

Developing a tool for Crowd-sourced Verification of a Radiation Oncology Ontology: a Summer Project

Project Description

Eleanor Pratt ^{1*}, Vishal Pandian ^{1†}, Evan Morrison ^{1‡}, Alexis Andrew Miller ^{1,2§}

1 Centre for Oncology Informatics, Faculty of Engineering & Information Science, University of Wollongong, Wollongong NSW 2515 Australia

2 Illawarra Cancer Care Centre, Wollongong Hospital, Wollongong NSW 2500, Australia

Abstract: We have been unable to find a verified, published Radiation Oncology Ontology. We undertook the process of verifying a Radiation Oncology Ontology with a mixture of crowd-sourcing and expert-based approaches to verify relationships in the ontology. We used a natural language based approach to portray concepts and relationships, surveying users to assess the relationships between concepts in the Radiation Oncology ontology. The work used a description of a patient's history expressed in XML. The natural language statements relating concepts are available on a website for verification, and readers are invited to complete the survey at <http://coi-hs-survey.appspot.com/> to contribute.

J Radiat Oncol Inform 2014;6:1:42-52

doi: 10.5166/jroi-6-1-23

The authors have no conflicts to disclose.

Keywords: Radiation Oncology • Ontology • crowd-sourcing • Natural language • Mechanical Turk

This article is licensed under a **Creative Commons Attribution 4.0 International License**.

1. Introduction

The creation of an ontology to describe an expert domain is a problematic and iterative process. The difficulties arise because the creation process covers two distinctly different domains rarely bridged by a single individual [1] - ontologists (the Informatics experts) and oncologists (the domain experts) share little. Unfortunately domain experts are rarely informaticians, and vice versa [2]. Another problem that adds complexity is the amount of time available for domain experts to guide the informatician, and since typically, domain experts use a knowledge structure developed by decades of study, identifying and transferring these structures in a time frame measured

* *E-mail: jp10010101010000@gmail.com*

† *E-mail: minimagic447@gmail.com*

‡ *E-mail: edm92@uowmail.edu.au*

§ *E-mail: alexisandrew@gmail.com*

in hours is unlikely [3].

The process is iterative in that the domain experts must describe knowledge structure to the informatician, and the informatician must then verify their developed knowledge structures with the domain experts. However once an ontology is produced, there are significant problems verifying its correctness [4]. Most ontologies will be defined in a format such as OWL and within software such as Protégé [5], neither of these are familiar to many domain experts. Until verification has been achieved, any domain-specific ontology must be regarded with suspicion and not used for practical purposes in the expert domain.

Radiation Oncology is a disease-centred (cancer) and modality-centred (radiotherapy) tertiary care medical specialty that utilises ionising radiation in the management of cancer [6], which only occupies a small position in a patient's lifelong care. Decision-making in Radiation Oncology follows a deductive path where verifiable and well documented information is applied to the precise problem of improved outcomes, while always reflecting the underlying biomedical model. The cancer problem is structured and defined by the disease (instantiated by a diagnosis and stage), a series of pertinent prognostic factors known to impact on the disease behaviour, and therapy options that have already been shown to provide an improved outcome in a similar group of patients. Other patient factors are considered, such as an assessment of the patient's ability to realise the benefits on offer [7].

Some research has investigated the level of knowledge required to verify specialist ontologies ([8, 9]). These previous findings suggest that domain experts are not necessary in the process and that crowd-sourcing from a range of backgrounds will provide similar verification results. The question remains whether this conclusion can be sustained in the setting of an expert domain that is not mainstream and only episodically encountered. The quality of the verification of non-standard ontologies, as well as their repeatability in comparison with SMEs needs to be established.

Crowd-sourcing is an on-line form of job completion drawing on the collective skills of a large group of people. The main concept of crowd-sourcing is that a crowd of unskilled individuals work on small tasks. Using crowds, project managers are able to reach more people, reducing risk of a single point of failure. This should improve the standard of job completion. The project manager now has a wider variety of responses; allowing them to find the best from a wider and more varied set of responses. *Microtasking* is a type of crowd-sourcing which involves breaking a large piece of work into many smaller tasks for completion rapidly by a crowd. Microtasks include proofreading passages, scanning images and transcribing audio files. Managers pay a small fee, for example, ten cents, to each participant if their response is worthy. This makes microtasking not only time efficient, but also mistake free and cost effective.

