Artif Intell Law (2008) 16:147–174
DOI 10.1007/s10506-007-9060-2

A transdisciplinary ontology of innovation governance

Wendy Ann Adams

Published online: 18 December 2007

Abstract Intellectual property law tends to be viewed as the only (or most significant) mechanism for achieving policy goals relating to innovation assets. Yet more creative and effective solutions are often available. When analysed from a transdisciplinary perspective, relying on the cooperative efforts of researchers from fields other than law, innovation governance is characterized not simply as the product of legal rules, but as a function of the interaction of legal rules, practices and institutions. When policy-makers seek to identify conditions under which the creation, use and exchange of innovation assets flourishes, care should be taken to focus on this combination of factors. This article describes the development of an ontology—a computerized method of representing knowledge as concepts and relations between concepts—to convey such understanding. Policy makers (and researchers) are provided with an organized, accessible representation of innovation governance that enriches their understanding and improves their decision-making.

Keywords Intellectual property · Transdisciplinary · Legal ontology · Semantic web

1 Introduction

Intellectual property law, at least in Western legal systems, is a complex instrument ostensibly[1] designed and implemented to achieve a simple objective: innovation assets are granted intellectual property protection on the basis that monopoly rights

[1] I used the term "ostensibly" here to indicate that a utilitarian rationale of intellectual property dominates in both doctrine and jurisprudence. Other theories of justification do exist. See Justin Hughes, "The Philosophy of Intellectual Property" (1988) 77 Geo. L.J. 287.

W. A. Adams (✉)
Faculty of Law, McGill University, 3644 Peel Street, H3A 1W9 Montreal, QC, Canada
e-mail: wendy.adams@mcgill.ca

in the short-term ensure a continual supply of these assets in the long-term.[2] Whether intellectual property law actually achieves this objective is a matter of considerable debate. Documentation of intellectual property's deficiencies, as well as recommendations for improvement, appears to increase in proportion to the economic value of the assets under consideration.[3] The submissions vary from suggestions for minor corrections to proposals for major reforms. What these submissions have in common, however, is a belief in the power of legal rules to bring about a particular policy objective in relation to innovation assets.

The difficulty with many of these submissions is that they tend to deal with intellectual property law in isolation, that is to say, as the only (or most significant) mechanism for structuring the behaviour involved in the creation and exchange of innovation assets, particularly in terms of providing the necessary incentives to create.[4] The fact, however, that innovation assets tend to be subject to intellectual property protection does not lead to the conclusion that related policy objectives are necessarily a matter of intellectual property law reform. Intellectual property rules are significant, but they must be viewed in context. When intellectual property is analysed from transdisciplinary perspective,[5] relying on the cooperative efforts of researchers from fields other than law, the necessary context is provided. A narrow focus on intellectual property rules in isolation is replaced with a broader focus on innovation governance, that is to say, on the full range of mechanisms available for achieving policy objectives. Viewed from this perspective, innovation governance is understood not simply as the product of legal rules but as a function of the interaction of legal rules, practices and institutions involved in the creation and exchange of innovation assets. Accordingly, when policy-makers and researchers seek to identify conditions under which the creation and exchange of innovation assets flourishes, care should be taken to focus on the interaction of these various factors.

The modest contribution of this article is to describe the development of an ontology of innovation governance as a conceptual framework for representing and communicating knowledge of intellectual property from this transdisciplinary perspective. The term "ontology" is used here not in its philosophical sense, as a study of being and the kinds of things that exist, but as the term used in the field of computer science to refer to the manner in which domain knowledge (a domain being a field of interest) is represented in the form of concepts and relationships between concepts. Ontological representation provides flexibility for representing and reasoning about the implications of concepts as they operate relative to each other.

[2] See e.g. U.S. Const. art. I, § 8, cl. 8; *Théberge v. Galerie d'Art du Petit Champlain*, 2002 SCC 34, [2002] 2 S.C.R. 336 at para. 30.

[3] A Quicklaw search of North American legal journals revealed 394 copyright or patent-related articles published in 2006, in comparison to 58 in 1982.

[4] Gold et al. (2004, pp. 301–306).

[5] "Transdisciplinary" is used here to indicate that researchers have collaborated to produce a methodology and perspective that transcends their individual disciplines, such as common vocabulary. Thus transdisciplinary work is distinguished from interdisciplinary work, where researchers' contributions to a collaborative project are based on discipline-specific methodologies and perspectives.

The development of the ontology is part of the work carried out by the Intellectual Property Modelling Group, a group of researchers associated with the Centre for Intellectual Property Policy at McGill University. Accordingly, Sect. 2 of this article briefly describes some of the work of the IPMG in order to provide the necessary background for understanding a transdisciplinary approach to analysing intellectual property and the decision to represent knowledge derived from this analysis in ontological form.[6] Section 3 describes the development of the ontology in detail, providing an outline of the structure of innovation governance as comprising the knowledge domains of legal rules, practices and institutions, as well as explanations for design decisions. The resulting ontology demonstrates that a transdisciplinary analysis of intellectual property provides policy-makers and researchers with a more sophisticated understanding of innovation governance. Focus shifts from a narrow consideration of legal rules in isolation to a broader focus on the interaction of the many legal rules, practices and institutions relevant to the creation, use and exchange of innovation assets.

2 A brief account of the work of the IPMG

The development of an ontological representation of innovation governance originates with the work of the Intellectual Property Modeling Group, a transdisciplinary group of scholars associated with the Centre for Intellectual Property Policy at McGill University. IPMG researchers from law, management, economics, ethics, philosophy, political science and the life and medical sciences are working together on a variety of projects that seek to expand our understanding of intellectual property protection. The IPMG is currently developing a methodology for modelling the manner in which laws, practices and institutions actually work together in creating or blocking the production of new knowledge in relation to biotechnological innovation.

The objective is to demonstrate that intellectual property rules do not operate in the abstract, and while policy-makers generally rely on legal reform to achieve policy objectives, such objectives may be better served by considering the interaction and complementary uses of legal rules, practices and institutions.[7] For example, policy makers may face constraints upon decision-making, whether legal, political or economic, and thus may need to consider alternatives other than legal reform for achieving a particular policy objective. In addition, certain practices and institutional contexts may present operational advantages in comparison with legal rules. In some circumstances, the best alternative may be for the state to provide suitable enabling conditions for market-based practices, or to encourage civil society organizations to play a more active role. For these reasons, policy-makers

[6] A description of the IPMG's methodology in its entirety is beyond the scope of this article. See ibid. at pp. 320–322; "Description of IPMG Research Approach", online: Centre For Intellectual Property Policy http://www.cipp.mcgill.ca/data/publications/00000047.pdf

[7] The IPMG does not recommend any particular policy objective. Instead, research is directed towards helping policy-makers and researchers in designing effective and creative strategies for using intellectual property to achieve desired objectives.

are encouraged to think more broadly about innovation as a governance issue rather than simply a matter of intellectual property law reform, and to give serious consideration to the choice of governing instrument.

Since the work of the IPMG focuses on innovation governance as a dynamic system rather than a static set of legal rules, systems theory forms the methodological basis of the IPMG's modelling process, which begins with a dynamic simulation of the process of biotechnological innovation. Dynamic simulation in (very) general terms is a method for modelling the value of variables in a given system over time. Each variable is related to one or more other variables in the system in the form of causal relationships. These causal relationships consist of probabilistic inferences in the form of either positive or negative feedback loops.

The terms "positive" and "negative" indicate whether the causal relationship is reinforcing or non-reinforcing. In a positive feedback loop, the relationship between the variables is reinforcing in that an increase or decrease in one variable results in a corresponding increase or decrease in another variable. In a negative feedback loop, the relationship between the variables is non-reinforcing in that an increase or decrease in one variable results in an inverse decrease or increase in another variable. Since positive feedback loops tend to increase or decrease indefinitely, a relatively stable system requires balance between positive and negative feedback loops.[8]

A common example used to explain the operation of dynamic simulation is the relationship between births, deaths, and population as positive and negative feedback loops, as indicated in the following diagram:

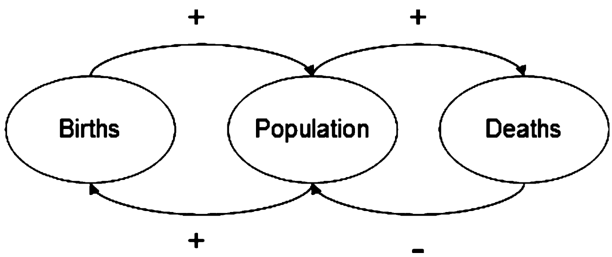

This diagram has two types of feedback loops. For example, the relationship between births and population is a positive feedback loop. As the number of births increases, the population increases. Population does not simply escalate out of control, however, due to the presence of a negative feedback loop between deaths and population. As the population increases, the number of deaths also increases, but then as the number of deaths increases, the population decreases. This leads to a corresponding decrease in the number of births. Thus stability in the system is maintained.

