal model of consent and the associated permissions on collection and distribution of specimens and data, integration of data across the translational spectrum, or from multiple banks and institutions will remain a difficult, manually intensive problem.

In this paper we briefly review current efforts toward such models, describe our own work toward a formal model for informed consent, and describe what we consider challenges and opportunities for supporting biorepository-based research with ontologies. A simple example that provides a motivation for this effort follows.

II. EXAMPLE OF THE CHALLENGE

Clinical or translational research often involves the extraction and usage of biospecimen from humans. Different biospecimens may be stored and processed differently, and

may be collected under different models of consent. A typical scenario might read something like this:

“For my study, I want to use samples from my organizations’s biobank, collected under a blanket biobank informed consent form. I discover that I will need more samples, so I contact another organization’s biobank to determine if they hold relevant and available specimens. That organization’s samples were collected under a tiered biobank informed consent form. While some samples are shared with me, I still need more samples to address the requirements of my study. I then collect additional samples using a consent form specific to my study.”

In this example there are three informed consents forms to account for – a blanket consent, a tiered consent, and the investigator’s single study consent. In an effort to support the expressed wishes of the donors, informed consent documents impose a series of legal and ethical restrictions, obligations, and permissions to biobank operators and research teams using the specimens and data collected in these banks. Often these rights, obligations, and permissions accrue from multiple sources of authority and are represented in multiple legal documents. Consequently, the biobanking domain presents a series of modeling challenges, including:

The operational model of the biobank. A biobank can be a single, dedicated resource that provides samples to single or closely allied groups of studies using a common consent model. It might be a virtual or distributed biorepository using precoordinated consent models. Another organization structure might be that of a shared biobank facility containing multiple sets of tissues from multiple projects and attempting to maximize use of these tissue resources by making them available to requestors.

The consent model used for the biobank. This can be opt-in or opt-out. In the case of an opt-in consent model, a tiered consent may be used to present the participant or volunteer with choices of the type of data the participant may want shared, and for what types of research or other constraints.

The protocol model the bank operates under. Typically a biobank serving more than one project would operate under one or more Institutional Review Board (IRB)-approved collection protocols and Health Insurance Portability and Accountability Act (HIPAA) authorizations. Researchers

* Corresponding author.

subsequently requesting specimens and data would operate under separate IRB-approved protocols, and depending on this protocol, separate consent and HIPAA authorization may be required for use of a previously banked specimen. Such a model is sometimes called a *two-protocol* model [4].

Rights, obligations, and permissions accrue from multiple sources and must be consistent across time. Properly modeling the decisions typically made by human review boards and regulatory personnel considering sample and data distribution for research requires modeling not just the consent documents, but the protocols, data use agreements, and possibly other information artifacts used in both depositing samples into a biobank, and withdrawing them for subsequent research.

In a research oriented university such as the University of Michigan, thousands of informed consent forms have been generated, and there are over 100 biobanks in the Medical School alone. Queries supporting appropriate use of banked biospecimens and data must be linked to the signed informed consent agreements with the biospecimen donor.

III. EXISTING EFFORTS

Several current efforts are evident, focused on modeling aspects of the biobanking domain. At least two BFO-aligned ontologies relate to biobanking. The Ontologized Minimum Information About Biobank data Sharing (OMIABIS) expresses data concepts in an ontology of biobank administration [5]. OMIABIS is based on work by Norlin and colleagues [6] to develop a minimum data set for eight countries participating in the EU Biobanking and Biomolecular Resources Research Infrastructure project. Limitations of this effort are that it is intended to serve only as a description of a biobank contents, and does not describe collection criteria, consenting, and protocol provenance of individual specimens. A group at the University of Pennsylvania is developing an ontology for the representation of biobanks, although the work is in early stages [7]. Similarly, we are aware that a group at Duke University is working on a collaborative effort to develop a normative set of data elements and terms to recommend as best practice to the International Society for Biological and Environmental Repositories (ISBER), although this work is not yet published [8, 9].

