tripleC 4(2): 293-303, 2006
ISSN 1726-670X
http://tripleC.uti.at

Interpretations of Ontologies for Breast Cancer

Srinandan Dasmahapatra and Kieron O'Hara

School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, UK
{sd, kmo}@ecs.soton.ac.uk

Abstract: There are increasing efforts directed at providing formal frameworks to consolidate the widening net of terms and relations used in medical practice. While there are many reasons for this, the need for standardisation of protocol and terminology is critical, not only for the provision of uniform levels of health care, but also to facilitate medical science research. In the domain of breast cancer pathology, a summary of current practice by the World Health Organisation states that the variability of the evidence archive (inconsistencies in describing microscopic appearances of phenomena, different diagnostic thresholds for working pathologists) is chief among the barriers to the medical understanding of the symptoms and development of early cancers. Such variability is acknowledged across specialist fields of medicine, motivating standardisation of terminologies for reporting medical practice. The desideratum of making these standards machine-readable has led to their formalisation as ontologies.

Ontologies are computational artefacts designed to provide representations of a domain of interest. Thus, the representation must be a formal description so that it can be encoded, and reused, allowing navigation of the key concepts recorded and retrieval of information indexed against it. This brings the required standardisation by offering a set of labelling options to record observations and events encountered by medical professionals.

Given the twin goals of ontologies -- representation and standardisation -- this paper will consider the key question of their design in the context of the use by experts, of information handling applications built around them. We build on our experience in developing ontologies for decision support software in the area of breast cancer diagnosis and treatment. We will also examine, from this perspective, the suggestion offered in the literature that a set of metaphysically motivated questions should form the basis of ontology building as guarantors of fidelity to reality.

We find that ontologies intended to support medical practice can only be understood within the context of their intended use. The declarative framework within which they are encoded generates the hope that their meaning transcends the specific application context. We show, however, that these declarative statements are to be understood as end products of chains of procedural engagements between humans, materials and communitarian norms. It is only when this scaffolding that brings this representation into existence becomes routine and consensual (within the community that exchanges information indexed against it) that the concepts stand in for physiological states with independent dynamics. However, as the state of biomedical knowledge is always in a state of flux, and different institutions and practitioners may be out of sync with respect to such modifications, the concepts embedded in the ontologies are constantly subject to reinterpretation within the context of specific institutional practices.

Given the fragmentation of the patient's body when viewed through various specialised lenses, ontologies can provide placeholders for co-ordinating disparate viewpoints to provide suitable medical interventions. The extent to which such interventions reflect any underlying reality, as manifest in measures of their efficacy, is closely wrapped up in the regulatory apparatus of protocol-guided consensus making. The value of ontologies lies in their reflection of, and support for, the sense-making activities that constitute expertise, not in their transparent access to a metaphysical reality.

Keywords: Ontologies, representation, reality, medical protocols, consensus.

Acknowledgement: Engineering and Physical Sciences Research Council (EPSRC) grants GR/N15764/01 (AKT), GR/R85150/01 (MIAKT).

1 Introduction

There are increasing efforts to provide formal frameworks to consolidate the widening net of terms and relations used in medical practice[i]. While there are many reasons for this, the need for

CC: Creative Common License, 2006.

standardisation of protocol and terminology is critical, not only for the provision of uniform levels of health care, but also to facilitate medical science research. In the domain of breast cancer pathology, a summary of current practice by the World Health Organisation states[ii] that the variability of the evidence archive (inconsistencies in describing microscopic appearances of phenomena, different diagnostic thresholds for working pathologists) is chief among the barriers to the medical understanding of the symptoms and development of early cancers. Such variability is acknowledged across specialist fields of medicine, motivating standardisation of terminologies for reporting medical practice. The desideratum of making these standards machine-readable has led to their formalisation as ontologies.

Ontologies are computational artefacts designed to provide representations of a domain of interest. Thus, the representation must be a formal description so that it can be encoded, and reused, allowing navigation of the key concepts recorded and retrieval of information indexed against it. This brings the required standardisation by offering a set of labelling options to record observations and events encountered by medical professionals.