The dynamic simulation developed by the IPMG to model biotechnological innovation involves a much larger number of variables, but the basic principles are the same.[9] The IPMG has currently identified hundreds of variables and inferential relationships and is in the process of gathering the empirical data necessary to perform a dynamic simulation of biotechnological innovation under a variety of

[8] Sterman (2000, pp. 12–14, 107–133).

[9] See Gold et al., *supra* note 4 at p. 321.

real-world conditions. The results of the simulation will provide useful information as to how the many factors involved in a system of biotechnological innovation combine to produce a particular outcome. For example, the current work of the IPMG involves the dynamic simulation of three different models of policy objectives: (a) maximization of the level of innovation; (b) maximization of the level of access to innovation; and (c) maximization of the amount of scientific infrastructure. The same methodology, however, can be used to model an unlimited number of desired outcomes.

Following dynamic simulation of a policy objective, a mechanism is required that will permit policy-makers and researchers to use the results to structure legal rules, practices and institutions in such a way that the creation, use and exchange of innovation assets is consistent with the desired outcome. In effect, the task is to develop a transdisciplinary approach to innovation governance, one that requires a different methodology than conventional methods of reform which focus almost exclusively on the modification of legal rules.

For example, assume the desired policy objective is to produce a system of biotechnological innovation that increases public access to innovation. Assume as well that a particularly relevant variable (although only one among many) identified in the process of dynamic simulation is that of government use, understood as the manner in which a state can compel the use of a patented invention on what it determines to be commercially reasonable terms in the circumstances. A positive relationship exists between government use and access to innovation; as government use increases, so too does access to innovation.[10] In what way can legal rules, practices and institutions be structured to operate in a manner that contributes towards the desired outcome?

A traditional approach to innovation policy examines the possibilities in relation to legal rules. For example, patent law may permit the government to issue compulsory licenses for patented inventions financed by public funding in order to ensure that the technology is available at a reasonable price.[11] Policy-makers could also rely on practices and institutions, however, alone or in combination, to achieve the same objective. For example, the state could place conditions on public funding of research grants to universities, setting out the terms under which technology can be transferred to the private sector. The state could also work in cooperation with university technology transfer offices to develop a culture of public access to publicly-funded research.

Each of these alternatives has various advantages and disadvantages, a discussion of which is beyond the scope of this paper. The important point to note is that legal reform is not the only method, and may not be the ideal method for achieving innovation policy objectives. An ontological representation of innovation governance

[10] This example ignores for the moment the possibility of negative relationships between government use and increased public access to information. For example, if government use is made available on terms that fall below market value, even though they are commercially reasonable, this may act as a disincentive to engage in research and development and thus would eventually decrease access to innovation.

[11] See e.g. *Bayh-Dole University and Small Business Patent Procedures Act of* 1980, Pub. L. No. 96-517, 94 Stat. 3015 (codified as amended at 35 U.S.C. § 200–212 (1994)).

derived from a transdisciplinary analysis of intellectual property law helps policy-makers and researchers develop creative solutions by encouraging them to analyse the manner in which legal rules, practices and institutions interact in relation to the creation, use and exchange of innovation assets. Each variable from the dynamic simulation is associated with legal rules, practices and institutions (as denoted by the arrows in the diagram below) that facilitate the status of the variable as dictated by the results of the dynamic simulation. Thus a policy-maker or researcher is able to analyse a number of possibilities for achieving a policy objective from a set of complementary legal rules, practices and institutions (represented by shaded elements in the following diagram).[12]

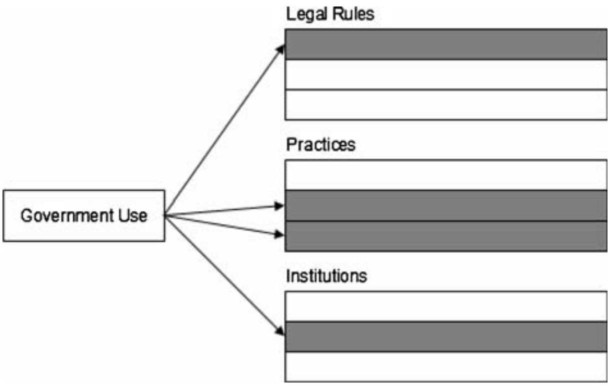

3 Developing a transdisciplinary ontology of innovation governance

3.1 Preliminary discussion regarding methodology

Before describing the process of developing the ontology, a few preliminary points should be noted concerning methodology. These points include the manner in which knowledge is extracted from a particular domain, the classification and organization of domain concepts, validation issues, development tools and the capacities and limitations of the finished product.

[12] Note that while the ontological representation of innovation governance is developed to implement the results of the dynamic simulation, the ontology can also be used as a standalone application. If the significance of a variable, such as government use, and the nature of its relationship to a desired policy outcome, such as increased access to technology, is uncontested (or validated through a method other than a dynamic simulation), then the ontology can be used independently of the analysis provided by a dynamic simulation. Policy-makers and researchers will still be able to benefit from the capacity provided by the ontology to understand and focus on innovation governance more broadly as a function of the interaction of legal rules, practices and institutions, and to use this understanding to identify alternative methods for achieving a particular objective in relation to the creation and exchange of innovation assets.

3.1.1 Extracting and organizing domain knowledge in ontological form

An appropriate source must be identified upon which to base a specification of a conceptualization of a knowledge domain. The source is likely to consist of documents that either constitute the entire knowledge domain (unlikely) or form a representative sample of the domain (more likely). The manner in which concepts are generated can vary from the most basic methods, in which human experts in the field identify concepts following a review of source documents, to more sophisticated methods using natural language processing software.[13]

Since the IPMG is a transdisciplinary academic research group, the most readily available and relevant source of documents from which to extract domain knowledge is that of academic commentary from various discipline-specific journal databases. As to the method of extraction, both time and financial constraints dictate that the ontology be generated using the basic method of physical (as opposed to automated) analysis of representative samples of the literature. Note that using a relatively straightforward method for data extraction does not affect the utility of the resulting ontology. The advantage of data obtained using more sophisticated, automated methods such as natural language processing is that the resulting product is likely to be more comprehensive.[14] On the other hand, depending upon the nature of the knowledge domain and the ontology's purpose, physical analysis of source documents by an expert in the field can result in superior identification and organization of concepts.

3.1.2 Design principles and validation issues

Once knowledge is extracted from a domain, the next task is to identify and organize the concepts and instances (or examples) of concepts in a manner that achieves the ontology's purpose. This task raises an important issue concerning the validity of the conceptual modelling process in two respects. The first concerns the validity of the classification process itself, that is to say, does the classification accurately represent domain knowledge? The second concerns the validity of the ontology as a final product; does the ontology serve the purpose for which it was designed?

The resolution of the second issue is outside the scope of this paper. The extent to which an ontology serves the purpose for which it was designed cannot be determined until the ontology has been implemented. Validity can then be assessed with reference to standardized metrics which are either subjective in nature, such as how the ontology has been reviewed by users,[15] or more objective, such as an assessment of how well the

[13] See e.g. the discussion of various methods of ontology generation in Ding and Foo (2002). For a discussion of ontology generation using natural language processing in the legal domain, see Lame (2004).

[14] The author is applying for additional funding to continue development of the ontology through the use of automated processes of ontology generation.