There are also non-BFO aligned ontologies in related areas, including a Permission Ontology used for development and evaluation of software tools for reasoning about consent permission, published by a group at the University of California San Diego (UCSD) [10]. Related work to build a Research Permission Management System was done at the Medical University of South Carolina (MUSC) to support a statewide research network [11]. A search of the term “consent” in the NCBO biportal identified the notion of informed consent at the class level in 19 different systems (<http://biportal.bioontology.org/search>).

Our efforts to develop a BFO-aligned informed consent ontology (ICO) emphasizes the broad domain of informed consent. Although motivated by a biobanking use case, initial development reported here is not restricted to that domain.

IV. THE INFORMED CONSENT ONTOLOGY (ICO)

Development of ICO, a BFO-based ontology represented in the Web Ontology Language (OWL2) [12], follows OBO Foundry principles of openness and collaboration. ICO is aligned with the BFO [13], making it possible to align and integrate with other BFO-based ontologies. The initial release of the ontology focuses on modeling informed consent documents. As for Aug. 14, 2014, ICO contains 471 terms including 137 ICO-specific terms and other terms imported from other BFO-aligned ontologies. Detailed ICO statistics can be found on the Ontobee ICO web page: <http://www.ontobee.org/ontostat.php?ontology=ICO>. ICO is released under an open Creative Commons 3.0 License.

The ontology was developed using a combination of top-down and bottom-up approaches. Protégé-OWL 4.2 was used for the ontology authoring and editing. To build the OBI-based framework of ICO we manually identified informed consent concepts from existing OBO Foundry library ontologies. These were imported to ICO using Ontodog [14] and OntoFox [15] which allowed for recursive inclusion of all defined axioms and related terms. The results were then manually reviewed for final approval before inclusion in the ICO framework.

Bottom up construction proceeded by manually identifying and extracting a list of candidate terms from two informed consent templates used at the University of Michigan (one from the Medical School Institutional Review Board, another from the Health Sciences and Behavioral Sciences Institutional Review Board). We also identified terms from a consent form used for the University of Michigan Medical School biorepository, and from World Health Organization (WHO) informed consent templates. The candidate terms identified from these templates were then enriched with metadata including definitions, concept identifiers, preferred terms, synonyms, and URIs extracted from three ontology repositories: the National Library of Medicine’s Unified Medical Language System (UMLS®) Metathesaurus [16]; the National Center for Biomedical Ontology (NCBO) BioPortal [17]; and Ontobee [18]. When textual definitions were not provided, other sources such as clinical research glossaries or the current literature were used. These enriched candidate terms were manually mapped to several pre-identified resources containing terms and definitions developed and vetted by the United States regulatory community. This process yielded candidate preferred terms containing definitions accepted as robust and well defined by that community. Resources used in this step included the National Cancer Institute Thesaurus (NCIt), the Biomedical Research Integrated Domain Group (BRIDG) [19], the Ontology of Clinical Research (OCRe) [20], the Consumer Health Vocabulary (CHV) and the University of California San Diego permission ontology [10].

The pool of enriched candidate terms was organized into categories of like terms according to their definitions. For example, the category ‘authorization’ included the terms ‘*authorization for medical records release*’, ‘*authorization documentation*’ or ‘*authorization*’. Enriched candidate terms grouped by categories formed to-be-included terms in ICO. The final set of categories (or modeling units) was then

mapped to branches of BFO. For example, terms categorized under ‘authorization’ were considered to be subclasses of BFO:process. Informed consent workflows in a typical clinical research study were modeled as three processes: (i) pre-informed consent processes, (ii) obtaining informed consent processes, and (iii) processes after signing informed consent documents. Relations between entities involved in the above processes were defined. Finally, all terms and relations were aligned with BFO.

V. DISCUSSION AND CONCLUSIONS

Modeling informed consent is a necessary but not sufficient part of the modeling needed to support responsible use of biospecimens and data in research. Biospecimen and data release is complex, and informed consent plays a major role in the regulatory and scientific governance used by biorepositories to release specimens and data. In follow on work we plan to examine the specific area of specimen and data release involving the longitudinal agreements of rights, permissions, and obligations. Other work is needed in the complex areas of protocol representation, data use agreements and material transfer agreements.

