



Food for Thought ...

A Toxicology Ontology Roadmap

Barry Hardy¹, Gordana Apic², Philip Carthew³, Dominic Clark⁴, David Cook⁵, Ian Dix^{5,6}, Sylvia Escher⁷, Janna Hastings⁴, David J. Heard⁸, Nina Jeliazkova⁹, Philip Judson¹⁰, Sherri Matis-Mitchell⁵, Dragana Mitic², Glenn Myatt¹¹, Imran Shah¹², Ola Spjuth¹³, Olga Tcheremenskaia¹⁴, Luca Toldo¹⁵, David Watson¹⁰, Andrew White³, and Chihae Yang¹⁶

¹Douglas Connect and OpenTox, Zeiningen, Switzerland; ²Cambridge Cell Networks, Cambridge, UK; ³Unilever, Sharnbrook, Beds, UK; ⁴EMBL-EBI, European Bioinformatics Institute, Cambridgeshire, UK; ⁵AstraZeneca, Macclesfield, Cheshire, UK; ⁶Pistoia Alliance; ⁷Fraunhofer Institute for Toxicology and Experimental Medicine, Hannover, Germany; ⁸Novartis, Basle, Switzerland; ⁹Ideaconsult Ltd, Sofia, Bulgaria; ¹⁰Lhasa Ltd., Leeds, UK; ¹¹Leadscope, Columbus, OH, USA; ¹²US EPA, Research Triangle Park, NC, USA; ¹³University of Uppsala, Uppsala, Sweden; ¹⁴Istituto Superiore di Sanità, Rome, Italy; ¹⁵Merck KGaA, Darmstadt, Germany; ¹⁶Altamira, Columbus, OH, USA

Summary

Foreign substances can have a dramatic and unpredictable adverse effect on human health. In the development of new therapeutic agents, it is essential that the potential adverse effects of all candidates be identified as early as possible. The field of predictive toxicology strives to profile the potential for adverse effects of novel chemical substances before they occur, both with traditional in vivo experimental approaches and increasingly through the development of in vitro and computational methods which can supplement and reduce the need for animal testing. To be maximally effective, the field needs access to the largest possible knowledge base of previous toxicology findings, and such results need to be made available in such a fashion so as to be interoperable, comparable, and compatible with standard toolkits. This necessitates the development of open, public, computable, and standardized toxicology vocabularies and ontologies so as to support the applications required by in silico, in vitro, and in vivo toxicology methods and related analysis and reporting activities. Such ontology development will support data management, model building, integrated analysis, validation and reporting, including regulatory reporting and alternative testing submission requirements as required by guidelines such as the REACH legislation, leading to new scientific advances in a mechanistically-based predictive toxicology.

Numerous existing ontology and standards initiatives can contribute to the creation of a toxicology ontology supporting the needs of predictive toxicology and risk assessment. Additionally, new ontologies are needed to satisfy practical use cases and scenarios where gaps currently exist. Developing and integrating these resources will require a well-coordinated and sustained effort across numerous stakeholders engaged in a public-private partnership. In this communication, we set out a roadmap for the development of an integrated toxicology ontology, harnessing existing resources where applicable. We describe the stakeholders' requirements analysis from the academic and industry perspectives, timelines, and expected benefits of this initiative, with a view to engagement with the wider community.

Keywords: toxicology, ontology, roadmap, interoperability, framework, risk assessment

1 Introduction

Foreign substances can have a dramatic and unpredictable adverse effect on human health. In the development of new therapeutic agents, it is essential that the potential adverse effects of all candidates be identified as early as possible. The field of predictive toxicology strives to profile the potential for adverse effects of novel chemical substances before they occur, both with traditional *in vivo* experimental approaches and increasingly through the development of *in vitro* and computational methods which supplement

and reduce the need for animal testing. To be maximally effective, the field needs access to the largest possible knowledge base of previous toxicology findings, and such results need to be made available in such a fashion so as to be interoperable, comparable, and compatible with standard toolkits. This necessitates the development of open, public, computable, and standardized toxicology vocabularies and ontologies that can be implemented in support of knowledge harnessing in applications and methods as well as related analysis and reporting activities. Such ontology development will support data management, model building, integrated



analysis, validation and reporting, including regulatory reporting and alternative testing submission requirements as required by guidelines such as the REACH legislation¹, leading to new scientific advances in a mechanistically-based predictive toxicology.

Through increasing access to relevant information and clarifying the meaning of terminology across different sub-disciplines, progress on public ontology development is likely to have scientific impact on the conceptual foundations and outcomes of research programs. It will also lead to cost savings through the avoidance of duplication of effort, and to increased quality of results and relevance of hypotheses through improving interoperability and easing the difficulties of data integration. To address this opportunity, a workshop was recently held at the European Bioinformatics Institute in Hinxton, UK, which brought together key academic and industrial stakeholders in predictive toxicology. In this communication we report on the outcome of that workshop, with particular focus on the objectives, stakeholders and their requirements, and present a roadmap for the development of a unified toxicology ontology. A parallel communication (Hardy et al., 2012) serves as a review and perspective describing recent advances in ontologies within the toxicology field and related projects which can be harnessed for re-use, such as the Gene Ontology (Ashburner et al., 2000).