[15] See e.g. Supekar (2005). Reference to how the ontology has been reviewed by users is characterized as subjective because the possibility exists that users may evaluate the ontology as serving its stated purpose, and yet be in error in making this assessment.

ontology meets the requirements of evaluation tests that establish specific criteria,[16] including an analysis of general characteristics such as consistency, completeness, conciseness and expandability.[17] Reference may also be made to theory-based frameworks to assess various aspects of ontology design and functionality.[18]

The resolution of the first issue is within the scope of this paper, subject to the important limitation that few methodologies have been devised to guide the conceptualization process, and no consensus concerning methodology exists within the research community.[19] A number of approaches have been suggested for rationalizing this process of identifying and classifying domain concepts.[20] Some emphasize the overall significance of a strong, theoretical foundation.[21] Others focus more on technical aspects in setting out procedures to be followed,[22] or specifying design criteria according to which the validity of the ontology can be assessed.[23]

Another important issue regarding validation is the extent to which research in ontologies has largely failed to take account of classification theory, particularly as developed in the field of information science. On the one hand, ontologies cannot be equated to traditional classification systems such as taxonomies, hierarchies, thesauri, controlled vocabularies or terminologies.[24] On the other hand, the use of ontologies arguably represents a new development in the use of classification systems, and commentators in the field of information science present a persuasive argument that ontological development could be improved by drawing upon established research in the areas of knowledge representation, semantic relations, facet analysis and other topics within classification theory.[25] In particular, the validity of the ontology as an accurate representation of domain knowledge may be more readily established in circumstances where concepts are identified and specified using an established methodology such as formal concept analysis.[26]

Given the lack of consensus regarding appropriate methodology, the preferred alternative may be to consider existing legal ontologies as a reliable source of design principles, albeit with the caveat the ontology design in general should refer to a greater extent than is currently the case to research in the field of information science, particularly that of classification theory. Research on legal ontologies has produced a number of specifications of different conceptualizations of the legal

[16] See e.g. Asuncion Gomez-Perez (2001).

[17] See e.g. Gomez-Perez (2002).

[18] See e.g. Gangemi et al. (2005) and Burton-Jones et al. (2005).

[19] See e.g. Jones et al. (1998) and Fernandez (1999).

[20] See e.g. the discussion of various methods for building ontologies in B.C. Vickery (1997). For discussions of methodologies for building legal ontologies, see Van Kralingen et al. (1999) and Corcho et al. (2005, p. 142).

[21] See e.g. Guarino (1995).

[22] See e.g. Uschold and Grüninger (1996).

[23] See e.g. Gruber (1995, p. 907, 908).

[24] Jacob (2003).

[25] See, for example, Vickery (1997). See also Soergel (1999). Although note that the problem of communication exists in both directions. See Hjorland (2000, 2002).

[26] See e.g. Chi et al. (2006).

domain.[27] Assuming that design principles used to develop existing legal ontologies have at least a presumption of validity on the basis of their acceptance within the research community, these principles can be extrapolated and transferred to a new research environment to establish a presumption of validity in favour of the classification of concepts and relationships for the ontology of innovation governance. In particular, as is discussed below, design of the IPMG ontology of innovation governance has been informed by the work of Valente and Breuker in their development of a functional ontology of law.[28]

3.1.3 Development tools

Principles and design guidelines for ontology development do not specify any particular design environment. Many different languages and tools exist for developing ontologies.[29] The primary consideration in the transdisciplinary environment of the IPMG, where not all members are equally familiar or comfortable with basic concepts in computer science, is ease of use. Of equal importance, however, is that the finished product can be readily adapted for use by policy-makers and researchers without the necessity of understanding the underlying syntax (or language) in which domain knowledge is specified. Also, since the IPMG intends to make its projects readily available to researchers interested in further study, the development environment must conform to an established standard.

A widely accepted format for ontology specification, as well as one with exceptional utility, is the set of design principles established by the W3C, or World Wide Web Consortium[30] for what is known as the "semantic web". Currently, much of the information available on the Internet is retrieved through the use of keywords established for individual web pages in the form of meta-tags. As a retrieval mechanism, the lack of standardized vocabulary in meta-tags limits the efficiency of searching through the available information for relevant content. Efficiency can be improved, however, if web content is expressed in a manner that can be processed by software. This vision of the Internet, in which content can be accessed, interpreted and processed by software agents as opposed to entering keywords in search engines, is characterized by the phrase "semantic web".[31] The enabling technology includes defined formats, such as the Web Ontology Language or OWL, for identifying and organizing domain knowledge in the form of concepts, instances of concepts and relationships.[32]

[27] For a brief overview citing a number of examples of legal ontologies, see Visser and Bench-Capon (1998). See also Valente (2005, p. 65). For individual approaches to legal ontologies, see e.g. Mommers (2004), Hage and Verheij (1999), Breuker et al. (1997).

[28] Valente and Breuker (1996).

[29] For a brief description of various ontology development languages and tools, see Su and Ilebrekke (2002).

[30] The W3C is an international standards organization for the World Wide Web.

[31] Berners-Lee et al. (2001).

[32] For information on the OWL web ontology language and the semantic web, see "Web Ontology Language (OWL)", online:W3C http://www.w3.org/2004/OWL/

Given the availability of the Internet as a platform for exchanging data and applications amongst multiple users and the significance of the semantic web for future developments in organizing Internet content, the IPMG adopted the OWL format for the ontology. Ease of use during the development process was addressed through the use of Topbraid Composer, a design environment implementing the OWL standard. Developers can organize concepts using Topbraid Composer without the necessity of understanding the underlying syntax.

Adopting the OWL format also provides the advantage of modular design. The IPMG's ontology consists of three main components: (a) the enumerated, hierarchical structure of type–subtype relationships between concepts drawn from the three knowledge domains of legal rules, practices and institutions; (b) what might be termed the ontology's "data", in the form of the modalities, or individual legal rules, practices and institutions that are classified as instances of concepts; and (c) a faceted classification structure in which modality attributes are used to specify (and infer) relationships between the modalities and influence diagram variables, as well as relationships between the modalities themselves. Researchers are able to retain relevant components while discarding the others. For example, the extensive list of modalities, disaggregated from the hierarchical classification of concepts, is likely suitable for use in a wide range of research projects. On the other hand, both the hierarchical and faceted classification structures which specify the relationships between concepts, modalities and influence diagram variables are more likely to be project-specific.

3.1.4 Capabilities and limitations

Ontological representation can greatly enhance understanding of concepts in a knowledge domain, but expectations of performance must be realistic. An ontology is a method for representing, communicating and reasoning about domain knowledge. The format of the representation may be as simple as a series of two-dimensional drawings suitable for human analysis or expressed in standard syntax suitable for automated processing. In either case, representation of knowledge in ontological form assists (and generally improves) but does not replace human decision-making.

The IPMG makes no claim that its ontology of innovation governance can be used by policy-makers and researchers to produce ready-made answers to questions concerning patent law and biotechnological innovation. The role of the ontology is to help policy-makers and researchers develop a more sophisticated understanding of innovation governance that focuses not simply on state-based legal rules, but on the interaction of these legal rules with practices and institutions. The intent is not to replace human decision-making but to enrich human understanding. Thus the ontology should be understood not as a mechanism for making decisions but as a method for improving the manner in which decisions are made.

3.2 A principled approach to ontology design

Valente and Breuker recommend four basic design principles. The first principle is that the ontology should reflect a clear theoretical basis for the identification of concepts within a domain and the manner in which they are related. The theoretical foundation adopted for this task is the IPMG's transdisciplinary perspective of innovation governance as comprising the knowledge domains of legal rules, practices and institutions.[33] Ontology design thus begins by specifying representative concepts for each of the three knowledge domains and classifying modalities as instances of these concepts.

Valente and Breuker's second principle of parsimony is implemented by achieving the appropriate degree of granularity (or specificity) in specifying the concepts that represent domain knowledge. No fixed rule determines the number of subdivisions required, but ontologies should represent rather than replicate domain knowledge. Specification begins with an initial abstract concept in each domain, along with successive subdivisions until the appropriate degree of granularity is reached. When modalities can no longer be distinguished from each other by any meaningful criterion, they are grouped as instances of what then becomes the lowest-order or most specific concept.

The third principle is that the use of categories of knowledge in the formation of concepts is preferable to terms of art, which are more likely to be examples of types of knowledge. This principle is implemented by specifying concepts in purposeful rather than merely descriptive terms. For the knowledge domain of legal rules, the concepts represented are the types of decisions to be made in designing patent law, such as what subject matter to protect, rather than a list of the patent rules of any particular jurisdiction. For the practices domain, the concepts specified represent the objectives of particular practices, such as increasing access to medicines, or improving technology transfer between the public and private sector. Concepts derived from the knowledge domain of institutions represent the types of functions that institutions perform in coordinating activity to reach a common goal, such as educational functions (e.g., a university) and lobbying functions (e.g., a consumer rights organization).