But the artefact's purpose has been stated in two ways: representation and standardisation. How do we resolve this ambiguity? Does it capture reality, or the state of the art of the biomedical domain? Does it make sense to distinguish between the two? Moreover, is the state of disorderly nomenclatural variability that is the state of medical practice to be merely described, or remedied? Given these natural questions, what should the underlying design principles be for ontology development?

Discussion of ontology and ontologies often brings out a number of tensions. It is agreed by all that the conceptual clarity of distinctions, and practitioner consensus are important aims of the development of a formal ontology. But that still leaves a lot of potential for methodological disagreement. This paper will consider the key question of ontology design in the context of their use by experts, of information handling applications built around ontologies. It will examine, from this perspective, the suggestion offered in the literature that a set of metaphysically motivated questions should form the basis of ontology building as guarantors of fidelity to reality.

We wish to clarify at the outset that it is not our position that bringing abstract principles to bear on ontology building necessarily makes the ontology useless, only that an understanding of the concrete context is a necessary pre-condition for such abstractions to provide any useful bite. Although the organisation of knowledge in an ontology does constrain the nature of the information one can draw from it, retrieval may also be insensitive to distinctions that stem from the a priori. While we do not endorse a merely `behaviourist' razor for entertaining only those distinctions which appear at the end-products of retrieval engines, the `cognitive' machinery organising the knowledge that needs representing should accommodate the ways in which revisions of knowledge take place. At the very least, the order imposed by a priori metaphysical concerns needs to be defended on the grounds of domain-specific relevance.

In this paper we examine the application and development of ontologies in real world domains. A bottom-up approach inevitably introduces a complication: any real-world application of an ontology will be relative to some task or group of tasks. However, as we shall show, the issues raised by task relativity are very strongly related to the issues raised by the question of usability in general, and these, we contend, cannot be ignored. Much of the work done by ontologists has addressed the lack of rigour in classification schemes used in medical practice or biomedical research, which is a worthy exercise. However, these have tended to not only take the starting point of their analysis – the existing thesaurus or compendium of commonplace terminologies –

CC: Creative Common License, 2006.

as fixed, they have also tended to ignore the kinds of applications these terminologies get used in. By ignoring the context within which this knowledge is gathered or reused, the tools of analysis left open for ontological scrutiny are the abstract principles of generic versus particulars, rigidity of properties, and so on.[iii] We will find these approaches ironically limited in the face of the richness of the domains they seek to capture, and whose *reality* they wish to certify.

2 The domain

We focus on the complex multidisciplinary area of breast cancer patient diagnostics and management, which raises many interesting challenges for the ontology designer. In particular, we will glean some of the lessons learned from a recent project MIAKT,[iv] which was intended to support the diagnostic decision-making process by means of ontology-mediated computational systems.

In many countries in Europe and the US, the decisions that direct management of patients with breast cancer are taken in a multi-disciplinary meeting. These involve radiologists (experts in interpreting imaging modalities such as X-ray, magnetic resonance images (MRI), ultrasound, etc.); oncologists who consider the molecular basis of disease progression and possible effects of intervention by specific drugs; histopathologists and cytopathologists who perform diagnoses based on tissues and cells extracted from patients and treated by various laboratory protocols; clinical nurses who know the medical history of the patient and their possible reactions to various treatment options; analysts of statistical information who present data on likelihood of treatment outcome and mortality; and surgeons who assume control over the interventions on the patient's body.

The ontologies developed in MIAKT were for the domains of ultrasound, X-ray and MR images, histopathology and general information to do with the patient's age, family history and so on. Ontologies are graphical structures, where nodes in the graph stand for concepts, labelled by words from the appropriate specialist jargon, and the edges are labelled by terms that express the relationships between the jargon terms. By embedding each of these terms in a graph, specific instances associated with a particular concept have to satisfy the relationships with which these concepts are associated. For each of the domains of expertise represented in ontologies, the multi-disciplinary meeting draws out possible relationships between these concepts for each case for which evidence is assimilated. Owing to the inherent uncertainty in mapping (say) visual descriptors on an ultrasound plate to a pathological condition, these multiple ontologies are coupled only in very general terms.