Limitations of our preliminary work will inform further development efforts toward a robust Informed Consent Ontology. First, the ICO is admittedly preliminary work and is currently focused on informed consent documents and processes. More work is needed to validate the coverage and completeness in the domain. Concepts from the US Common Rule and the EU Prior Informed Consent legislation need to be included. Our current models of informed consent processes likely lack the richness and complexity of real-life informed consent processes, and they need validation with research study teams from a variety of domain areas. Aspects of rights, obligations, permissions, and ethics must be modeled and used to extend the ontology. Finally, axioms must be developed and competency validation of the ICO must be conducted using a series of still to be defined use case derived competency questions.

We have described our work on ICO, a preliminary ontology of informed consent that provides general classification of content contained in general informed consent documents. It requires expansion, revisions and collaboration to build a robust model, and to move toward a representation of the complex area of biobank data sharing and specimen release. We hope to collaborate with the broader community in this effort.

ACKNOWLEDGMENT

This research was supported by a University of Michigan interdisciplinary research award (MCubed) and by the National Center for Advancing Translational Sciences of the National Institutes of Health under Award Number 2UL1TR000433-06. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

REFERENCES

- [1] A. Cambon-Thomsen, E. Rial-Sebbag, and B. M. Knoppers, "Trends in ethical and legal frameworks for the use of human biobanks," *Eur Respir J*, vol. 30, pp. 373-82, Aug 2007.
- [2] R. E. Hewitt, "Biobanking: the foundation of personalized medicine," *Curr Opin Oncol*, vol. 23, pp. 112-9, Jan 2011.
- [3] G. E. Henderson, R. J. Cadigan, T. P. Edwards, I. Conlon, A. G. Nelson, J. P. Evans, A. M. Davis, C. Zimmer, and B. J. Weiner, "Characterizing biobank organizations in the U.S.: results from a national survey," *Genome Medicine*, vol. 5, no. 1, p. 3, Jan. 2013.
- [4] F. J. Manion, R. J. Robbins, W. A. Weems, and R. S. Crowley, "Security and privacy requirements for a multi-institutional cancer research data grid: an interview-based study," *BMC Med Inform Decis Mak*, vol. 9, p. 31, 2009.
- [5] M. Brochhausen, M. N. Fransson, N. V. Kanaskar, M. Eriksson, R. Merino-Martinez, R. A. Hall, L. Norlin, S. Kjellqvist, M. Hortlund, U. Topaloglu, W. R. Hogan, and J.-E. Litton, "Developing a semantically rich ontology for the biobank-administration domain," *J Biomed Semantics*, vol. 4, p. 23, 2013.
- [6] Norlin, M. N. Fransson, M. Eriksson, R. Merino-Martinez, M. Anderberg, S. Kurtovic, and J.-E. Litton, "A Minimum Data Set for Sharing Biobank Samples, Information, and Data: MIABIS," *Biopreservation and Biobanking*, vol. 10, no. 4, pp. 343-348, Aug. 2012.
- [7] J. Zheng, Personal Communication to Yongqun He, 2014.
- [8] H. Ellis, Personal Communication to Frank J. Manion, 2014.
- [9] ISBER. (2014). ISBER - International Society for Biological and Environmental Repositories. Available: <http://www.isber.org/>
- [10] A. Grando and R. Schwab, "Building and evaluating an ontology-based tool for reasoning about consent permission," *AMIA Annu Symp Proc*, vol. 2013, pp. 514-23, 2013.
- [11] J. S. Obeid, K. Gerken, K. C. Madathil, D. Rugg, C. E. Alstad, K. Fryar, R. Alexander, A. K. Gramopadhye, J. Moskowitz, and I. C. Sanderson, "Development of an Electronic Research Permissions Management System to Enhance Informed Consents and Capture Research Authorizations Data," *AMIA Summits Transl Sci Proc*, vol. 2013, pp. 189-193, 18 2013.
- [12] W3C, "OWL 2 Web Ontology Language document overview," pp. <http://www.w3.org/TR/2009/REC-owl2-overview-20091027/>. Accessed on March 1, 2014, 2009.
- [13] P. Grenon and B. Smith, "SNAP and SPAN: Towards dynamic spatial ontology," *Spatial Cognition and Computation*, vol. 4, pp. 69-103, 2004.
- [14] J. Xiang, Z. Xiang, C. J. Stoeckert, and Y. He, "Ontodog: a web-based ontology community view generation tool," *Bioinformatics*, vol. 30, no. 9, pp. 1340-1342, May 2014.
- [15] Z. Xiang, M. Courtot, R. R. Brinkman, A. Rutenberg, and Y. He, "OntoFox: web-based support for ontology reuse," *BMC Res Notes*, vol. 3, p. 175, Jun. 2010.
- [16] NLM, "U.S. National Library of Medicine (NLM). Unified Medical Language System (UMLS®) Metathesaurus [version 2013AA][Internet]. Bethesda (MD): National Library of Medicine," 2013.
- [17] P. L. Whetzel, N. F. Noy, N. H. Shah, P. R. Alexander, C. Nyulas, T. Tudorache, and M. A. Musen, "BioPortal: enhanced functionality via new Web services from the National Center for Biomedical Ontology to access and use ontologies in software applications," *Nucleic Acids Res*, vol. 39, no. Web Server issue, pp. W541-W545, Jul. 2011.
- [18] Z. Xiang, C. Mungall, A. Rutenberg, and Y. He, "Ontobee: A Linked Data Server and Browser for Ontology Terms.," in *Proceedings of the 2nd International Conference on Biomedical Ontologies (ICBO)*, Buffalo, NY, USA, 2011, pp. 279-281.
- [19] D. B. Fridsma, J. Evans, and C. N. Mead, "The BRIDG Project: A Technical Report," *Journal of the American Medical Informatics Association*, vol. 15, no. 2, p. 130, 2008.
- [20] I. Sim, S. W. Tu, S. Carini, H. P. Lehmann, B. H. Pollock, M. Peleg, and K. M. Wittkowski, "The Ontology of Clinical Research (OCRe): An informatics foundation for the science of clinical research," *Journal of Biomedical Informatics*, Nov 13 2013.