Finally, Valente and Breuker's fourth design principle requires coherence in that the specified concepts, taken as a whole, comprise a framework describing the nature of the domain. Valente and Breuker do not specify any particular criteria for assessing coherence, but reference is made to a relationship between coherence and a functional approach to ontology design. For example, a functional characterization of medicine would refer to diagnosis and treatment of diseases, and thus the identified concepts should reflect these functions. The ontology adopts a functional and transdisciplinary approach by referring policy-makers and researchers to a range of modalities across the three domains of legal rules, practices and institutions from which to choose an ideal arrangement for implementing a particular policy objective. In other words, just as a functional ontology of medicine focuses on

[33] This may appear to be a self-referential assessment of validity, that is to say, that the measure of whether the appropriate theoretical basis has been chosen is a measure of the degree to which the IPMG members consider their theoretical approach to innovation governance to be valid. Given that ontology design, however, is a constitutive process involving an explicit (and shared) specification of a conceptualization (T. Gruber, http://www.ksl.stanford.edu/kst/what-is-an-ontology.html), to the extent that the IPMG members agree that the theoretical account accurately represents and explains domain knowledge, the self-referential nature of the classification process at the pre-implementation stage is unavoidable.

diagnosis and treatment, a functional ontology of innovation governance focuses on the implementation of policy objectives.

3.2.1 Identifying and organizing concepts in the knowledge domain of legal rules

Legal reasoning is typically used to determine the law applicable to a given set of circumstances. Accordingly, legal knowledge tends to be represented as a taxonomic classification of legal rules along with propositional logic concerning their application. The objective of the work of the IPMG, however, is not to identify legal rules applicable to a given situation, but rather to identify "modalities", that is to say, the types of legal rules, practices and institutions that can be used to ensure that the creation, use and exchange of innovation assets is consistent with a desired policy objective. Thus concepts representing the knowledge domain of patent law focus not on the legal rules themselves, but on questions of patent law design.

Fortunately, the traditional taxonomy of intellectual property law, which roughly divides the topic into subject matter to be protected, the nature and scope of protection granted, acts which amount to infringement and available defences, can be adapted to suit this approach. The task is to express these functions in the form of design decisions for which possible responses are specified as instances of these concepts. For example, if the type of design decision specified is determining the appropriate start date for the term of protection, instances of this type of decision would include the date of invention or the date of filing for patent protection.

Identifying the full range of modalities within the knowledge domain of patent law is a formidable task. A survey of existing patent law across jurisdictions will not suffice for two reasons. First, much of current patent law has been subject to the harmonizing influence of the universal minimum standards set out in the TRIPS Agreement,[34] and thus existing laws do not present wide scope for diversity. Second, knowledge of existing law does not represent all available knowledge in the domain; knowledge regarding the full range of possibilities (as opposed to actual legal rules) must be captured and represented. Thus the most appropriate source from which to extract knowledge is the work of jurists in the field of patent law. These commentators are likely to address not only law as it exists but also speculate as to the ways in which law might exist or ought to exist. This work is available in the form of written commentary in legal databases.[35]

Having located an appropriate source of relevant knowledge, the next challenge is determining a suitable method for extracting this knowledge from the large body of literature available, none of which is in format compatible with the IPMG's research methodology. Fortunately, law review articles dealing with patents in the field of biotechnology tend to follow a standard format whereby a problem in the area is first identified, followed by a proposed solution and discussion of its merits.

[34] *Agreement on Trade-Related Aspects of Intellectual Property Rights*, 15 April 1994, 1869 U.N.T.S. 299, 33 I.L.M. 1197 (1994) [*TRIPS*].

[35] The representative sample is currently limited to North American publications. Further development of the ontology will address this limitation by drawing upon a wider variety of sources.

Since the solution can be characterized as a possible response to a design decision, then logically each suggestion for reform, isolated from the context of the problem being analysed, identifies a potential modality within patent law.

Accordingly, a representative data sample is selected from the available commentary and the articles within the data set are analysed.[36] For each article, the proposed solution(s) to a problem is isolated from the surrounding discussion.[37] Once the various modalities of change within patent law are identified, the next task is to classify this knowledge as instances of concepts, which are specified and organized in hierarchical fashion from the abstract to the specific. To use the language of legal reasoning, the process is iterative in nature in that classification of the modalities is a product of both inductive and deductive reasoning; at times, the modalities suggest a specification of a concept, while at other times a specified concept suggests a grouping of modalities. To use the language of ontology design, three methods have been used to generate the classification structure: bottom-up, from the specific to the general; top-down, from the general to the specific; and middle-out, from the most important concepts to generalization and specialization.[38]

The final product is too complex to represent here in its entirety, but an example is provided in order to demonstrate both the process of classification and the end result. Throughout the development process in each of the three domains, complex knowledge is organized in the form of abstract and simplified concepts to facilitate communication. For example, at the top-most level of the knowledge domain of legal rules, decisions of regulatory design are divided into two basic types, those dealing with the design of the patent system itself (internal regulatory decisions), and those dealing with the manner in which patent law interacts with other areas of law, such as competition law (external regulatory decisions). A third top-level classification represents the fact that all regulatory design decisions are subject to constraints on the decision-making process. Thus the top-level of the knowledge domain of patent law is organized as follows:

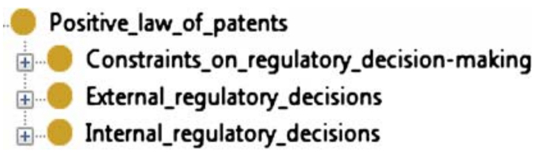

Working with the example of internal regulatory decisions, these decisions are further classified into type–subtype relations in accordance with the general

[36] Random sampling of the relevant databases was supplemented by tacit knowledge of the IPMG in selecting the works of authors acknowledged as influential in their fields of expertise. Random sampling was also limited to post-TRIPS commentary, given the radical nature of the changes introduced by this international agreement concerning universal minimum standards of protection (see *TRIPS*, *supra* note 34).

[37] Knowledge concerning the particular problem to which the solution is addressed is discarded as this knowledge does not contribute to representing the possible modalities through which changes can be made to patent law. This does not mean that the information is invalid (this is most certainly not the case), but merely that the information is not suited to the ontology's purpose, as the authors of the various articles are not working within the same conceptual framework as the IPMG.

[38] See, for example, *Foo, supra* note 13, p. 125.

regulatory structure of all intellectual property systems. Decisions must be made as to what type of asset to protect, the rights granted, the term of protection, exceptions to infringement, etc.:

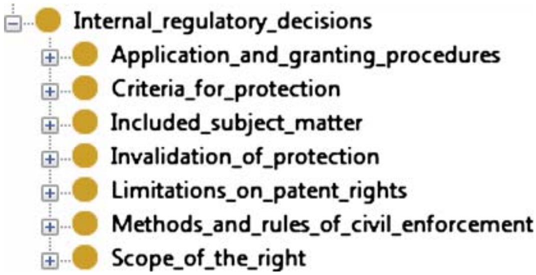

As the initial classifications are further sub-divided, design decisions move from abstract and general principles to provide increasingly specific information. For example, decisions concerning the scope of the right can be further sub-divided into those dealing with the term of protection and those dealing with the type of protection granted:

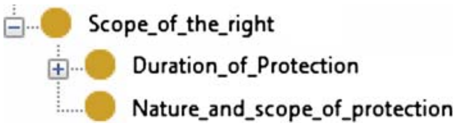

The difference between the classification of types of decisions and the range of possible decisions, or modalities, is critical to understanding the manner in which knowledge in the domain is represented. For example, the duration of protection is classified as a type of decision that can be made. This classification can be further sub-divided to represent decisions concerning the start of the term and decisions concerning the duration of the term:

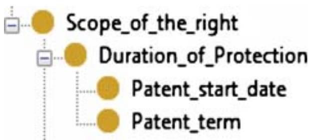

At this point, classification now reflects the appropriate degree of granularity. No meaningful criterion exists for distinguishing between modalities such as the date of the invention or the date of filing for the start date, or a period of 20 years or more than 20 years for the patent term. All that remains is to classify individual modalities in accordance with the lowest-order specification of concepts. For example, the modalities classified as possible decisions to be made regarding patent term are as follows:

◆ A._Permit_de_facto_patent_extension_techniques
◆ B._Do_not_permit_de_facto_patent_extensions_techniques
◆ C._Permit_de_facto_patent_extension_techniques_but_only_for_licensing_practices_after_term_expiry
◆ D._Provide_different_patent_terms_based_on_subject_matter
◆ E._Do_not_provide_different_patent_terms_based_on_subject_matter
◆ F._Provide_different_patent_terms_when_patents_are_based_on_higher_life_forms
◆ G._Provide_provisional_patent_rights
◆ H._Do_not_provide_provisional_patent_rights
◆ I._Provide_for_provisional_patent_rights_when_pending_applications_are_disclosed
◆ J._Provide_20_year_term_for_patent_rights
◆ K._Provide_20_year_plus_term_for_patent_rights

Moving in this manner from the most general characterization of the decision-making process to the most specific, the hierarchical type–subtype organizational structure in which the knowledge domain is represented provides a conceptual framework for understanding the design of patent law as a series of decisions to be made concerning the nature and scope of protection. Examined on their own, any one of the hundreds of patent law modalities derived from the literature review would communicate very little knowledge other than the fact of its own existence. When organized, however, as instances of specified concepts derived from the knowledge domain of patent law, the modalities can be understood as examples of particular design decisions to be made in any patent law regime.