Given these characteristics of the domain within which our ontologies are put to work, we return to the key questions addressed in this paper: Do the concepts encoded in the ontologies represent "elements of reality" or are they merely some units of "conceptualisation by domain experts?" And how does the answer to this question determine the dynamics of ontology development?

3 Deliberative Dynamics of Ontology Development

If one structure, A, which is subject to change is to be compared with another, B, one formal way of verifying this parity is to verify it first at some stage, and then consider the conditions under which A and B undergo changes. Let us assume for the sake of argument that we have a scenario in which two different approaches to ontology construction have produced two ontologies, which can be mapped on to each other (in the extreme case, this map may even be the identity map). If these two ontologies undergo changes and remain structurally congruent, then that is a reason for thinking that the different criteria for ontological revision should be related by some means.

CC: Creative Common License, 2006.

With this motivating principle in mind, we do not only consider the structure of concepts and what they stand for. In this paper, we mostly base our discussion on possible *changes* in the ontological structures which would be the result of applying a methodology for making the ontology better suited for the domain it seeks to represent. Moreover, it is the organisation of concepts that remains the key focus of ontology-builders, since the choice of labels is determined by common usage of practitioners in the domain of interest. Ontologists are typically not averse to using "microcalcification" as a concept label if that is the term used by x-ray radiologists, glossing over the question of whether calcification is an object or a process.

In fact, this is the sort of fundamental question that philosophically minded ontology builders pose. Is the entity in question an endurant or a perdurant – does it exist as a persistent object in time, or does it temporally unfold? Is it the case that the term in question denotes a class or an instance of a class? Does concept A subsume concept B (usually interpreted extensionally by both computationally hurried as well as philosophically reflective practitioners[v])? Is an epithelial cell whose malfunction is linked to the cancerous condition a cell? A normal cell? Would a defined attribute possessed by an entity also be possessed by it in all possible worlds?

On the other hand, the concepts that are (often informally) exchanged within the medical community to anchor findings and projections also go through a process of rationalised development and structuring, but of a different kind. The sets of questions posed are of the form: Do architectural features on a histopathological slide for likely DCIS (ductal carcinoma in situ) correlate with prognostics? Are the morphological features of microcalcifications evident on x-ray slides of the type that would be interpreted as an indication or contraindication for biopsy? Do epithelial cells possess receptors for hormones estrogen or progesterone?[vi]

Unsurprisingly, the dynamics of (often non-formalized) ontology development in the biomedical profession are answerable to the need for clarity and effectiveness of decision-making on interventions. There exist professional exercises of "consensus conferences"[vii] which decide on standards of nomenclature and protocols for generating them reproducibly. For breast cancer, for instance, such consensus conferences establish the necessity or superfluousness of applying criteria such as architectural distribution of cellular units and organelles for verbal summaries of the appearance of biopsy slides. While the adequacy, or reality, of the descriptors to the tissue is not in question in such a case, their influence on the usefulness of decision-making and prognostic judgement is the deciding factor.

In contemplating modes of intervention, it is the assessment of prognosis that guides decision-making, thus rendering some conceptual labels inherently dispositional. It is thus not only what a given concept denotes in an extensional sense that is relevant; the underlying dynamics that are likely to unfold and bring about a future manifest condition are also incorporated.[viii]

To sum up, even if two ontologies developed in accordance with two sets of governing principles, one focused on the philosophically-informed description of reality (top down), the other on task-related communication and efficacy of intervention and practice in the domain (bottom up), were isomorphic, they are likely to diverge under development and maintenance. Different sets of principles are used to detect faults and propose revisions.

A truism: for ontologies to form the framework for a software architecture that supports decision making in the biomedical domain, they should provide the kinds of handles to information that medical practitioners are comfortable with. By indicating the diverging demands of the two modes of ontology modification, philosophical scruple should be subservient to the representational choices made that

CC: Creative Common License, 2006.

organise knowledge in a manner that is most effective for biomedical reasoning for interventions – *if* an ontology is intended primarily for practical use in decision-making.