A Specimen-based View of the World

Using the Biological Collections Ontology to Model Biodiversity Collections

Ramona Walls

iPlant Collaborative
University of Arizona
Tucson, AZ, USA
rwalls@iplantcollaborative.org

Robert Guralnick

Ecology and Evolutionary Biology
University of Colorado at Boulder
Boulder, Colorado, USA
robert.guralnick@colorado.edu

John Deck

Berkeley Natural History Museums
University of California
Berkeley, California, USA
deck@berkeley.edline

Andréa Matsunaga

Dept. of Electrical and Computer Engineering
University of Florida
Gainesville, Florida, USA
ammatsun@ufl.edu

Abstract—Application ontologies for biodiversity and biomedical specimen data are being developed within the OBO Foundry framework. Both the Biological Collections Ontology (BCO) and the Ontologized Minimum Information About Blobank data Sharing (OMIABIS) are ontologies rooted in specimens and the need to track, share and query data about specimens. In this paper, we briefly describe the structure of the BCO and the way that it is used to annotate and reason over biodiversity data. We conclude with a discussion of the relationship of the BCO to bio-bank ontologies and areas of potential collaboration through the Ontology for Biomedical Investigations (OBI).

Keywords—ontology; biodiversity; specimen; material sample; Darwin Core; MlXS

I. INTRODUCTION

Museum specimens and the data associated with them are a critical foundation of biodiversity knowledge. They provide evidence of an organism's occurrence at a particular place and time and are source material for genetic, genomic, and metagenomic sequence data as well as for morphological, physiological, and biochemical trait measurements. Environmental data and field notes taken at the time of specimen collection or observation provide needed context that can form the basis of ecological studies into species' distributions and interactions (e.g. [1–3]). As our ability to measure and record scientific data has grown through technologies such as sequencing and digital data capture, so too has the need to store, track, access, and understand new types of specimens and their associated data.