In the remainder of this introduction, we give more detail about the purpose and objectives of this communication and the proposed toxicology ontology development. Section 2 describes the results of our analysis of different stakeholder perspectives. A detailed list of requirements is laid out in Section 3, and finally we specify a roadmap for the ontology development in Section 4 before concluding.

1.1 Roadmap objectives

The motivation and objectives of this perspective and roadmap are to:

- Raise awareness of existing ontology development activities. Specifically, we integrate profiles of current project activities (scope, objectives, indicative time-lines and use cases) in order to map existing activities to scope statements to develop a roadmap for toxicology ontology development;
- Define a Toxicology Ontology Roadmap including product, scope, and phases for implementation;
- Identify the key existing ontology contributions that differing organizations, projects, and existing ontology programs could make to such a roadmap;
- Identify needs of different stakeholders and establish the key ontology requirements of stakeholders in this field;
- Identify opportunities for collaboration and agreed area(s) of focus for exchange of best practices, developing an action plan for increased coordination;
- Communicate to a broader community, highlighting value and needs, engaging with all stakeholders including policy makers, industry, research institutions, regulators, and solution providers.

Ontology purpose

The purpose of a toxicology ontology is to show practical utility in supporting applications and satisfying user needs in use cases and analysis scenarios. Several current high level goals are of relevance:

- Universal access to high quality experimental data is a major pre-requisite for the successful implementation of the main principles of the Three Rs Declaration of Bologna, adopted by the 3rd World Congress on Alternatives and Animal Use in the Life Sciences (Bologna, Italy, August 31st 1999) – namely Reduction, Refinement and Replacement Alternatives (3Rs);
- The improved storage, exchange and use of information from experiments already carried out is essential to avoid unnecessary repetition of experiments and to support increased access to data and associated metadata including experimental protocols and scientific concepts;
- The integrated use of physical, biological and chemical techniques and data support more accurate predictions and ultimately facilitate the development of more robust human-oriented toxicology models;
- The integration of interdisciplinary and translational concepts facilitates an enriched toxicology science that protects human and environmental health;
- Standardized ontologies enhance data retrieval, easing the satisfaction of regulatory demands in risk assessment as required by legislation such as REACH and the Cosmetics Directive, both of which require a reliable integration of weight of evidence of different types;
- Computational research and prediction in toxicology will harness an increase in the development and application of complex computer-based models and systems in toxicology, including modeling of structure-activity relationships, molecular simulation, computer graphics, and the linked modeling of biochemical, pharmacological, physiological, toxicological, and behavioral processes;
- Since it is likely that, in many circumstances, an animal test cannot be replaced by a single replacement alternative method, the development, evaluation and optimization of stepwise testing strategies and integrated testing schemes should be encouraged.

The OpenTox² (Hardy et al., 2010) data facilities (made publically accessible through a semantic web services framework) provide a solid basis for addressing the above mentioned goals in a more efficient, technically sound and integrated way compared to the existing uncoordinated practices and fragmented resources. Unfortunately, even today, more than half a century after Russell and Burch's original publication and more than 10 years after the adoption of the Three Rs Declaration of Bologna, the "state-of-the-art" is characterized by highly fragmented and unconnected life sciences data related to toxicology (both from a physical and ontological perspective). The capability to more easily integrate and use data, models and concepts within a public ontological framework would create significant value in use cases for all sectors.

¹ REACH Guidance on Information Requirements and Chemical Safety Assessment, http://guidance.echa.europa.eu/docs/guidance_document/information_requirements_en.htm

² OpenTox, <http://www.opentox.org/>



2 Perspectives

2.1 The user perspective on data analysis

Perhaps the most important near term high impact use for a toxicology ontology is enabling integrated data analysis of toxicology-related data through the creation of a capacity to capture, access, reuse, and analyze data in a variety of user contexts not possible or easy to achieve today.

Toxicology sits at the intersection of many different branches of science within discovery and development. Since it uses data from different branches, the impact of standards (or lack thereof) is very acute for this discipline. The key needs are:

- Understanding what data and what standards for data including vocabularies are needed;
- Access to data including retrieval and interoperability as well as an understanding of access permissions and “who needs to know”, possibly at the level of the individual data, and including part of the metadata;
- Storing the data so it is easy to retrieve, but more importantly, linking to the correct data and metadata needed to answer the user’s questions.