4 Logic of Class

In medical practice, and consequently, in the storage and representational medium that supports this practice, particular patient cases are slotted into class labels on the basis of a generic pattern general enough to be recognised as such in other cases. When the instance and the class it belongs to are recorded in the ontology in a specific representation language, the semantic apparatus is declarative. Thus, the procedural roots of declarative representations are concealed if one focuses solely on the reification implicit in the ontology and the instances catalogued within it.

The necessity of a semantics such as model theory has been amply demonstrated in the history of knowledge representation in artificial intelligence in order that sound conclusions can be drawn from codified premises by a mechanised reasoning apparatus. However, in order to accommodate changes to the structure of information that may be necessitated by the changing state of biomedical knowledge and practice it is important to remind ourselves of the chain of evidence-based reasoning as well as the institutional sanction that cements particular modes of description and intervention. It is in this sense that the application of metaphysical apparatus to bear on the constructive endeavour of building ontologies[ix], valuable though it is, hinders an appreciation of how they come into being within even the biomedical framework, let alone its larger social matrix (which is not addressed in this paper either).

To take a simple example, one domain of medical expertise (partially) represented in one of the ontologies in MIAKT seeks to provide a framework for storing reports of pathological studies of slide preparations of biopsy tissue. When cells are extracted by means of a needle (fine needle aspiration cytology, FNAC), the customary practice of basing interpretations of cell morphology leads to a common false positive that the pleomorphic appearance of apocrine cells is designated malignant. The need to avoid the stress caused by a false positive diagnosis has led to the National Health Service (NHS) in the UK issuing the following guidance[x]:

> Recognition of the dusty blue cytoplasm, with or without cytoplasmic granules with Giemsa stains or pink cytoplasm on Papanicolaou or haematoxylin and eosin stains coupled with a prominent central nucleolus is the key to identifying cells as apocrine.

Knowledge representation support is required for the keeping of records and indexing the adherence to procedure or the need to reinterpret the evidence. Let us represent a laboratory protocol conducted on sample specimen x in a particular context t as L(x, t). Such a protocol might be the process by which tissue is extracted, fixed and stained, while relevant contextual issues could include the time, the particular conditions present at the time which might influence the outcome of the analysis, and so on. If that procedure L(x, t) elicits a behavioural response B(x, t), a predicative attribute P(x) is identified as the property encoded in the ontology, which, following Carnap[xi], we can connect with B using a biconditional (and not a conditional, since the falsity of the antecedent of a conditional statement makes the sentence logically true):

$$\forall x, t, L(x, t) \rightarrow (P(x) \leftrightarrow B(x, t)),$$

which we adopt as a key ingredient in all cases when a procedural route to the specification of a concept is required.[xii] Once the particular attribute P has been identified, it can participate in the rational decision making procedures independently of the procedure L or the behavioural characteristics B.

However, it is not clear how deliberations on whether P might be a rigid property or satisfies some other metaphysical requirement can be rationally constructed, still less validated in medical practice, without understanding and description of that practice informing them. Indeed, it is the practice here that is certain, while the metaphysical principles are uncertain and disputed.

CC: Creative Common License, 2006.

5 Routine contexts erase epistemological traces

Declarative naming is a shortcut that circumvents a variety of conditions and stages in both the passage from a coherent set of behavioural manifestations under standardised protocols, and the explicit adherence to protocols for, in the case of identification of malignancies from mammograms, assessing visual patterns in order to issue a name that summarises complex physiological states. In medical analytical laboratories for specific suspected diseases, the choice of fixative agents and stains can be routinely specified, and thus relegated to the cognitive background when reasoning over the implications of results for patient care. In such cases, the routine context makes it possible to associate (say) responses to particular immunophenotypic assays with specific states or classes; in effect, the experts sacrifice precision for speed and communicability of assessment. In the case of specifying explicit rules for composing morphological and other related cues to generate a label that is stable against interpretive variability, stability of reference is made equivalent to the establishment of inter-expert consensus.