The Biological Collections Ontology (BCO) [4] is a semantic model that describes and links both traditional and novel types of biodiversity data. While observations play a key

role in biodiversity research and the BCO, the need to track and describe relationships between specimens, their origins, and their derivatives continues to be the BCO's primary driving use case. Although the BCO models basic biodiversity domain knowledge, it is primarily an application ontology that relies on imports from and coordination with other ontologies such as the Ontology for Biomedical Investigations (OBI) [5], the Environment Ontology (ENVO) [6], and the Population and Community Ontology (PCO) [4]. Coordination with ENVO is crucial for describing the environments in which specimens are collected, and coordination with PCO allows the BCO to describe multi-organism specimens, such as metagenomic samples.

The BCO grew out of a series of workshops [7,8] aimed at harmonizing traditional museum collection data, typically described using Darwin Core (DwC) [9], and genomic-based biodiversity data, typically described using Minimum Information for any (x) Sequence (MlXS) [10,11]. Although MlXS is a standard for sequence data, there is overlap with specimen-based standards given that many of the MlXS terms describe the specimen that was sequenced or the conditions under which it was collected. The first term developed for the BCO was *material sample*¹, which was defined as a *material entity* (from the Basic Formal Ontology or BFO) [12,13] that **realizes a material sample role** by being the output of some *material sampling process*. As the BCO matured, it became apparent that BCO classes for *material sample* and *material sampling process* were very similar to *specimen* and *specimen collection* in OBI. Because these and many of the other concepts needed to describe biodiversity data were already present in OBI, a decision was made to coordinate BCO development with OBI. Classes that can be applied to biological investigations generally, such as *specimen*, should be housed in OBI, while only those specific to biodiversity studies, such as *museum specimen*, should be maintained in the BCO.

Funding was provided by EAGER: An Interoperable Information Infrastructure for Biodiversity Research (NSF-IIS-1255035), by RCN4GSC: A Research Coordination Network for the Genomic Standards Consortium (NSF-DBI-0840989), and by Collaborative Research: BiSciCol Tracker: Towards a tagging and tracking infrastructure for biodiversity science collections (NSF-DBI: 0956371, 0956350, 0956426).

1. Ontology class names are shown in italic and relations in bold.

II. ONTOLOGY DESIGN

The BCO is being developed according to OBO Foundry principles [14]. It is organized around two of the key processes that generate biodiversity data: *specimen collection* and *observing process* (Fig. 1). Both have a *material entity* as an input, but the key difference is that *specimen collection* generates a *material entity* (a *specimen*) as output while *observing process* generates an *information content entity* (from the Information Artifact Ontology or IAO) [15]. The BCO interprets *specimen collection* in a broad sense to include collection of museum or herbarium specimens, subsampling processes such as tissue sampling or DNA extraction, and collection of environmental (e.g., metagenomic) samples. Sequence generation and its output data are also crucial to biodiversity studies, but they have been modeled in OBI and the Sequence Ontology [16], so we import classes as needed for those concepts.

An essential functionality of the BCO is the ability to trace data through a series of processes. For example, one may have organismal sequence data stored in GenBank [17] and metagenomic data stored in another database and want to determine if those sequences came from the same museum sample (Fig. 2). To make queries like this, we needed a transitive property chain that links inputs and outputs of planned processes, which was not available in the Relation Ontology (RO) [18] or BFO. We created two relations using property chains, defined as follows:

is_specified_input_of o 'has output' subPropertyOf 'derives from by planned process'
(http://purl.obolibrary.org/obo/BCO_0000067)

is_specified_output_of o 'has input' subPropertyOf 'is derived into by planned process'
(http://purl.obolibrary.org/obo/BCO_0000068)

The **is_specified_input_of** and **is_specified_output_of** relations are from OBI and **has output** and **has input** are from the RO. We are working with curators of RO to develop more broadly applicable definitions and names for these relations.