Note that the language of the modalities does not reflect the manner in which the modalities are discussed by authors of the reviewed articles. Each modality, in order to suit the purpose for which the ontology is to be used, must be translated into a particular format. Since each modality represents an instance of a design decision, the modalities are drafted in imperative form as actions to be taken. Furthermore, the language of the modalities is standardized to reflect three types of patent law decisions: (a) those that permit (or prohibit) an action; (b) those that provide for (or do not provide for) a given feature; and (c) those that require (or do not require) a condition to be met. Often, these standard terms are qualified, usually by the circumstances required for their application, e.g., "Permit de facto patent extension techniques but only for licensing practices after term expiry".

3.2.2 Identifying and organizing concepts in the knowledge domain of practices

As with the knowledge domain of legal rules of patent law, the knowledge domain of practices is expressed in purposive terms. For example, what types of activities, such as licensing, does a firm initiate for the purpose of acquiring or exchanging intellectual assets? What types of activities does the state perform for the purpose of funding research and development in the public and private sector? What types of activities are undertaken in institutions, such as technology transfer offices, for the purpose of transferring intellectual assets from the public to the private sector? These are all factors that are as relevant, or in some cases more so, than the legal rules that grant patent protection to innovation assets.

Capturing the knowledge to be represented in the domain of practices is a simpler task than is the case for the domain of patent law. Unlike the representation of patent law modalities, where legal rules must be recast as design decisions, the modalities of relevant practices can be addressed on their own terms; each practice is, quite simply, a modality, given that each practice is a purposive action. As to the source from which knowledge of practices is extracted, many of the modalities are derived from management literature, reflecting the degree to which firms dominate the creation and exchange of assets produced by biotechnological innovation. Other modalities are derived from the data set used to specify the knowledge domain of legal rules.

Determining how best to represent the extracted knowledge is a challenging task, primarily because the domain of practices is inherently transdisciplinary in nature. Thus the organization of concepts cannot borrow from an established, discipline-specific taxonomy. One possibility would be to specify and organize practices based on the identity of the initiating actor. The selection of any actor, however, among several participants is inevitably an arbitrary choice. For example, even in the simplest of sales transactions, classifying the initiating actor is difficult. A seller who offers to sell does not operate in the abstract, but is responding to the perceived existence of a buyer who will accept the offer. The inverse is true in relation to the buyer. Which of these two actors can be said to be the initiating actor? On the other hand, to recognize all participants as initiating actors would lead to duplicate results.

Another possibility would be to specify and organize practices based on the sphere of activity as being public or private, state or market, but again this produces undesirable results. Many practices seem to bridge these spheres, leading once again to arbitrary choices and duplicate classifications. Taking again the example of an ostensibly commercial transaction, the identity of the parties may complicate the classification process. The seller could be an entity representing a public–private partnership between the state and the private sector, organized to perform what is normally a state function, such as the provision of military facilities and services. Should a practice carried out in these circumstances be characterized as public or private, state or market?

The specification and organization of concepts in the practices knowledge domain is ultimately based on the fundamental research objective of the IPMG. Reliance on theoretical inquiry in this instance is not misplaced, and does not call into question the validity of the classification as an accurate representation of the knowledge domain of practices. Classification theory from the field of information science provides an account of a pragmatic approach to classification. If no natural or best classification scheme can be presumed (which is the case in a constitutive exercise such as modelling domain knowledge from a particular perspective, for a particular purpose), and if the classification is required for a purpose, then consideration of that purpose is the most important part of the methodology.[39]

The work of the IPMG is directed towards capacity-building for policy-makers and researchers. Thus the fundamental objective is to enhance understanding of how laws, practices and institutions interact in the creation, use and exchange of innovation assets. Thus an appropriate initial classification is to divide the

[39] Hjorland and Pederson (2005).

knowledge domain of practices into those related primarily to the creation of innovation assets, and those related primarily to the use and exchange of these assets. This top-level classification does present some difficult decisions, given that practices improving the benefits of using or exchanging innovation assets are likely to encourage the creation of innovation assets, and visa-versa. Provided this reinforcing relationship is acknowledged, however, classification can proceed without the necessity of arbitrary choices and duplicate classifications:

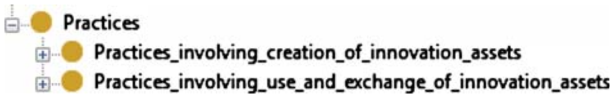

As with the knowledge domain of patent law, specifying and organizing concepts and classifying instances involves a combination of deductive and inductive reasoning, with concepts expressed in purposive terms. Following the initial subdivision, concepts are organized in type–subtype relations according to the function to which a practice is directed, such as product sales, technology transfer and access to medicines. Specification of these second order concepts in purposive terms avoids the problem of duplicate results that may occur when practices are specified in terms of the initiating actor or whether the practice is characterized as public or private in nature:

Further subdivision maintains this purposive approach by describing the relationship between the practice and a particular outcome. Logically, three possibilities exist: the practice may have a positive effect on the outcome, a negative effect, or no effect whatsoever. Practices having no affect on a particular outcome were few in number and were quickly discarded as irrelevant; why engage in a practice that does not affect an outcome? As to the remaining possibilities, practices having a negative effect on a particular outcome did not require further refinement, given that the effect in every case is to interfere to some degree with the outcome; these practices were thus described as impeding an outcome.

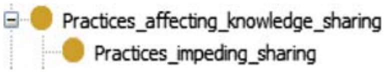

Practices having a positive effect, however, are capable of further refinement in that the positive effect can be characterized as improving, encouraging or increasing a particular outcome, depending upon the nature of the activity. For example, practices are characterized as encouraging an outcome if the practice establishes more favourable conditions as an inducement for actors to engage in a particular activity:

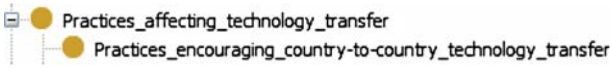

Practices are characterized as improving an outcome if the practice establishes conditions resulting in a qualitative improvement in a particular activity:

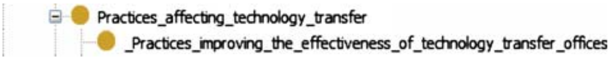

And finally, practices are characterized as increasing an outcome if the practice establishes conditions resulting in a quantitative improvement in a particular activity:

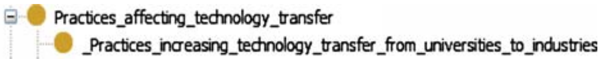

Once practices are classified by the effect of the practice in relation to a particular outcome, further subdivision (with a few minor exceptions) as to the type of practice no longer communicates meaningful information. All that remains is to classify the modalities as instances of the lowest-order concepts:

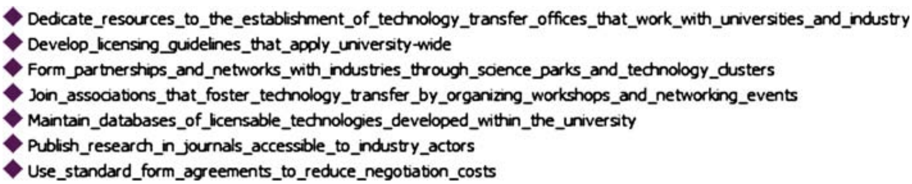

As with the knowledge domain of legal rules, care is taken to maintain the purposive nature of the representation by drafting the modalities in imperative terms. Thus each modality expresses a practice as an action taken to further a desired objective.