From a strategic point of view we must understand the questions that users need to have answered, and this will drive the evolution of data models and ontologies and their details. To begin, we also need to have a survey of what has already been done, as there are likely to be challenges associated with incorporating legacy data. As we will need to handle and translate from many different formats, data exchange formats and vocabularies will be important challenges to solve. We also need to set expectations on the historical data: it is not just the standard problem of controlled vocabularies, harmonized units of measure, or observer variability. Advances in technology and changes in medical science have occurred; thus, data may be missing in historical datasets. We would not expect to have Single Nucleotide Polymorphism (SNP) data for studies done years ago, nor would we realistically be able to go back and get consent even if we did have the samples and the resources required to generate the data to fill in the blanks. The types of tests and measurements have changed as well as the accuracy and the ways the values are reported; in some cases the tests have been renamed and scientific advances have given us additional markers that were not known 10 years ago. It is also important to understand other systematic reasons for missing data, like noncompliance or lack of signal because the patient recovered. So, we need to set the correct expectations for how the available historical data can be used. Optimally, prospective studies offer the most control over the type of data generated and made available but there is also great value in using what has already been generated.

There are real challenges in integrating data and understanding how the data will be used can drive some potential solutions. Scientific and business questions will dictate the types of data to be linked as well as the permissions on accessing and sharing the data. We also need to understand what data is needed and the relationships between the data, i.e., is one data type merely a direct transformation of another, thus adding little information?

Without good negative data and balanced large and high quality datasets, we cannot develop good models.

The data needs to be accessible to those that will retrieve and use it. The access permissions could be part of the metadata to expedite retrieval and additional metadata, such as experimental protocols and inferences, needs to be available as well.

Looking ahead to uses of data in an integrated data analysis supported by ontology, we would need to focus on supporting methodology that is robust to many of these issues, e.g., models that tolerate missing data. Critical validation should be an inherent part of any model building activity; the more transparent the model, the closer we are to being able to carry out a proper scientific evaluation. The output of the models must be usable; simply providing a score or answer with a lack of context is insufficient. Ontology enables us to place diverse data in a framework with biological and functional meaning that is more suited to the discovery and application of safety biomarkers.

2.2 Market needs

The integration of information and data on both hazard and exposure to deliver a transparent, effective risk assessment for environmental and human health is fundamental for toxicological safety. The approaches and capabilities to deliver this from a technical, scientific, and computational capacity have advanced rapidly in the last 10 years. This has combined with a shift in the needs and drivers coming from industry and consumers and also the regulatory environment to push for a transition in our toxicology science towards greater mechanistic understanding, aided by computer-based approaches.

Across industry, the need to bring innovative, beneficial products to the market rapidly continues to grow both for competitive and sustainability reasons, while the resources necessary to deliver this have to remain affordable. In addition, each industry sector has its own needs, for example, within the pharmaceutical industry there is a continued requirement to improve the quality of candidate compounds and lower attrition rates. This requires effective assessment of efficacy balanced with improved early stage toxicity prediction. At the same time, a strengthening of pharmacovigilance approaches is highlighted through optimizing available epidemiology data sources and maximizing their use. Across other industries, such as the consumer goods and chemical industries, regulations such as REACH in the EU, TSCA³ reauthorization in the USA, and EU Directives such as the Cosmetics Directive⁴ have, in combination with 1) addressing the ethical concerns of consumers around animal welfare and 2) the debate on the relevance of high dose animal testing to low dose human exposure, shifted the emphasis to identifying novel alternative approaches to animal testing.

The continued advances made in molecular, cellular, and computational biology for both basic and clinical research have provided a renewed momentum for toxicological research generating a range of new and improving tools. A stimulus for these changes was highlighted in the report on Toxicity Testing in the 21st Century (NRC, 2007). In parallel, and as a result, a range of initiatives and database developments from the USA and Europe have com-

³ Toxic Substance Control Act <http://www.epa.gov/agriculture/lasca.html>

⁴ Cosmetics Directive, http://en.wikipedia.org/wiki/Cosmetics_Directive



menced. US initiatives include The Human Toxicology program⁵, RISK21⁶, and the US EPA ToxCast⁷, ExpoCast⁸ and database developments. European efforts include joint public-private partnerships between the European Commission and industry bodies such as EFPIA⁹ and COLIPA¹⁰. These include components of the IMI¹¹ and SEURAT-1¹² programs, respectively, alongside the ongoing efforts of the EPAA¹³, AXLR8¹⁴ and associated EU framework projects, and the CEFIC toolbox development¹⁵. This momentum has provided an opportunity for a constructive dialog to emerge involving key stakeholders covering NGOs, industry, regulatory bodies and academia to build a broad scientific base to further the development of credible approaches for risk assessment. The underlying principle within each of these approaches requires the grouping of diverse datasets with *in silico* structure relationships and mechanistic relevance to relevant human dosimetry. In order for this to be satisfactorily achieved and transparent to all stakeholders from within the industry, academia, and regulatory bodies, there is a requirement for a clearly defined ontology that enables the consistent capturing of relationships between a chemical structure, or an *in vitro* assay, or a text-based finding and a physiological endpoint that signifies both the type and also the extent of the adverse response.