The extra-linguistic origins of concepts have their roots in the materiality of laboratory practices, the signature traces on prepared tissue slices, on x-ray emulsion plates, in written transcripts of patient responses. To make demands on the reality of concepts apart from the empirical underpinnings of biomedical reason, as advocated by Smith, appears to miss out on the richness of phenomena and instead promote a reductive understanding of drastic simplifications brought about in designing ontologies. Merely endorsing the existence of the Real does not guarantee that a priori analysis of concept labels (e.g. whether the verb form or the noun form illustrates the property of a concept being an endurant or a perdurant) anchors the concept to a metaphysical reality. The appeal to modes of analysis that ignore epistemological trajectories, in both the material as well as institutionally dictated conventional forms, is feasible only to the extent that the contextual protocols that produce, or make manifest, the traces of that reality become routine, are no longer subject to recalibration.

It is this forgetting of the unremarkable epistemological scaffolding, such as the play of shadows on an ultrasound screen, that enables stable ontological characterisations to emerge. Such re-description is built upon prior ontological descriptions, ones that situate the contextual factors of the current crop of concepts emerging into ontological light within a stable context. This prior ontological framework is often (but not always) typified in the upper echelons of the hierarchical structure. In such cases, coarse-grained characterisations precede the refinements brought about by more detailed studies that are usually necessitated by the heterogeneity of a population lumped together under the coarse label. The multiplicity of fates, hence the difficulty of prognoses, of conditions such as DCIS, make many of the standard interventions rather shotgun-like. The set of biomarkers that need to be developed to sift through diverging phenotypic trajectories then has to go through the procedures of standardisation and widespread deployment and hence provide the means for forgetting. Only then will the characterisations of the subtypes as real cease to be evanescent or idiosyncratic.

6 Rule following and reproducibility

In order that reasoning with medical concepts can have empirical bite across a range of cases it is necessary that they be named in the same way by different practitioners, that concepts be recognised as such by different experts in the same specialist area when they are exposed to evidence elicited by *standardised* biomedical protocols. This extends the reach of procedural definitions and highlights its impact on the stability of denotation, which is fundamental to the semantics of the entire enterprise. In capturing domain knowledge in a representational formalism that guarantees the consistency of mechanized inference, Tarskian semantics, a model theoretic interpretation sealing the denotation of reference to referent, has been shown to be the only viable alternative. The parallel demands of

CC: Creative Common License, 2006.

uniformity of reference by medical practitioners and the model theoretic interpretation required for computational representation schemes have to be reconciled.

The difficulty of interpreting the complex visual features of medical images by different specialists – be they radiologists or pathologists – has been demonstrated by the presence of inter-expert variability in diagnosis. For example, variation in interpretation of the morphological appearance of microscopic slides is evident in studies on the reproducibility of diagnostic labels across pathologists, which has prompted the issuance of rule-following props as recommended practice by histopathologists in recognising particular syndromes. The following table demonstrates the inter-expert variability in diagnosing ductal carcinoma in situ (DCIS) and atypical ductal hyperplasia (ADH) under specially organised test conditions[xiii]:

With standardised protocols		Without standardised protocols	
# in agreement	% of cases	# in agreement	% of cases
6/6	58	5/5	0
5/6	71	4/5	20
4/6	92	3/5	50

When provided explicit protocols to adhere to in a particular session with a predetermined corpus of images to interpret, all six pathologists agreed on the same condition in 58% of cases that they looked at. In contrast, the pathologists not given explicit rules to go by, never achieved unanimity.

Expert histopathologists in everyday practice variably interpret the visual characteristics of a customarily fixed and stained tissue specimen on a microscopic slide. Yet, when they are presented with standardised protocols, typically those they may have themselves explicitly articulated as providing relatively unambiguous diagnostics, the degree of consensus, hence the degree of intersubjective coherence of terms of reference is greatly increased. The act of interpretation is thus responsive to the tacit or explicit ways in which rules for pattern stabilisation by the interpretive community are followed. The extent to which consensual knowledge can be retrieved when case notes are indexed by an ontology (itself a construct intended to encode consensual knowledge) is dependent on a set of antecedent processes of recognition which are demonstrably dependent on an accepted set of rules are explicitly followed. To the procedures in the laboratory or imaging apparatus that underlie stable descriptive labelling as described above, we must add the rule-following props that point the way to denotative coherence.