We considered the notion of verifying an ontology of medical terms by making a survey available to domain experts and the general public. To explore the the benefits of crowd-sourcing the opinions of domain experts and the general public in examining an ontology, the functionality of the AMAZON MECHANICAL TURK¹ was considered an acceptable vehicle to provide an insight into the level of specificity and verifiability of the knowledge of domain experts.

The tight coupling of cancer work flow with cancer diagnosis and stage followed by resultant cancer treatment with objective end points [10] indicates a medical specialty that is well suited to the development of a domain ontology. Furthermore, radiotherapy for benign conditions is uncommon, so it is only the patient with cancer and their families who will become intimately aware of Radiation Oncology and so be likely to have any detailed knowledge of the expert domain. The rest of the crowd should be unaware of any Radiation Oncology ontology. The question of the usefulness of crowd-sourcing in ontology verification will be addressed by allowing the public and radiation oncologists to assess the quality of a Radiation Oncology ontology. The survey will characterise people by medical qualification, oncological specialty, and experience with cancer.

The general public users of the Mechanical Turk survey will provide an insight into uninformed perceptions of the Radiation Oncology knowledge structure, which should differ greatly from the survey results of radiation oncologists. An adequate sample size should identify the authentic and the invalid or inadequate answers. The deliberate use of deliberately reversed and disjointed statements will assist identification.

Since the Mechanical Turk attracts its participants by payment, the allocated funding will finance 175 participants. Our experimental design calls for crowd-sourced workers to be surveyed who, on completion of the first 20 questions, will be given an option to continue and so collect the full payment.

We believe that excessive data cleansing and small population samples have not provided a conclusive view of crowd-sourcing capabilities [9]. By using a similar experimental protocol, we hope to demonstrate a contrary view using a non-standardised ontology written by a domain expert in the area of Radiation Oncology.

2. The Motivation for the Project

The validation of medical expert domain knowledge structures is a problem because medical domain experts are unfamiliar with the production and structure of the ontologies produced by informaticians [11]. The only known

¹ The history of “The Turk” can be found [here](#). The modern electronic version is found [here](#).

example of an attempted knowledge structure by a domain expert did not fit the characteristics expected of an ontology [1] but experience with an ontology editor, Fluent Editor 2² demonstrated that Controlled Natural English could be used to produce an ontology. The statements, while not plain English, were sufficiently clear to bridge the semantic gap between informatician and oncologist [12]. Following is a series of English Controlled Natural Language (ECNL) knowledge statements about Radiation Oncology:

COMMENT: HOSPITAL AND MEDICAL.

Every doctor is a human.

Every doctor has-license[medicallicense].

Every medicaloncologist is a doctor.

Every radiationoncologist is a doctor.

Every radiationoncologist has-license[radiationlicense].

Every doctor is an institution.

Every hospital is an institution and has-license[hospital] and has-employed[doctor].

COMMENT: PATIENT.

Every patient is a human.

Every patient has-anatomy[anatomy].

This possibility of an ECNL bridge became the motivation for the project, where the development of a crowdsourcing system of polls could allow the validation of any ontology, but particularly this Radiation Oncology ontology.

This experiment was undertaken as a summer project at the University of Wollongong with a structure of the project that called for the authors to develop the industry standard skills of SQL, Google AppEngine, Memcache HTML forms and CSS while developing the methodology to :

- investigate and quality assure the simple XML structure with the domain expert, parse items and sub-items into relationships that maintain the hierarchy of levels and develop a list of tuples
- use the tuples of relationship in combination with text to develop a series of possible and reversed plain ECNL statements that are displayed for human readers to select as most appropriate
- developed a Google Apps site which maintains a database for the recording of some simple characteristics (professional, doctor, oncologist, radiation oncologist, cancer patient, radiotherapy patient) and the on

² *Fluent Editor 2* is provided for educational use by *Cognitum Sp. z o.o.*

line survey tool with log in screen, memory for user return and completion, data storage of choices and integration into Mturk.com

- the on-line community, Mechanical Turk, is used as a base for the attraction and registration of participants while the website records their demographic details and answers
- data recording was constructed to attempt answers for some questions - what relationships are validated by the radiation oncologist? does any other group have a similar appreciation? Are there relationships that have a contrary meaning in the non-radiation oncology community? The results of the surveys will be reported separately.