Industry drivers

“First: do not harm!” is an ethical guideline for any social or environmentally-sensitive activity. Industry, defined as a set of entities delivering products and services to the society, therefore has to ensure that hazard risks be identified, specified and predicted in the best possible way. For these applications, a toxicology ontology is a core component for knowledge capturing and formalization, and an enabler of mechanistic understanding of toxicological behaviors. Beyond compliance with legal, regulatory and ethical requirements, toxicological profiles of existing products define opportunities for new products and thus are key for definition of new business opportunities. Industry drivers need to be able to maximize available public data and to integrate it with internal studies and place it reliably in context. Therefore, not only is the highly-regulated pharmaceutical industry a major driver for toxicology ontology, but many other industries, such as agrochemicals, veterinary, chemical, food, consumer products and consumables, share similar drivers. Software and content vendors enable the industrial use of toxicological knowledge and can contribute greatly to the needed standardization effort.

Technical components impacting industry R&D and risk assessment

For a toxicology ontology to be effective it must meet the requirements of its stakeholders. An ontology should act as an enabler to

optimize the use of the available data resources that are currently being developed and to reinforce the link for human relevance. Two further factors remain clear issues now and for the future:

- 1) The amounts of biological data now available to the industrial toxicologist continue to grow at a significant rate. Such data, both from proprietary and public sources, includes literature and datasets that are vast, diverse and of variable quality. The resources needed to maintain and internalize these within a company environment are limited, making the use of distributed data resources with strong integration across both proprietary and public sources a more realistic scenario. The use of multiple proprietary standards can be a barrier and a cost to industry. However in areas of non-competitive advantage, as supported for example by the Pistoia Alliance¹⁶ and OpenTox, the engagement of industry to develop open source standards in combination with tool and service providers can be seen.
- 2) The levels of evidence required to support a risk assessment are different across the different industries, e.g., regulation of chemicals under REACH, pesticides, biocides, pharmaceuticals or cosmetics. It also will vary depending on the requirement, e.g., toxicology research versus regulatory risk assessment. It is therefore essential that the computational toxicologist is able to stratify the data, utilizing the underlying information about the experimental conditions to extract only the data that are relevant to the model required. For example, a hazard ranking for prioritization in early screening might be obtained from a combination of (Quantitative) Structure-Activity Relationship ((Q)SAR) models available in the literature and bioinformatics. Here the extraction of the maximum information to allow clustering of compounds with biological effects may be sufficient to aid the screening. However, as the requirement changes towards a risk assessment, the context of the information incorporating Absorption, Distribution, Metabolism, Excretion, and Toxicity (ADMET) profiles, mechanistic information, exposure, and dosimetry become more critical.

In each case, for such systems to be effective, a structured, formalized toxicology ontology becomes a necessity for both combining different levels of evidence and making this transparent within a risk assessment. With the growing capability for predictive simulations, the current focus for toxicology research is on understanding and modeling human “toxicity pathways” or pathways leading to an adverse endpoint. This highlights the need to ensure that links to clinical data sources and population data are made available. In each case ontologies and standards have been proposed or implemented and their use or incorporation where suitable for toxicology should be encouraged to

⁵ The Human Toxicology program, <http://toxicology.grad.uiowa.edu/>

⁶ RISK21, <http://www.hesiglobal.org/>

⁷ US EPA ToxCast, <http://www.epa.gov/ncct/toxcast/>

⁸ US EPA ExpoCast, <http://www.epa.gov/ncct/expocast/>

⁹ EFPIA - European Federation of Pharmaceutical Industries and Associations, <http://www.efpia.org/>

¹⁰ COLIPA - Cosmetics Europe, <http://www.colipa.eu/>

¹¹ IMI – Innovate Medicines Initiative, <http://www.imi.europa.eu/>

¹² SEURAT-1, <http://www.seurat-1.eu/>

¹³ EPAA - European Partnership for Alternative Approaches to Animal Testing, http://ec.europa.eu/enterprise/epaa/index_en.htm

¹⁴ AXLR8, <http://www.axlr8.eu/>

¹⁵ CEFIC - European Chemical Industry Council toolbox, <http://www.cefic-iri.org/iri-toolbox>

¹⁶ Pistoia Alliance, <http://www.pistoiaalliance.org/>

ensure cross fertilization between toxicology and biomedical research.

Application to safety and regulatory toxicology

One of the major developments in chemical safety evaluation over the last few years has been the White Paper for the REACH process and implementation of this policy in Europe, to ensure the safety of human exposure to chemicals. This process has made chemical safety evaluation for human exposure more in line with the process used for ecotoxicology evaluation of chemicals, and the process used to ensure the safety of chemicals and ingredients in food.