The advantage afforded by model theoretic semantics carries with it the declarative framework in which adherence to rules are obtained `for free,' in the model itself. The marked contrast to the required cognitive attention that seems to affect reference highlights the difference between knowledge constructs as computational artefacts and those that are expressed in medical practice. The class labels that are the core ingredients in any biomedical ontology appear at the end of a process of deliberation and projection onto possible future manifestations or physiological developments which are coarsely mapped out by associative means, and which are prone to inter-subject variation. This entails that the reality of specific classes in medical ontologies is not given a priori and hence will only problematically be available for top down analysis. The communitarian protocols and agreements that stable reference warrants are only partially mapped out by explicit encoding, say of the input features of these rule-following methods. For stable denotation, not only does the contextual preliminaries need to become routine and forgettable, the rules and dispositions need to be sanctioned within the community of practitioners as well.

CC: Creative Common License, 2006.

7 Norms and dispositions

The assessment of consensus, the routinising of protocols and their standardised uptake among health care institutions and the specification of recommended procedures all point strongly to the role of social frameworks and norms in the emergence of concepts. Routine protocols and ways of assessing findings from such assays determine the assignment of class labels to patient instance data. These classes are thus dispositional proxies of groups of instances that respond in coherent ways to such patterns of questioning. The set of questions is circumscribed by the state of knowledge within the medical community.

One institutional apparatus for assessing the state of this knowledge in the context of viable modes of intervention, recording and evaluation, is the consensus conference. These are periodic meetings held by experts in a particular subfield within pathology or radiology who examine past evidence to determine the need for retaining some descriptors and dropping others in order to track the extent to which refinements or coarse-graining of descriptive labels correlate with future development of such cases.

Once the set of descriptors is agreed upon even if provisionally, the task of recognising instances as belonging to these classes requires the normative rule-following steps addressed in the previous section to ensure unambiguous denotation. Extensional semantics requires an "interpretation function" that maps symbols in the representation scheme to sets, whose elements are the instances of the concept thus labelled[xiv]. However, this interpretation function is nowhere specified. We have isolated some of the key determinants of this function, and we find that much of it requires a communitarian, rule-following component. Without needing to profess a denial of reality, we highlight the normative dimension, itself based on interventional efficacy, as being key to sifting "good ontologies from bad."[xv]

It is in this context, that we claim that merely appealing to 'reality' as arbiter of such distinctions misses the dynamics of this complex process. It may well be that the adequacy of the descriptors' account of the world is the best explanation for their adoption, but in the absence of an analysis of their grounding in practice and consensus (for example by the consensus conferences) it cannot be ruled out that other factors, such as increased biomedical understanding in terms of micro- (*eg.*, molecular biomarkers) or macro-scopic (*eg.*, epidemiology), speed and consistency of application, communicability, the quality of instrumentation, or integration into other problem-solving processes, may also affect their adequacy. The grounding process needs to be understood, if only to rule out such procedural, non-metaphysical underpinnings for the discourse.

8 Formal demands on ontologies

The advocates of analytical metaphysics are not alone in demanding that the conceptual schemas encoded in ontologies be respectful of philosophical scruple. An older argument from the work of Fodor[xvi] might be examined in this light in which compositionality is deemed the least common component of what it means to be a concept. A breast carcinoma is most definitely a carcinoma which affects the breast; on this account we can put the concepts of "breast" and "carcinoma" together in the conceptual machinery in our brains and understand what "breast carcinoma" means. The formal apparatus of logic-based formalisms that are typically used to build ontologies quite naturally maintain compositionality by means of logical connectives like "and". However, as we have pointed out, despite this encoding, the epistemological means by which these concepts are reached are important for understanding what these ontologies are about. In addition, by paying attention to recording parts of this process, we can get a handle on the dynamics of ontology evolution, extremely important for ontology maintenance as the state of biomedical knowledge is undergoing rapid change, with completely new apparatuses (genomics, for example) and hence representational relations coming online.

CC: Creative Common License, 2006.

Core to this process is the act of recognition of instances to belong to concept classes. Typically, what are recorded are cases where cases are "good instances" of a class, prototypical elements exhibiting a set of characteristic attributes which determine class membership. Fodor points out that prototypes don't compose – good instances of "pet fish" (such as goldfish) are not good instances of "fish" (such as salmon), neither are they good instances of "pet" (dog). Thus, the concepts that we encode in our ontologies may formally, by virtue of their encoding, be concepts in the sense of being compositional, but as distillates of elaborate recognitional procedures, concepts they ain't.