3. The XML patient specification

The patient specification used was a hand crafted XML specification of a particular patient's oncology history. The specification is available in Appendix J of the thesis document available on-line (<http://ro.uow.edu.au/theses/3744> [1]). The specification was developed by a domain expert, but remains unverified.

Listing 1. XML Specification of Patient Oncology History [1]

```

1 <Organisation>
2   <Department>ILLAWARRA CANCER CARE CENTRE
3     <Doctor>ALEXIS ANDREW MILLER</Doctor>
4   </Department>
5   <Patient>
6     <PatientID>
7       <LastName>LOLLIPOP</LastName> * a pseudonym
8       <FirstName>LARRY</FirstName>
9       <Gender>
10        <Male>YES</Male>
11      </Gender>
12      <UniqueID>99999999
13      <DateOfBirth>32/13/1949</DateOfBirth>
14    </UniqueID>
15  </PatientID>
16  <Demographics>
17    <Race>CAUSACIAN</Race>
18    <Education>TERTIARY</Education>
19  </Demographics>
20  <BiologicalMilieu>
21    <Physiology>
22      <RenalFunction>
23        <CreatinineClearance>80

```

```
24 <CreatinineClearanceAlgorithm>COCKCROFT-
    GAULT</CreatinineClearanceAlgorithm>
```

4. Producing tuples from the XML specification

The XML code for the patient specification was parsed to exclude the instantiation of semantic entities and reveal the imputed relationships between various entities.

Listing 2. Parsed XML Specification of Patient Oncology History [1]

```
1 <Organisation>
2   <Department>
3     <Doctor>
4   <Patient>
5     <PatientID>
6       <LastName>
7       <FirstName>
8       <Gender>
9         <Male>
10      <UniqueID>
11      <DateOfBirth>
12      <PlaceOfBirth>
13    <Demographics>
14      <Race>
15      <Education>
16      <Employment>
17      <Housing>
18    <BiologicalMilieu>
19      <Physiology>
20      <RenalFunction>
21      <CreatinineClearance>
22      <CreatinineClearanceAlgorithm>
```

This hierarchical list was then parsed into tuples. A sample list of tuples in the form “PARENT:CHILD” is provided in Table 1. These tuples are used to construct the statements in the survey.

5. Constructing the survey questions from the tuples

These tuples were then used to prepare a set of statements to be used as the multiple choice answers for the survey. The answers are similar to ENCL statements. The tuples (PARENT:CHILD) are fed into a predetermined list of text statements (e.g., for the first statement in Figure 1, “**CHILD**” IS A TYPE OF “**PARENT**”) for the survey with appropriate substitution. This allows the hierarchical objects to be assessed in many different types of relationships in the survey. All the multiple choice answers were generated and saved as text files to be uploaded

Organisation:Department
 Organisation:Patient
 Department:Doctor
 Patient:PatientID
 Patient:Demographics
 Patient:BiologicalMilieu
 Patient:Diagnosis
 PatientID:LastName
 PatientID:FirstName
 PatientID:Gender
 Demographics:Race
 Demographics:Education
 Demographics:Employment
 Demographics:Housing
 BiologicalMilieu:Physiology
 Physiology:RenalFunction
 RenalFunction:CreatinineClearance
 CreatinineClearance:CreatinineClearanceAlgorithm

Table 1. Examples of Tuples derived from the patient specification

Which relationship best fits a Patient and an Organisation?