III. USING THE BCO

One of the main uses of the BCO is to query over data sets that have metadata associated with both *specimens* and *specimen collection*. We have held several workshops in which we mapped column headings to ontology terms and specified relations among columns in order to convert datasets from the typical spread-sheet format to RDF [19]. Work is ongoing to develop tools that can automate mapping of data in common formats (such as Darwin Core archives or MIXS spread sheets) to RDF using BCO and other ontologies (see the BiSciCol Triplifier [20] and Biocode FIMS tools [21]). Three major challenges in this endeavor are that few researchers distinguish between specimens and specimen collection processes when they are recording data, the information content of many spreadsheets is ambiguous, and the lack of a clear standard for instance identifier assignment for biodiversity data.

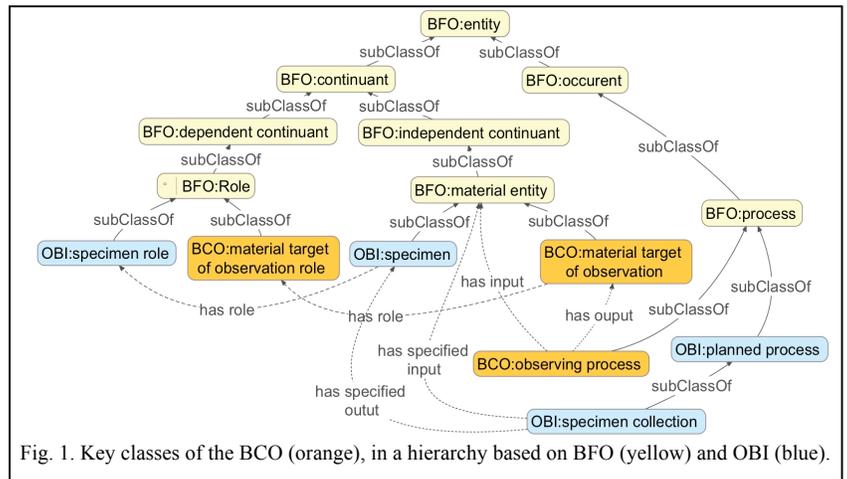
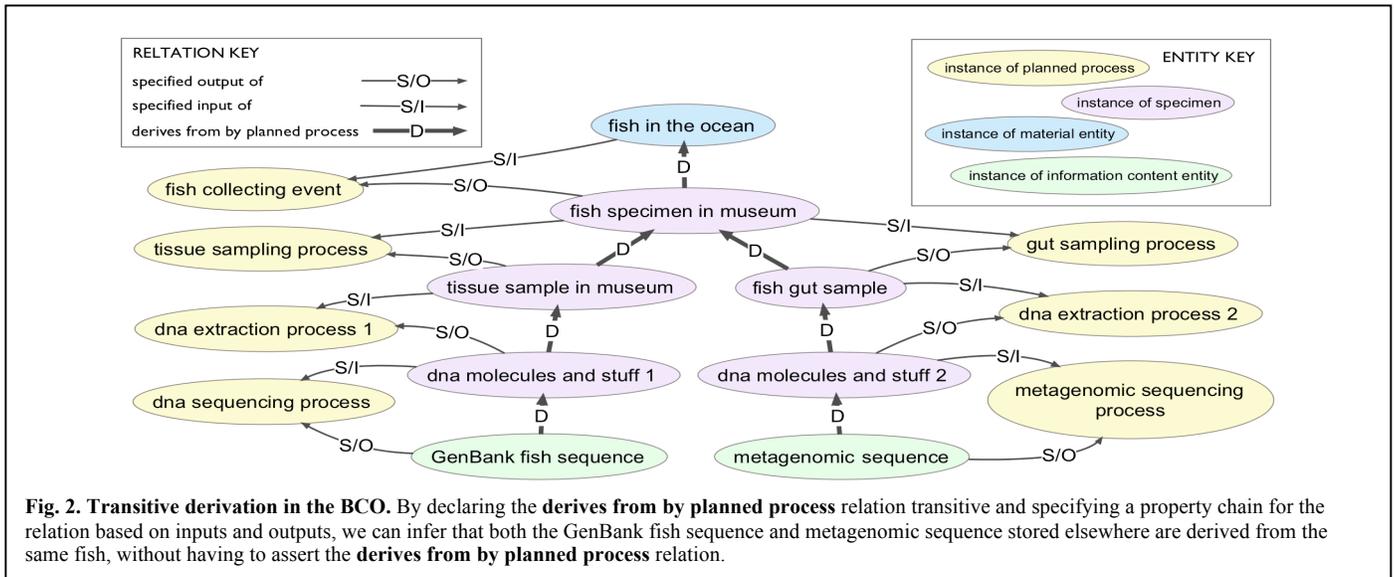


Fig. 1. Key classes of the BCO (orange), in a hierarchy based on BFO (yellow) and OBI (blue).