REACH still uses the paradigm of determining the threshold for toxicological adverse effects, the no adverse effect level (NOAEL) in animal toxicity testing, and then comparing this to the human exposure scenarios for differing uses of chemicals, both in the occupational and consumer context.

REACH has introduced the new concept of the derived no effect level, or DNEL, for thresholded effects and the concept of derived minimum effect levels (DMEL) for non-thresholded effects (such as for genotoxic carcinogens). The DNEL is derived from a NOAEL, identified from an animal study, and is effectively a “human NOAEL”, calculated by incorporating assessment (uncertainty and extrapolation) factors into its derivation, so that it becomes the level of no appreciable concern for human exposure to a chemical.

Because the implementation of REACH could have resulted in significant amounts of additional animal testing being triggered under the tiered tonnage approach used in safety evaluation where, as the tonnage level increases additional animal testing is required, special attention was given to the provision for using integrated testing strategies. Options such as reducing or waiving the testing requirements can be considered on the basis of, for instance, low exposure, or an understanding of the mechanism of toxicity in animals, and whether it is relevant to human exposure.

This has allowed the potential application of a more detailed understanding of pathways critical to the mechanisms of action to be used in safety evaluation to minimize animal testing. Obviously such an approach relies on the development of an understanding

at the genomic, metabolomic and proteomic level of responses to toxic insults. The integration of this information via an approach that uses ontologies can help to secure an understanding of the toxicological anchorage of effects seen in animal studies and *in vitro* to the likely effects in man.

The most effective use of such an approach will rely heavily on the integration of the existing ontologies for descriptive histopathology, defining pathological effects and pathways, with the responses seen at the genomic and protein expression level both *in vivo* and ultimately in *in vitro* test systems as they are developed and validated for specificity and sensitivity, to detect common human response pathways in toxicity.

3 Requirements analysis

3.1 Requirements for risk assessment

Although there are some general requirements for risk assessment, the specifics depend on the particular use case, e.g., regulatory, drug safety assessment, etc. We consider the ontological needs for the three main phases of risk assessment: prioritization, targeted testing and quantitative assessment (see Fig. 1). In a tiered-testing paradigm, these three phases are highly interconnected, with the output of one feeding into another. Prioritization involves ranking dozens or even thousands of chemicals based on some combination of exposure and hazard. This activity can utilize actual data about household or other human exposure scenarios, and laboratory tests in rodents or cell culture models. For most commodity chemicals, however, information on exposure and bioactivity is generally limited so chemical structure activity relations must be utilized. Once the subset of chemicals posing the greatest risk has been identified, targeted tests are conducted to evaluate specific predictions of hazard or of exposure. Ultimately, for a small number of chemicals it may be necessary to evaluate the dose-dependent risk of human toxicity, which is the purpose of quantitative assessment. Ontologies can play a valuable role in organizing information at different stages of risk assessment, which can streamline workflows and increase transparency in support of product development and regulatory decisions.