Patient is a type of Organisation

Patient is a part of Organisation

Patient is a part of an Organisation

Each Organisation has Patient

Each Organisation has a Patient

Patient is served by an Organisation

Patient is not a type of Organisation

Patient is not a part of Organisation

Patient is not a part of an Organisation

No Organisation has Patient

No Organisation has a Patient

Patient is not served by an Organisation

None of these

Figure 1. Web page appearance with standard options for the tuple “Organisation:Patient”

to the Google App Engine where the multiple choice answer pages were automatically generated and served to users to await an answer. The user is required to select the most accurate statement describing the relationship between the tuple. An example of the page is provided in Figure 1.

Answer a short survey

We are conducting an academic survey about the general understanding of medical terms. We need to understand what you know about medical terms. Select the link below to complete the survey. At the end of the survey, you will receive a code to paste into the box below to receive credit for taking our survey.

Survey link: <http://coi-hs-survey.appspot.com/signup?mturk=yes>

Provide the survey code here:

Template Note: In order to verify that the Worker completes your survey, you may require the Worker to enter the code usually provided at the end of the survey from the survey website you use. You can also ask the Worker to provide their Worker ID in the survey on the external site to verify the Worker completed the survey.

Figure 2. Example of the survey used on mTurk

6. Setting up the Crowd-sourcing environment

We opted to use the crowd-sourcing platform, Mechanical Turk. There are many other crowd-sourcing platforms that offer similar features. In future work, we may assess the variability in results obtainable from different platforms. The process of developing a microtask is straightforward. We designed a survey task on mTurk that asked users to visit our custom survey application. On completing the survey users receive a unique code used on the mTurk website to confirm that the user had completed the survey. An example of the survey has is in figure.2.

A number of combinations of worker payment and request types were tried. Initially we opted to used master workers. Master workers are workers that have been verified as being competent through the completion of other jobs. We found through experimentation that these workers required higher payment, though often spent longer completing tasks. Table.2 shows the average time and effective hourly rate of mTurk workers completing this task.

Survey	Reward per task	No. Tasks	Tasks Complete	Tasks Rejected	Masters	Average Time	Average Rate
1	11c	50	3	0	Yes	2m 57s	\$2.237
2	25c	50	3	0	Yes	5m 2s	\$2.90
3	50c	10	10	0	Yes	6m 9s	\$4.87
4	35c	20	6	0	Yes	3m 17s	\$6.39
5	40c	20	0	0	Yes	NA	NA
6	45c	20	7	0	Yes	9m 51s	\$2.74
7	45c	50	2	0	Yes	4m 39s	\$5.806
8	30c	50	50	2	No	6m 42s	\$2.687
9	40c	50	50	1	No	6m 55s	\$3.47
10	20c	50	50	2	No	6m 14s	\$1.92

Table 2. MTurk Survey

7. Conclusion

From an educational point of view, summer projects have value for the research produced, for the development of students and for the propagation of the academic area [13, 14].

This report describes a student project based on the use of crowd-sourcing to assist in the verification of an expert domain ontology which is expressed in ECNL after derivation from a hierarchical specification of medical knowledge in a crude XML format.

The website will be able to demonstrate the extent and agreement of public knowledge in the specialist domain of Radiation Oncology. It also provides a mechanism to assess the public's understanding of domain specific terms to inform medical literature. The use of payment on Mechanical Turk enables the procurement of many participants for a small outlay. In our acquisition phase, increasing the sum offered to 35 cents per survey seemed to increase the uptake rate significantly. The framework developed can also be deployed to engage domain experts without informatics knowledge to verify knowledge statements. While this is expected to be the ontologically useful activity, it was hypothesised at the start that the offer of remuneration was unlikely to influence completion by oncologists. Since this is a crowd-sourcing activity, involvement is voluntary and so no ethics approval was required.

It is envisioned that the work will be extended in several ways. At present, the use of parent:child pairs is simple. In future, we will assess responses to see whether eliminating unselected choices will increase the survey process. The automation will also allow the concatenation of more levels to reveal more reasoning than the single relationship seen already. This work is likely to engage domain experts more closely.

While entities such as ORGANISATION, PATIENT and DEMOGRAPHICS are easily defined and can be seen to be connected, in a formal ontology there is a need for a relationship. In this work, the term `is_a` has been applied generically to denote a connection, without dwelling on a more precise nature of the relationship. Modification of the web site could ask for a user to define the relationship term that they believe connects two entities. This would support another round of polling to determine the most acceptable terms.