3.2.3 Identifying and organizing concepts in the knowledge domain of institutions

Knowledge of institutions is extracted from a different type of source and in a different manner than is the case for the knowledge domains of legal rules and practices. Many of the modalities are drawn from the datasets of academic literature used to develop these two domains, but by far the most significant source of information was the Internet; most if not all of the relevant institutions now possess an online presence. Data collection involves a process whereby the web pages of the most well-known institutions associated with each of the three domains are noted and further reviewed for lists of links to other institutions, which are in turn reviewed for additional links to other relevant institutions. This reiterative process is followed until a point is reached where additional searching of referred links fails to yield significant new institutional information and instead merely refers back to links and institutions identified earlier in the search process.

Once the data is collected, the knowledge domain of institutions, like those of legal rules and practices, is organized in purposive terms. The objective is to identify the type of activity performed by an institution as it coordinates individual efforts to achieve a common goal. The particular goal is less relevant as an organizing principle than the activity performed, given that a particular goal may be carried out by many different types of institutional activities. The exchange of scientific information, opposition to patent applications and harmonization of intellectual property laws are all specific goals, but these goals can be achieved through the coordinated efforts of individuals working through institutions performing a number of different types of activities, such as adjudication, administration, standard-setting, etc.:

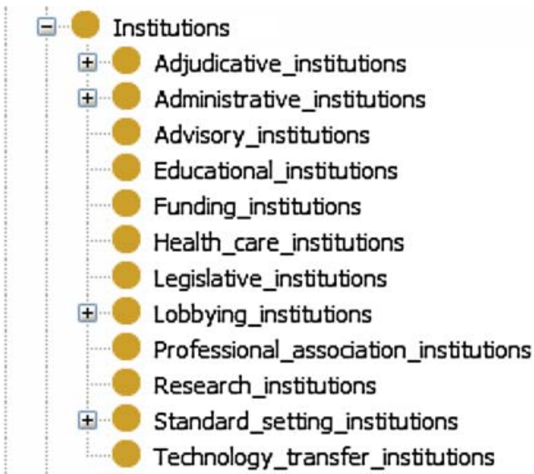

Of the three knowledge domains, that of institutions has the "flattest" design in that only four of the first-order classifications are further sub-divided into type–subtype relations to communicate meaningful distinctions. For example, "Lobbying institutions" is further sub-divided to communicate the significance of differences in membership as follows:

The complexity of a classification structure, however, or lack thereof, is irrelevant so long as the design does not sacrifice clarity for brevity. If further sub-division fails to communicate meaningful distinctions between modalities, this unnecessarily complicates the ontology design. Concepts derived from the knowledge domain of institutions simply reflect fewer meaningful distinctions between modalities, at least for the purposes for which the ontology is designed.

As with the knowledge domain of legal rules and practices, the individual modalities within each class are drafted in purposive terms. Since institutional activity generally coordinates individual efforts, the modalities reflect facilitative actions, as in the example of modalities classified as instances of public interest lobbying organizations:

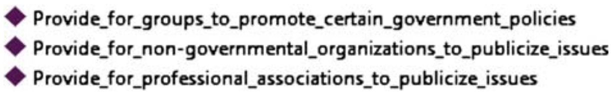

3.3 A transdisciplinary approach to achieving policy objectives

3.3.1 Identifying relationships between variables and modalities

The individual knowledge domains of legal rules, practices and institutions provide the relevant concepts and instances of concepts for describing and communicating knowledge of the actual legal rules, practices and institutions involved in the creation, use and exchange of innovation assets. The next step is to develop a method for representing in ontological terms the interactions between the dynamic simulation variables relevant to a particular desired policy objective, such as maximizing the level of access to innovation, and the various legal rules, practices and institutions through which a desired objective can be achieved. Accordingly, the first task is to identify relevant dynamic simulation variables. Following this, the second task is to specify relationships between these variables and various modalities, as well as between the modalities themselves.

3.3.1.1 Identifying relevant variables Recall that the current work of the IPMG involves the dynamic simulation of different models of policy objectives, the purpose being to identify variables in a system of biotechnological innovation that, when manipulated in a particular manner, are most likely to achieve a desired outcome. Of the hundreds of possible variables, the IPMG chooses the most relevant for a specific desired outcome by using a methodology known as the nominal group method.[40]

The method is ideally-suited to capturing the tacit knowledge of a transdisciplinary working group such as the IPMG. Each member of the group is asked to characterize all variables individually as being highly relevant, relevant or not relevant to achieving the policy objective being modelled. Variables which are consistently ranked by most or all of the group members as not relevant are then excluded, as are those for which no strong consensus emerges among group members concerning relevancy, e.g., approximately one-third of the group characterizing the variable as highly relevant, one-third as relevant, and one-third as not relevant. The final selection of variables includes those consistently ranked by the largest number of group members as being highly relevant.

Once this list is compiled, each selected dynamic simulation variable is linked to associated modalities across the three knowledge domains of legal rules, practices and institutions. To remove any possible bias in favour of legal rules, and to demonstrate the IPMG's emphasis on innovation governance as a function of the interaction of legal rules, practices and institutions, this exercise is completed by temporarily separating modalities from the hierarchical structure of concepts in which they are organized. The modalities are viewed as a single list.

Identifying the associated modalities and establishing the links between these modalities and the relevant dynamic simulation variables requires a change in focus in terms of the organizational structure of the ontology. Up to this point, the ontology reflects an enumerative, hierarchical classification process, with the structure provided by type–subtype relationships between concepts, and the specification of instances at the level of the lowest-order concepts. The focus now changes to reflect a faceted classification structure, whereby domain knowledge is organized not on the type–subtype relationship between concepts, but in mutually-independent groupings based on shared attributes of the instances of the lowest-order concepts.[41]

Arguably, the range of relevant attributes is much narrower than is typically the case in faceted classification schemes, but this characteristic is not a design deficiency but rather a function of the purpose of the ontology, which is to provide a knowledge base and reasoning capabilities for implementing the results of a dynamic simulation model of biotechnological innovation. The dynamic simulation provides a set of relevant variables and inferential relationships that when manipulated in a particular manner is likely to achieve a desired outcome. Achieving the desired state of affairs is a matter of choosing appropriate modalities, i.e., those modalities that achieve the desired increase or decrease in the incidence of a variable as determined by the results of the dynamic simulation. Accordingly, the only attributes of interest (at this time) are a modality's effect on a variable and the type of modality, i.e., a type of legal rule, practice or institution.

The conceptual organization is reintroduced after the associations between variables and modalities are identified in order to analyse the nature of a particular

[40] See Van de Ven and Delbecq (1972).

[41] See, for example, Uddin and Janecek (2006). Note that the mutually independent categories may subsequently be organized in the form of a type–subtype hierarchy.

modality and examine its function in relation to its originating knowledge domain. Thus the ontology is organized both as an enumerated hierarchy of type–subtype relationships and as a faceted classification structure. Although the range of attributes is rather limited, the ontology's expressive functionality is nonetheless robust. The reasoning process described in the next section makes use of both specified and inferred relationships between modalities to provide an optimised combination of legal rules, practices and institutions for implementing a desired policy objective.

3.3.1.2 Identifying modalities suitable for implementation Relations between a given variable and relevant modalities are defined by whether an association expresses a causal relationship in either a positive or negative direction.[42] In a positive relationship, adoption of a particular modality increases the incidence of the variable while the inverse is true for a negative relationship. If the results call for an increase in the incidence of a particular variable, then the modalities which would be desirable to implement are those in which a positive relationship exists between the modality and the variable. On the other hand, if the results of the dynamic simulation call for a decrease in the incidence of a variable then the modalities which would be desirable to implement are those in which a negative relationship exists between the modality and the variable. Not every modality needs to be linked to every variable, although each variable must be linked to at least one (and preferably more) modalities.

Taking the example of the variable of government use (as discussed above), assume that the results of the dynamic simulation call for an increase in the incidence of government use in order to achieve the policy objective of maximizing the level of access to innovation. A number of modalities are the subject of a positive causal relationship, such as providing for compulsory licensing of publicly-funded patented inventions (a legal rule modality) and placing conditions on research grants to publicly-funded universities that engage in technology transfer to the private sector (a practices modality).