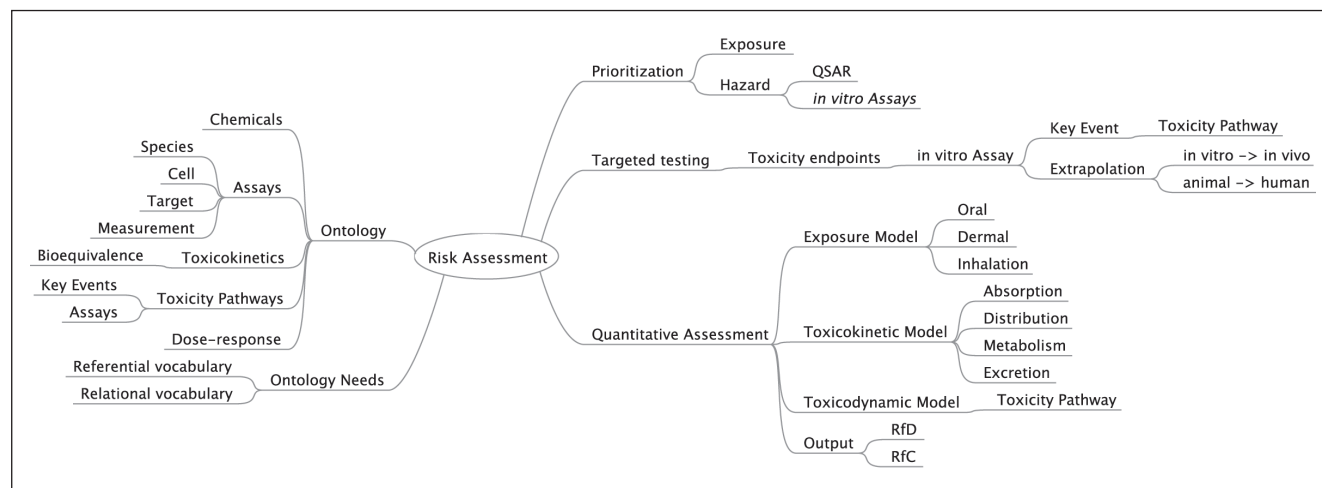


Fig. 1: Mapping risk assessment needs to ontology-based solutions



3.2 Requirements from a consumer goods industry perspective

Risk assessment needs are driven by the requirement for a scientifically robust framework focusing on human relevance and increasingly utilizing non animal alternative approaches, be they in chemico, *in silico*, or *in vitro*. To this end the ontology should act as an enabler aiming to capture the relationship between hazard characterization, dose response, pharmacokinetic parameters and exposure models. Currently there is an information requirement and a need to integrate disparate datasets which relates to the mechanistic approaches currently being developed. As such the vocabulary must provide the granularity to underpin developments in new *in vitro* assays and show the relationship between them and the adverse effect. The emphasis in this case is on the associated metadata to enable transparency and relevance to be assessed. This must also then define the sequential relationship with the biological target and downstream events such as mode of action via biological pathways, disease state progression from a clinical knowledgebase and known toxicology endpoints in order to provide the capability to combine datasets at a mechanistic level. However to progress this towards quantitative risk assessment there is a need for *in vitro* to *in vivo* dose extrapolation, pharmacokinetic data and metabolism parameters and exposure models to be defined without which the assays remain limited to hazard characterization. Secondly it must also provide the scope to incorporate chemical descriptors based on physicochemical properties and structural features to enable predictive *in silico* modeling through the clustering of chemical and biological datasets. This should also further the generation of datasets to build on current exposure-based waiving approaches such as the Threshold of Toxicological Concern (TTC). A further area for risk assessment that should be considered as part of a toxicology ontology would be around population studies incorporating susceptibility factors at both molecular and phenotypic levels at one end while providing context for History of Safe Use risk assessment approaches at the other.

Finally, it should be highlighted that for the above to happen, the quality of the content in terms of simple and consistent formats, aligned to current ontology approaches, such as Open Biological and Biomedical Ontologies' (OBO), to map across current ontologies with well-populated synonyms and use of approved identifiers, are required as a priority. As previously stated, ontologies already exist that provide some coverage of toxicology terminology so care should be taken not to reinvent the wheel. It should also be recognized that the ontology is likely to initially be used more as a structured controlled vocabulary, rather than for performing complex logical reasoning. A lot can be learned from the success of the Gene Ontology and why, despite a number of

shortcomings, it was readily adopted by the bioinformatics community.

3.3 Requirements from a regulator's perspective

Regulatory agencies around the world have a significant interest in the use of toxicology ontologies to support their daily activities as well as enabling new approaches to safety assessment. The ability to look up historical studies on chemicals is critical to avoid requiring sponsors to replicate studies and ontologies should support the integration of such legacy databases. Predictive approaches, including read across and (Q)SAR models, are also important especially as more recent regulatory and guidance documents (such as REACH) have included the submission of *in silico* results as an alternative to *in vitro* or *in vivo* tests. For example, the US FDA recently introduced guidelines for the use of (Q)SAR model results for drug impurities where exposure is below a specific threshold¹⁷. Look up, read across, and (Q)SAR models can also support internal review of compounds including industrial chemicals, pesticides, food contact materials, drugs, metabolites, contaminants, excipients, degradants, and so on. A number of agencies are also grouping legacy databases for prioritization of chemicals for testing, including the Domestic Substances List (DSL) project of Health Canada¹⁸. The use of standardized vocabularies and ontologies has an important role in supporting the electronic submission of dossiers to the agencies. One example is the use of the SEND format for submission of preclinical data to the US FDA¹⁹. The electronic submission of structured information not only streamlines the processing of this data but also avoids data entry errors and supports future *in silico* modeling, since this information can flow directly into an electronic database to support *in silico* assessments. Making non-proprietary data available outside the agencies in a standardized and controlled manner will also support external analyses such as the development of (Q)SAR models. For example, the US FDA has been collaborating with Leadscope, Inc. to develop a series of structured databases based on the ToxML standard that covers genotoxicity, rodent carcinogenicity, acute toxicity, as well as reproductive and developmental toxicity (Arvidson, 2008). The use of toxicology ontologies will also support research within regulatory agencies. This includes research to support the FDA Critical Path Initiative which includes the generation of a "new product development toolkit – containing powerful new scientific and technical methods such as animal- or computer-based predictive models, biomarkers for safety and effectiveness, and new clinical evaluation techniques." Such new tools are urgently needed to improve predictability and efficiency along the critical path from laboratory concept to commercial product.²⁰