The results provided by the website will be reported separately.

The natural language statements relating concepts are available on a website for verification, and readers are invited to complete the survey at <http://coi-hs-survey.appspot.com/> to contribute.

References

- [1] A. A. Miller, “Developing an Ontology for Radiation Oncology,” *University of Wollongong Thesis Collection*, Jan. 2012. <http://ro.uow.edu.au/theses/3744>.
- [2] A. A. Miller, “Electronic medical record,” in *Informatics in Radiation Oncology, Imaging in Medical Diagnosis and Therapy*, pp. 27–38, CRC Press, Aug. 2013. <http://dx.doi.org/10.1201/b15508-6>.
- [3] S. Spreeuwenberg, J. van Grondelle, R. Heller, and G. Grijzen, “Using CNL techniques and pattern sentences to involve domain experts in modeling,” in *Controlled Natural Language*, pp. 175–193, Springer, 2012.
- [4] P. Rocca-Serra, A. Ruttenberg, M. J. O’Connor, P. L. Whetzel, D. Schober, J. Greenbaum, M. Courtot, R. R. Brinkman, S. A. Sansone, R. Scheuermann, and B. Peters, “Overcoming the Ontology Enrichment Bottleneck with Quick Term Templates,” *Appl. Ontol.*, vol. 6, pp. 13–22, Jan. 2011. <http://dl.acm.org/citation.cfm?id=1971674.1971679>.
- [5] N. F. Noy, M. Crubézy, R. W. Fergerson, H. Knublauch, S. W. Tu, J. Vendetti, and M. A. Musen, “Protégé-2000: An open-source ontology-development and knowledge-acquisition environment: Amia 2003 open source expo,” in *AMIA Annual Symposium Proceedings*, vol. 2003, p. 953, American Medical Informatics Association, 2003.
- [6] A. A. Miller, “Clinical Information Systems in Oncology - Making a Difference to Patient Outcomes,” *Health Care & Informatics Review Online. Open Access*, vol. 7, no. 1, 2003. <http://www.hinz.org.nz/journal/2003/12/Clinical-Information-Systems-in-Oncology---Making-a-Difference-to-Patient-Outcomes/888>.
- [7] Medical Oncology Group of Australia, “Understanding the cancer treatment process.” <http://www.moga.org.au/patients-carers>.
- [8] A. Dumitrache, L. Aroyo, C. Welty, R.-J. Sips, and A. Levas, ““Dr. Detective”: combining gamification techniques and crowd sourcing to create a gold standard in medical text,” *CEUR Workshop Proceedings*, (Aachen), pp. 7–22. <http://ceur-ws.org/Vol-1030/paper-02.pdf>.
- [9] J. Mortensen, M. Musen, and N. F. Noy, “Developing Crowd sourced Ontology Engineering Tasks: An iterative process,” *CEUR Workshop Proceedings*, (Aachen), pp. 46–55, 2013. <http://ceur-ws.org/Vol-1030/paper-06.pdf>.
- [10] A. A. Miller and A. K. Phillips, “A contemporary case study illustrating the integration of health information technologies into the organisation and clinical practice of radiation oncology.,” *The HIM journal*, vol. 34, no. 4, p. 136, 2006.
- [11] S. Spreeuwenberg, J. van Grondelle, R. Heller, and G. Grijzen, “Design of a CNL to involve domain experts in modeling,” in *CEUR Workshop Proceedings*, vol. 622, (Marettimo Island, Italy), 2010. <http://ceur-ws.org/Vol-622/paper13.pdf>.

- [12] T. Kuhn, “How to evaluate controlled natural languages,” *arXiv preprint arXiv:0907.1251*, 2009. <http://arxiv.org/abs/0907.1251>.
- [13] D. Lopatto, “Survey of undergraduate research experiences (sure): first findings,” *Cell biology education*, vol. 3, no. 4, pp. 270–277, 2004.
- [14] D. Lopatto, “Undergraduate research experiences support science career decisions and active learning,” *CBE Life Sciences Education*, vol. 6, no. 4, p. 297, 2007.