The expressive capability of the completed ontology is not limited to the specification of positive or negative relationships between variables and modalities, but provides for inferred relationships as well. For example, when a relationship (either positive or negative) is specified between a variable and a particular modality, an inverse relationship between the modality and the variable is inferred. Accordingly, one can analyse the knowledge communicated by the ontology either by examining all modalities appropriately related to a particular variable, which would represent the range of possibilities in the form of legal rules, practices and institutions for a given variable, or all variables related to a particular modality,

[42] The proposed methodology for specifying all of the relationships between presumptively positive and negative modalities in relation to each variable is once again that of the nominal group method, the benefit being that this method capitalizes on the transdisciplinary nature of the knowledge acquired by the IPMG. At the current stage of development, however, the ontology represents a proof of concept and thus the relationships have not been subject to the nominal group method.

which would represent all the dynamic simulation variables affected by the implementation of a particular modality.

Being able to locate all variables affected by a particular modality is necessary for determining which modality to adopt among the various possibilities associated with a particular variable. Many decision factors will be external to the ontology, as they will relate to the political, legal, social and economic conditions in existence at the time the analysis takes place. The choice of a particular modality, however, is also determined by constraints inferred by the ontology. An appropriately negative or positive relationship may exist between a variable and a modality, but this does not mean that an appropriately positive or negative relationship exists between that modality and each of the variables used in the dynamic simulation.

For example, a number of modalities are inappropriately negatively related to government use in that adoption of the modality would decrease rather than increase the incidence of this variable. Such modalities include not permitting states to assert immunity in relation to patent infringement suits (a legal rules modality) as well as restricting public funding of research and development (a practices modality) and focusing research efforts in the private as opposed to public sector (an institutions modality). These modalities may be appropriately related, either positively or negatively, to any number of relevant variables, but since they are inappropriately related to the variable of government use, adopting any one of these modalities will work against the desired policy objective.

The necessity of selecting appropriately related modalities across the entire range of variables used in the dynamic simulation demonstrates the value of being able to rely on the expressive capacity of the ontology design for drawing inferences. When a number of modalities are specified as being appropriately related, either positively or negatively, to a variable, a rule can be implemented to infer a complementary relationship between the modalities themselves. Following this inference, all of these modalities are now related in the sense that they belong to the set of modalities appropriately associated with a particular variable and thus are suitable for implementation. Selecting the appropriately related modalities for each variable is now a matter of identifying the modalities that are members of the greatest number of sets of complementary modalities, and are therefore compatible across the widest range of variables.

This analysis could be further refined. For example, the results of the dynamic simulation could be used to specify rules for inferring which modalities are appropriately related, either positively or negatively, to one or more variables and are not inappropriately related to any variable. This would automatically exclude from consideration modalities that are not compatible across the entire range of relevant variables. When implemented for each variable, the inferences would produce a set of truly complementary modalities in that any given modality would be appropriately related to any given variable, in which case the effect of the modality is reinforced across the range of variables, or would not be related inappropriately with any variable and thus would not work against the desired policy objective.

A set of truly complementary modalities, however, is unlikely. Each variable is linked to one or more modalities, and the greater the number of linked modalities, the greater the possibility that one or more of these modalities is the subject of an inappropriate relationship with one or more variables. Accordingly, one may be unable to choose any modality facilitating a given variable that is not also contradictory with at least one other variable. For example, assume that each of the modalities that increase the incidence of the variable "government use" also decrease the incidence of another relevant variable, "private sector commercialization". Assume as well that the results of the dynamic simulation indicate that that in order to achieve the desired policy objective, the incidence of both of these variables must increase.

Choosing between modalities in this case is a matter of selecting the most facilitative (and least non-facilitative) modalities. Relying on inferences drawn regarding the relationship between the modalities themselves across the range of variables, one could identify modalities with the greatest number of complementary modalities (membership in the same set in relation to a variable) and the least number of what might be termed "non-complementary" modalities (a member of the set of modalities which are inappropriately related to a variable). In circumstances requiring the implementation of a non-complementary modality, care must be taken to calibrate the modality to achieve the minimal degree of implementation necessary to facilitate an associated variable, thus minimizing the degree to which the modality inappropriately affects the incidence of other variables.

3.3.2 Analysing modalities as instances of specified concepts

Implementation of the modalities cannot take place in the abstract. Accordingly, once the relationships between the variables and the modalities have been established, the enumerated hierarchical structure of type–subtype relationships between domain concepts is reintroduced. This permits analysis of the selected modalities as instances of a concept in relation to other concepts. In effect, the selection of modalities forms a pattern of instances located throughout the visual layout of the ontology.

The hierarchical structure of the ontology, with its abstract and simplified representation of significant concepts derived from the knowledge domains of legal rules, practices and institutions, provides the conceptual framework for a contextual analysis of innovation governance. Examination of the modalities in relation to their ontological classification achieves two objectives. First, analysing the modalities as instances of concepts across the three knowledge domains provides the necessary transdisciplinary perspective by demonstrating the manner in which policy objectives can be achieved using a combination of legal rules, practices and institutions. Viewing the modalities as instances of concepts in an ontological structure operates as a proof of concept to policy-makers and researchers that policy objectives involving innovation assets can be achieved through other than legal means.

This information is useful if a policy-maker or researcher is operating under decision constraints, such as mandatory universal minimum standards of intellectual property protection set out in the TRIPS Agreement.[43] The degree of flexibility

[43] *TRIPS*, *supra* note 34.

available to states in implementing the minimum universal standards of the TRIPS Agreement is a matter of debate. Developed states typically prefer more restrictive interpretations of its terms, while developing states support less restrictive interpretations. Additional flexibility can also be achieved, however, by considering alternative methods for achieving policy objectives. For example, a state cannot refuse to patent an invention on the basis of subject matter alone; so long as an invention is new, useful and non-obviousness and does not fall within an enumerated subject-matter exception, a patent must be granted. A state can, however, regulate the market conditions under which the patented invention is commercialized,[44] just as non-governmental actors such as NGOs can influence the market behaviour of corporations.[45]

Second, understanding the modalities as instances of concepts that are themselves related in a particular manner assists in the necessary deliberations involved in determining the precise manner in which to implement a modality. This is particularly important when determining the implementation of non-complementary modalities as described above. Implementation of a modality requires more than transforming a prescription such as “place conditions on research grants” into an action whereby state funding agencies are directed to change their practices accordingly. The IPMG methodology is designed to assist, not replace, human decision-making. Implementation should be the product of the kind of thoughtful discussion and analysis that takes place when a modality is analysed as part of a specification of a conceptualization.

Thus the value added by analysing the modality as an instance of a concept is that a policy-maker or researcher can make use of the expressive nature of the ontology to understand the nature of the modality in relation to the concepts representing relevant knowledge of each of the three domains. Informed discussion and analysis is facilitated when participants in policy development or research can readily understand the most significant concepts of a complex body of knowledge communicated in accessible terms, can work from a common conceptualization that prevents misunderstandings and can discuss issues using a common vocabulary.

Issues are also more readily defined when working from a common conceptualization that focuses analysis and discussion by framing the terms appropriately. Specified concepts, along with the hierarchical type–subtype relation used to organize the concepts, represent and communicate what is known (from a given perspective) about a particular topic in relation to a modality. Taking once again the example of placing conditions on research grants, analysis of the modality is facilitated by the knowledge communicated by the structure of the specified concepts. Conditional research grants are an example of a practice involving federally-funded innovations, which is one of four types of practices that affect technology transfer (which in turn is one of the many types of practices affecting the use and exchange of innovation assets):

[44] The manner in which a patented invention is sold, for example, may violate competition law.

[45] For example, after global outcry from NGOs, Monsanto announced that it would abandon its controversial plan to market “terminator” seeds (Lambrecht 1999); “Genetic Engineering: Monsanto to Abandon Terminator Seeds Plan” *Europe Agri* (8 October 1999) (QL).

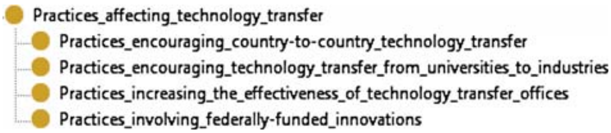

The examination of a modality within the context of an established, credible classification structure permits the ontology to be used in a formative manner. The knowledge communicated is used to generate new knowledge; analysis proceeds from an organized conceptual framework that serves as a foundation for the development and discussion of implementation issues in relation to a particular modality. The formidable task of analysing implementation issues for a large number of modalities is also simplified. The classification structure provides for a narrow focus when addressing issues relating to a specific modality and for a broader focus when viewing the functioning of the modalities in the aggregate across the three domains of legal rules, practices and institutions.