¹⁷ USFDA / CDER Impurities Guidance (draft 12/2008)

¹⁸ Domestic Substances List (DSL) project, Health Canada
<http://www.hc-sc.gc.ca/ewh-semt/contaminants/existsub/screen-eval-prealable/summary-sommaire-eng.php>

¹⁹ SEND format for submission of preclinical data to the US FDA, <http://www.fda.gov/ForIndustry/DataStandards/StudyDataStandards/ucm155320.htm>

²⁰ FDA Critical Path Initiative 2004,
<http://www.fda.gov/ScienceResearch/SpecialTopics/CriticalPathInitiative/CriticalPathOpportunitiesReports/ucm077262.htm>

3.4 Requirements from the pharmaceutical industry perspective

The inability of the pharmaceutical and biotechnology industry to accurately predict toxicity is one of the main obstacles to improving drug discovery productivity. By incorporating faster and cheaper methods that have more human relevance across the entire discovery process, the industry could start to reduce the enormous attrition costs associated with late stage compound failures. This industry is increasingly applying *in silico* methods such as (Q)SAR modeling, read across, look up, and consensus modeling approaches throughout discovery and development. This includes the application of these methods in (1) discovery – both prior to synthesis and for prioritization of lead series as well as to understand toxicity mechanisms, (2) development – to evaluate impurities, synthetic intermediates, metabolites as well as exploratory toxicology to understand mechanisms of toxicity, and (3) occupational safety – including the assessment of genotoxicity. Access to legacy toxicity data, associated PBPK data and biological information (including toxicity pathways) data through controlled and linked resources is critical to support this *in silico* assessment.

3.5 Requirements for alternative methods development

The development and validation of new approaches is complex and time-consuming, requiring a multidisciplinary approach. For example, the new EU FP7/COLIPA SEURAT-1 program is developing alternative methods for chronic systemic toxicity and includes over 70 organizations covering many disciplines and types of organizations including academic laboratories, small businesses and government agencies²¹.

A critical component in the development and validation of alternative methods is the construction of databases of reference compounds that contain information on historically tested compounds with specific toxicity data that directly matches the readout of the new assay or other method. Information on the specific toxicity mechanisms is usually required, which can be challenging since these mechanisms can often only be inferred from these historical studies. Linking the data to other omics or adverse event information, particularly with human origins, can provide valuable support to pinpoint these mechanisms or more fully understand the toxicity profile or human relevance. Another challenging aspect of these reference databases is the identification of negative controls since this information is rarely reported in historical toxicity study databases, which highlights the importance of capturing all historical data (including negative findings) in a consistent manner. The toxicology data is only one component necessary for these reference compounds, since their physicochemical properties, their availability, any handling or regulatory restrictions, as well as stability data are also needed.

Toxicology ontologies as well as ontologies from related chemical and biological disciplines can play an important role as a framework to support the development of these new *in vitro*

methodologies as well as supporting an integrated data analysis. These methods will invariably require the development and integration of new cellular biology and engineering methods, as well as the identification of new toxicity biomarkers and use of functional assays. This complex and heterogeneous data must be assimilated to support the development of these new methods. The assessment of toxicity hazards and risks based on these new methods must look across consistent doses and time points to enable an understanding of safe dose levels and suggest biological mechanism for the toxicity. The development of these new approaches as well as the integrated data analysis will be accelerated through the use of linked open ontologies.

3.6 Use cases

We identify the following general use case types as ones which could be prioritized to show the benefits and impact of a toxicology ontology:

1. Support combination of data from safety reports from different studies;
2. Support safety liability assessment of biological targets or “evidence of on-target safety issues”²²;
3. Provide information on what alternative testing methods are potentially matching user’s compound/endpoint situation and can be linked with assays;
4. Provide insight on mode of action, biological mechanisms, kinetics related to data;
5. Enable systematic classification of assays and link to mode of action and mechanism (link to MIABE)²³;
6. Help remove uncertainties with synonyms (disambiguation);
7. Provide guidance on compound features for directions in chemistry to avoid toxicity;
8. Reduce time finding and integrating data;
9. Support reliable merging of toxicology information into resources and applications;
10. Support solutions that constrain user behavior to inputting data in a more reliable and useful manner through template data input tools matching specified ontology.

3.7 Existing ontologies and harmonization

An extensive review of existing ontologies of relevance to a predictive toxicology framework is provided in (Hardy et al., 2012).