IV. THE BCO AND BIOBANK ONTOLOGIES

Similar to biodiversity specimen repositories, biomedical specimen repositories (biobanks) need to track and share data on specimens, their sources and their derived products or data. Here we compare the BCO to the Ontologized MIABIS (OMIABIS), named after the Minimum Information About Biobank data Sharing (MIABIS) [22]. OMIABIS is not the only existing biobank ontology, but is the only published, freely available one of which we are aware. We also considered the ontology described by [23] but do not have access to a current version for direct comparison. BCO and OMIABIS not only share a similar focus on specimens but also reuse many of the same terms from OBI and IAO. Some aspects that appear to differ between the two ontologies are in fact simply domain-specific differences in terminology. For example, descriptions of the environment in which a sample was collected in the BCO are, from a knowledge modeling perspective, very similar to a description of patient disease status and history in biobank ontologies, as both describe the conditions under which a specimen was collected. We recommend that developers of both ontologies work to develop a shared set of design patterns that can be used to model the environmental context of specimens, other aspects of specimen collection processes, and relations such as **derives from by planned process**, described earlier.

Both the biodiversity and biobank domains have standards for describing data (MIABIS for biobanks and DwC and MIXS for biodiversity specimens) and infrastructure for aggregating relevant data – the European Biobanking and Biomolecular Resources Research Infrastructure (BBMRI) [24] for biobanks and the Global Biodiversity Information Facility (GBIF) [25] for biodiversity data. Nonetheless, the diversity of specimen and data types, the distributed nature of collections, and the novelty of informatic approaches to many researchers in both fields lead to uneven application of standards and additional challenges for semantic reasoning, particularly across legacy data sets. These challenges call for tools that make the ontologies easier to work with, on top of ontology development. We see an opportunity for BCO and OMIABIS developers to collaborate on tool development in areas such as universally unique specimen identifiers, data integration across legacy data sources, and reasoning over large data sets.



It is clear that there are many areas of overlap and potential collaboration in modeling biodiversity specimen collections and biobanks. We are interested in discussing re-use of OBI terms without importing the entire OBI logic chain. Much of OBI's logic is not necessary for BCO's use cases and is likely to put off potential users, and we would like to learn if OMIABIS faces a similar situation. The only major conflict we found during this comparison was a difference in the version of BFO used by the two ontologies, and this is a conflict we think can be easily resolved. Concepts from OMIABIS, such as the **owns** and **administers** relations are highly useful and important to the BCO to capture the administrative data and relationships among biodiversity collections, may be better housed in the more general OBI. We recommend that curators from both domains work with the OBI to develop common terminology wherever possible.