4 Conclusion

Intellectual property law is often viewed as a mechanism for achieving policy objectives involving innovation assets. The fact, however, that innovation assets tend to be subject to intellectual property protection does not lead to the conclusion that related policy objectives are necessarily a matter of intellectual property law reform. Legal rules are significant, but must be viewed in context as but one factor among many. When innovation governance is analysed from transdisciplinary perspective, relying on the cooperative efforts of researchers from fields other than law, the necessary context is provided; innovation governance is characterized not simply as the product of legal rules, but as a function of the interaction of legal rules, practices and institutions relevant to the creation and exchange of innovation assets.

Accordingly, when policy-makers and researchers seek to identify the conditions under which the creation, use and exchange of innovation assets flourishes, care should be taken to focus on the interaction of these various factors. The modest contribution of this article is to describe the development of an ontological representation of innovation governance as providing the necessary conceptual framework. The term “ontology” is used here not in its philosophical sense as a study of being and the kinds of things that exist, but as a term used in the field of computer science to refer to the manner in which knowledge of a particular domain is represented in the form of concepts and relations between concepts.

The ontological form provides an ideal knowledge structure for representing a transdisciplinary perspective of innovation governance. The final product, with its abstract and simplified representation of significant concepts derived from the knowledge domains of legal rules, practices and institutions, as well as relations between dynamic simulation variables and modalities, provides policy-makers and researchers with a range of alternative arrangements for achieving public policy objectives as well as a context in which to understand their operation. Achieving a given objective, such as maximizing the level of access to innovation, is a matter of selecting appropriate modalities from a number of possible legal rules, practices and institutions.

The ontology is not a form of automated decision-making but is designed to provide policy-makers and researchers with an improved understanding of innovation governance from a transdisciplinary perspective. Informed discussion and analysis is facilitated when participants in policy development or research can readily understand the most significant concepts in a complex body of knowledge communicated in accessible terms, can work from a common conceptualization that prevents misunderstandings and can discuss issues using a common vocabulary. The expressive nature of knowledge represented in ontological form thus assists policy-makers and researchers in designing effective and creative strategies to achieve desired objectives in relation to innovation assets.

Acknowledgements This research project benefited from funding provided by the Social Sciences and Humanities Research Council of Canada and the Canadian Institutes of Health Research. I would also like to thank the many students who have worked on this project over the past 3 years, including Andrew Brighten, LL.B.-B.C.L. 2009 (McGill).

References

Agreement on Trade-Related Aspects of Intellectual Property Rights, 15 April 1994, 1869 U.N.T.S. 299, 33 I.L.M. 1197 (1994)

Asuncion Gomez-Perez (2001) Evaluation of ontologies. Int J Intell Syst 16:391, 393–409

Bayh-Dole University and Small Business Patent Procedures Act of 1980, Pub. L. No. 96-517, 94 Stat. 3015

Berners-Lee T, Hendler J, Lassila O (2001) The semantic web. Sci Am 284:34–43

Breuker J, Valente A, Winkels R (1997) Legal ontologies: a functional view. In: Visser P, Winkels R (eds) Proceedings of the first international workshop on legal ontologies, Melbourne, 1997. University of Melbourne, Melbourne, p 23

Burton-Jones A, Storey VC, Sugumaran V, Ahluwalia P (2005) A semiotic metrics suite for assessing the quality of ontologies. Data Knowl Eng 55(1):84–102

Chi Y-L, Hsu T-Y, Yang W-P (2006) Ontological techniques for reuse and sharing knowledge in digital museums. Electronic Libr 24(2):147–159

Corcho O et al (2005) Building legal ontologies with METHONTOLOGY and WebODE. In: Benjamins VR et al (eds) Law and the semantic web. Springer, New York

Description of IPMG research approach. Online: Centre for intellectual property policy http://www.cipp.mcgill.ca/data/publications/00000047.pdf

Ding Y, Foo S (2002) Ontology research and development. Part 1 – a review of ontology generation. J Inf Sci 28(2):123–136

Fernandez M (1999) Overview of methodologies for building ontologies. In: Benjamins VR, Chandrasekan B, Gomez-Perez A, Guarino N, Uschold M (eds) Proceedings of the IJCAI-99 workshop on ontologies and problem solving methods, August 1999

Gangemi A, Catenacci C, Ciaramita M, Lehmann J (2005) A theoretical framework for ontology evaluation and validation. In: Semantic web applications and perspectives (SWAP) – 2nd Italian semantic web workshop, Trento, Italy, 2005

Genetic engineering: Monsanto to Abandon terminator seeds plan Europe Agri (8 October 1999) (QL)

Gold ER et al (2004) The unexamined assumptions of intellectual property. 18 P.A.Q. (4):299–344

Gomez-Perez A (2002) Evaluating ontological decisions with ontoclean. Commun ACM 45(2):61–65

Gruber TR (1995) Toward principles for the design of ontologies used for knowledge sharing. Int J Hum Comput Stud 43(5–6):907–928

Gruber T http://www.ksl.stanford.edu/kst/what-is-an-ontology.html
Guarino N (1995) Formal ontology, conceptual analysis and knowledge representation. Int J Hum Comput Stud 43(5–6):625–640
Hage J, Verheij B (1999) The law as a dynamic interconnected system of states of affairs: a legal top ontology. Int J Hum Comput Stud 51(6):1043–1077
Hjorland B (2000) Book Review, Rita Marcella and Arthur Maltby, The future of classification. Gower Publishing, Hampshire, England
Hjorland B (2002) Book review: the future of classification. J Am Soc Inf Sci Technol 53(1):57–62
Hjorland B, Pederson KN (2005) A substantive theory of classification for information retrieval. J Doc 61(5):582–597
Jacob EK (2003) Ontologies and the semantic web. Bull Am Soc Inf Sci Technol 29(4):19–22
Jones DM, Bench-Capon TJM, Visser PRS (1998) Methodologies for ontology development. In: Proc. IT&KNOWS conference of the 15th IFIP world computer congress, Chapman-Hall
Lambrecht B (1999) Monsanto will wait for studies of disputed new gene technology. St. Louis Post-Dispatch (23 April 1999) A4
Lame G (2004) Using NLP techniques to identify legal ontology components: concepts and relations, 12 A.I. & L. (4):379–396
Mommers L (2004) Searching for arguments: applying a knowledge-based ontology of the legal domain, 13 I. & Comm. T. L. (1):75–91
Soergel D (1999) The rise of ontologies or the reinvention of classification. J Am Soc Inf Sci 50(12):1119–1120
Sterman JD (2000) Business dynamics. McGraw-Hill, Boston
Su X, Ilebrekke L (2002) A comparative study of ontology languages and tools. In: Pidduck AB et al (eds) Advanced information systems engineering: proceedings of the 14th international conference, CAiSE 2002, Toronto, Canada, 27–31 May 2002, Springer, New York, p 761
Supekar K (2005) A peer-review approach for ontology evaluation. In: 8th Int. Protégé Conf., Madrid, Spain, July 2005, pp 77–79
Théberge v. Galerie d'Art du Petit Champlain, 2002 SCC 34, [2002] 2 S.C.R. 336
U.S. Const. art. I, § 8, cl. 8
Uddin MN, Janecek P (2006) Faceted classification in web information architecture. Electronic Libr 25(2):219–233
Uschold M, Grüninger M (1996) Ontologies: principles, methods and applications. Knowl Eng Rev 11(2):93–136
Valente A (2005) Types and roles of legal ontologies. In: Benjamins VR et al (eds) Law and the semantic web. Springer, New York
Valente A, Breuker J (1996) Towards principled core ontologies. In: Proceedings of the tenth knowledge acquisition for knowledge-based systems workshop, Banff, 1996. Online: University of Calgary http://ksi.cpsc.ucalgary.ca/KAW/KAW96/KAW96Proc.html
Van de Ven AH, Delbecq AL (1972) The nominal group as a research instrument for exploratory health studies. Am J Public Health 62(3):337–342
Van Kralingen RW et al (1999) A principled approach to developing legal knowledge systems. Int J Hum Comput Stud 51(6):1127–1154
Vickery BC (1997) Ontologies. J Inf Sci 23(4):277–286
Visser PRS, Bench-Capon TJM (1998) A comparison of four ontologies for the design of legal knowledge systems, 6 A.I. & L. (1):27–57
Web Ontology Language (OWL). Online: W3C http://www.w3.org/2004/OWL/