The pragmatic and observational ways in which we have approached the building of ontologies in the breast cancer domain seem to miss the requirements of those who engage in analysis and who advise that ontologies be (a) informed by top down philosophical principles which are arrived at either a priori or as deep abstractions from a range of analyses of a range of disparate domains, and (b) understood as descriptions of an underlying reality that provides the backdrop to problem-solving. The question then arises of how unfortunate a lacuna this is.

9 Ontologies as sovereign or servant

In the wake of genomic sequencing, there has arisen an articulation of a new kind of science conducted in increasingly automated forms. In this context, the gene ontology and related representational projects have been approached by many as requiring a level of description that requires unambiguous semantics for autonomous machine-based science and discovery. While our paper does not address those issues, we feel that the level of clarity required for such a free standing computational enterprise has been confusingly aligned with the philosophical (metaphysical) drive to afford a transparent access to reality.

The ontologies that we have designed on MIAKT are intended to support medical practice. Success depends on their functioning in their intended context. The declarative framework within which they are encoded generates the hope that their utility transcends the specific application context. However, as we have indicated, these declarative statements are to be understood as end products of chains of procedural engagements between humans, materials and communitarian norms. It is only when this scaffolding that brings this representation into existence becomes routine and consensual (within the community that exchanges information indexed against it) that the concepts stand in for events and (physiological, for example) states with independent dynamics. However, as the state of biomedical knowledge is always in a state of flux, and different institutions and practitioners may be out of sync with respect to such modifications, the concepts embedded in the ontologies are constantly subject to reinterpretation within the context of specific institutional practices.

Given the multi-disciplinary setting, the strongest claim that can be made of the role of our ontologies is that they are bridges between the different worlds of expertise – but this is not a trivial task. Given the fragmentation of the patient's body when viewed through various specialised lenses, they provide a significant placeholder for the co-ordination of disparate viewpoints to provide suitable medical interventions. The extent to which such interventions reflect any underlying reality, as manifest in measures of their efficacy, is closely wrapped up in the regulatory apparatus of protocol-guided consensus making.

More to the point, if ontologies do reflect reality, then (a) that is a by-product of their meeting their primary purpose, and (b) not a direct measure of their value. An ontology that reflected reality well but couldn't be mapped onto expert problem-solving practices in the domain would surely be of limited utility – leaving aside the tricky question of how, given the lack of mappings onto practice, we could recognise that the ontology described an independent reality successfully. The value of ontologies lies in their reflection

CC: Creative Common License, 2006.

of, and support for, the sensemaking activities that constitute expertise, not in their transparent access to a metaphysical reality.

References

Bodenreider, O and Burgun, A (2005) "Biomedical Ontologies," In: Chen H, Fuller S, McCray AT, (eds.) *Medical informatics: Advances in knowledge management and data mining in biomedicine*, Kluwer.

Carnap, Rudolf (1936), "Testability and Meaning," *Philosophy of Science*, 3:420-468 and (1937) *Philosophy of Science*, 4:1-40.

Fielding, J. M, Simon, J., Ceusters, W., Smith, B. (2004) "Ontological Theory for Ontological Engineering," *Proceedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning* (KR2004), Whistler, BC.

Fodor, Jerry A. (2004) "Having Concepts: a Brief Refutation of the Twentieth Century." *Mind & Language*, 19:2947.

NHSBSP (2001) "Guidelines for Non-operative Diagnostic Procedures and Reporting in Breast Cancer Screening," Publication No 50 June 2001, available at http://www.cancerscreening.nhs.uk/breastscreen/publications/nhsbsp50.pdf

Page DL, Rogers LW (1992) "Combined histologic and cytologic criteria for the diagnosis of mammary atypical ductal hyperplasia." *Hum Pathol* 23: 1095–97.

Schwartz GF, Solin LJ, Olivotto IA, and Consensus Panel (2000), "Treatment of In Situ ductal carcinoma of the breast," *Cancer* 88: 946-954.