REFERENCES

- [1] C.H. Graham, S. Ferrier, F. Huettman, C. Moritz, and A.T. Peterson, "New developments in museum-based informatics and applications in biodiversity analysis," *Trends Ecol Evol*, vol. 19, pp. 497-503, 2004.
- [2] C. Moritz, J.L. Patton, C.J. Conroy, J.L. Parra, G.C. White, and S.R. Beissinger, "Impact of a century of climate change on small-mammal communities in Yosemite National Park, USA," *Science* vol. 322, pp. 261-264, 2008.
- [3] G.H. Pyke, P.R. Ehrlich, "Biological collections and ecological/environmental research: a review, some observations and a look to the future," *Biol Rev Camb Philos Soc*, vol. 85, pp. 247-266, 2010.
- [4] R.L. Walls, J. Deck, R. Guralnick, S. Baskauf, R. Beaman, S. Blum, P.L. Buttigieg, et al., "Semantics in support of biodiversity knowledge discovery: an introduction to the Biological Collections Ontology and related ontologies," *PLoS ONE*, vol. 9, p. e89606, 2014.
- [5] R. Brinkman, M. Courtot, D. Derom, J. Fostel, Y. He, P. Lord, et al., "Modeling biomedical experimental processes with OBI," *J Biomed Semant*, vol. 1(Suppl 1), p. S7, 2010.
- [6] P.L. Buttigieg, N. Morrison, B. Smith, C.J. Mungall, S.E. Lewis, "The environment ontology: contextualising biological and biomedical entities," *J Biomed Semant*, vol. 4, p. 43, 2013.
- [7] EÓ. Tuama, J. Deck, G. Dröge, M. Döring, D. Field, R. Kottmann, et al., "Meeting Report: Hackathon-Workshop on Darwin Core and MixS Standards Alignment (February 2012)," *Stand Genomic Sci*, vol. 7, pp.166-170, 2012.
- [8] J. Deck, K. Barker, R. Beaman, P.L. Buttigieg, G. Dröge, R. Guralnick, et al., "Clarifying concepts and terms in biodiversity informatics," *Stand Genomic Sci*, vol. 8, pp. 352-359, 2013.
- [9] J. Wieczorek, D. Bloom, R. Guralnick, S. Blum, M. Döring, R. Giovanni, et al., "Darwin Core: an evolving community-developed biodiversity data standard," *PLoS One*, vol. 7, p. e29715, 2012.
- [10] P. Yilmaz, J.A. Gilbert, R. Knight, L. Amaral-Zettler, I. Karsch-Mizrachi, G. Cochrane, et al., "The Genomic Standards Consortium: bringing standards to life for microbial ecology," *ISME J*, vol. 5, pp. 1565-1567, 2011.
- [11] P. Yilmaz, R. Kottmann, D. Field, J.R. Cole, L. Amaral-Zettler, J.A. Gilbert, et al., "Minimum information about a marker gene sequence (MIMARKS) and minimum information about any (x) sequence (MIXS) specifications," *Nat Biotechnol*, vol. 29, pp. 415-420, 2011.
- [12] P. Grenon, B. Smith, "SNAP and SPAN: towards dynamic spatial ontology," *Spat Cogn Comput*, vol. 4, pp. 69-103, 2004.
- [13] B. Smith, "On classifying material entities in Basic Formal Ontology," in: *Interdisciplinary Ontology. Proceedings of the Third Interdisciplinary Ontology Meeting*. Keio University Press, Tokyo, 2012, pp. 1-13.
- [14] <http://www.obofoundry.org/wiki/index.php/Category:Accepted>
- [15] <https://code.google.com/p/information-artifact-ontology/>
- [16] M. Bada, K. Eilbeck, "Efforts toward a more consistent and interoperable Sequence Ontology," in: *Proceedings of the 3rd International Conference on Biomedical Ontology (ICBO 2012)*. KR-MED Series. Graz, Austria; 2012.
- [17] <http://www.ncbi.nlm.nih.gov/genbank>
- [18] <https://code.google.com/p/obo-relations/>
- [19] R.L. Walls, R. Guralnick, J. Deck, A. Buntzman, P.L. Buttigieg, N. Davies, et al., "Meeting report: Advancing practical applications of biodiversity ontologies," *Stand Genomic Sci*, unpublished.
- [20] <https://code.google.com/p/triplifier/>
- [21] <https://code.google.com/p/biocode-fims/>
- [22] M. Brochhausen, M.N. Fransson, N.V. Kanaskar, M. Eriksson, R. Merino-Martinez, R.A. Hall, et al., "Developing a semantically rich ontology for the biobank-administration domain," *J Biomed Semant*, vol. 4, p. 23, 2013.
- [23] A. Andrade, M. Kreutzhaler, J. Hastings, M. Krestyaninova, S. Schulz, "Requirements for semantic biobanks," *Stud Health Technol*, vol. 80, pp. 569-573, 2012.
- [24] <http://bbmri.eu/>
- [25] <http://www.gbif.org>