Schnitt SJ, Connolly JL, Tavassoli FA, Fechner RE, Kempson RL, Gelman R, Page DL. (1992) "Interobserver reproducibility in the diagnosis of ductal proliferative breast lesions using standardized criteria," *Am J Surg Pathol* 16:1133-43.

Schwartz GF, Solin LJ, Olivotto IA, and Consensus Panel (2000), "Treatment of In Situ ductal carcinoma of the breast," *Cancer* 88: 946-954.

MJ Silverstein, M.J. (*ed.*) (1997) *Ductal Carcinoma In Situ of the Breast*, Baltimore: Williams & Wilkins.

Smith, Barry (2004) "Beyond Concepts: Ontology as Reality Representation." In A. Varzi and L. Vieu (eds.), *Proceedings of FOIS 2004, International Conference on Formal Ontology and Information Systems*, Turin.

US Department of Health and Human Services, National Institutes of Health, National Library of Medicine (2002). Unified Medical Language System (UMLS).

World Health Organisation (2002) *Breast Cancer Screening*. IARC Press.

---

[i] US Department of Health and Human Services, National Institutes of Health, National Library of Medicine. Unified Medical Language System (UMLS), 2002.

See also O. Bodenreider and A. Burgun, "Biomedical Ontologies," In: Chen H, Fuller S, McCray AT, (eds.) Medical informatics: Advances in knowledge management and data mining in biomedicine, Kluwer (2005).

[ii] World Health Organisation. Breast Cancer Screening. IARC Press (2002).

[iii] B. Smith, "Beyond Concepts: Ontology as Reality Representation." In A. Varzi and L. Vieu (eds.), Proceedings of FOIS 2004, International Conference on Formal Ontology and Information Systems, Turin, 4-6 November 2004.

[iv] http://www.aktors.org/miakt

[v] $\forall x\ B(x) \rightarrow A(x)$, or the set of objects identified by A, B, which is often denoted A^I, B^I satisfy $B^I \subseteq A^I$. In implementations, it is necessary to add an existential quantifier as well, in order to prevent difficulties with empty sets.

[vi] These receptors stimulate ER/PR positive (E for estrogen, P for progesterone, R for receptor) cells to grow, and are therefore likely to respond to systemic therapy like administering tamoxifen, which blocks ER.

[vii] See, for example, "Consensus Conference on the Classification of Ductal Carcinoma of the Breast, April 22–25, 1999." Cancer, Volume 88, Number 4, 946-954 (2000).

[viii] This dispositional nature of some medical ontological concepts carries serious political implications, which we do not go into in this paper. The sharp increase in the number of cases identified as "ductal carcinoma in situ" following improvements in imaging technology have resulted in women being classified as having such a "diseased" state, yet in a sizable fraction of such cases, they never go on to actually contract cancer.

[ix] J. M. Fielding, J. Simon, W. Ceusters, B. Smith, "Ontological Theory for Ontological Engineering." Proceedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning (KR2004), Whistler, BC, June 2004.

[x] NHSBSP (2001).

[xi] Carnap (1936-1937), "Testability and Meaning," Philosophy of Science, 3:420-468 and 4:1-40.

[xii] We skip the details of translating the NHS specification into the first-order logical formalism above, and the subsequent translation of that expression into the syntax of description logic, which is the basis for the representation language in which the ontology is encoded, as it is not necessary for any of the arguments in this paper.

CC: Creative Common License, 2006.

[xiii] Criteria of Page et al (*Cancer* 1982; 49:751-758; *Cancer* 1985; 55:2698-2708; see also *Hum Pathol* 1992, 23:1095-1097), reported by Fechner in MJ Silverstein, Ductal Carcinoma In Situ of the Breast, (1997), original article Schnitt *et al* (1992).

[xiv] In endnote v, predicates are denoted extensionally by sets which carry a superscript I.

[xv] B. Smith, ref. iii.

[xvi] J. A. Fodor. “Having Concepts: a Brief Refutation of the Twentieth Century.” Mind & Language, 19:2947, (2004).

CC: Creative Common License, 2006.