Introduction of an ontology approach to evolving safety evaluation will require the promotion of the common language that is implicit to the usefulness of any ontological approach. It is one thing to define ontology as providing a shared vocabulary to describe a formal representation of knowledge, as a set of concepts within a domain, in this case a Toxicology Ontology Roadmap, but quite another to integrate this with existing vocabularies. There are existing glossaries of terms agreed over the years in toxicological pathology, for instance, by the working groups in the professional societies that represent this area of expertise. This is because societies that represent the professional language of medicine and pathology have developed a

²¹ SEURAT-1 Research Program, <http://www.colipa.eu/news-a-events/news/185-unique-research-initiative-launched.html>

²² <http://www.apuit.com/Services/Preclinical-Technologies/Toxicology-and-Pathology.aspx>

²³ MIABE - Minimum Information About a Bioactive Entity, <http://mibbi.org/index.php/Projects/MIABE>



commonly understood series of descriptive terms to define their diagnostic output, which is useful for clinical prognosis in medicine, or defining thresholds for adverse effects in toxicological pathology and subsequently risk assessment.

Historically the development of glossaries of terms in diagnostic pathology has been descriptive, based on microscopic observations documenting features such as the architectural appearance, or tinctorial stain of a feature in a tissue section. In more recent years this has included a mode of action understanding, or even sometimes a receptor-mediated mechanism of action, but the changes and standardization of terms has been slow, due to the need to achieve acceptance and harmonization across the world. Most textbooks and glossaries have long historical lists of alternative names for the feature being described morphologically used to recognize the progression of a pathological process.

The development of an ontological approach to toxicology will rely heavily on the acceptance of an approach that integrates the visual and morphological with the biochemical and molecular mechanistic understanding of either the omic signature and/or critical pathway that defines a pathology, or toxicity for a chemical, or mixture. Some descriptive terms may have to be sacrificed in this vocabulary to acknowledge the fact that measuring and describing toxicological responses will not necessarily be done at a microscopic level, for instance, in the future.

It will be particularly important to expose the more established societies that govern training and best practice in understanding and describing human and toxicological pathology to an ontological approach that is more focused on the identification of molecular signatures and critical pathways for the prognosis for toxic injury, in both man and animals, and how they relate to the process of safety and risk assessment, in particular.

4 The Toxicology Ontology Roadmap

Here we describe the proposed Toxicology Ontology Roadmap.

4.1 Needs

The following critical needs require the development of a public open toxicology ontology standard:

- Need to support an integrating mechanistic framework for research activities and test system development;
- Need to access, integrate and interpret an increasing volume and type of toxicology data;
- Need for improvements in interoperability between toxicology resources;
- Need for improved knowledge systems supporting R&D and risk assessment;
- Need to support communication and collaboration in an increasingly complex, translational and interdisciplinary science involving a large number of diverse stakeholders;
- Need for improved communications of scientific-related health and safety knowledge to consumers.

4.2 Benefits

The following critical benefits will be delivered by the development of a public open toxicology ontology standard:

- Greatly increased capability to reliably and efficiently combine toxicology data and metadata from different sources;
- Significantly increased interoperability between systems increasing industry competitiveness, reducing costs, and providing information benefits for toxicology R&D and risk assessment activities;
- Improved communications and collaboration between stakeholders on toxicology data, meaning, models and knowledge;
- Reduced cost and increased effectiveness of toxicology infrastructure;
- Increased number of applications providing superior toxicology knowledge to both professionals and consumers;
- Support for safer products in the marketplace with improved risk assessment and management.

4.3 Product and scope

The product of this initiative will be a public open toxicology ontology. The scope of the toxicology ontology is defined by its support of toxicology-related use cases and scenarios delivering practical utility and value to the activities of industry, regulators, academia, and consumers.

4.4 Stakeholders

The following critical stakeholders will be involved in the development of a public open toxicology ontology standard:

- Industry groups involved in toxicology research, product development, safety, risk assessment;
- Academic groups involved in toxicology research including emerging areas of new toxicology science supporting the development of alternative testing methods;
- Regulators involved in the assessment of the safety of products such as chemicals, drugs, food, agrochemicals, biologicals, cosmetics, and other consumer products;
- Resource providers including publishers, information suppliers, database and software developers, service organizations, and system integrators;
- Programs involving collaboration, coordination and cluster activities related to toxicology R&D and risk assessment.

4.5 Methods

How should a toxicology ontology be formulated and managed? We suggest the use of ontology development best practices as have been previously developed by the OBO Foundry²⁴. The ontology development should be driven by reviewed and accepted use cases generated by the user community with the combination of existing and new ontologies enabling the successful implementation and testing of the use cases. Online collaborative resources should be used to maintain and provide open access to the ontology and its ongoing updates.

²⁴ <http://obofoundry.org/>

4.6 Phases

Four main phases are envisioned in the roadmap: a) Roadmap Development, b) Piloting, c) Ontology Development and Implementation, and d) Sustainable Development and Maintenance.

Roadmap Development (1 year)

In this phase the following activities will be carried out:

- Establishment of needs, benefits, scope, vision;
- Engagement of stakeholders; supporting their input into roadmap;
- Identification and elaboration of requirements, use cases, scenarios, and storyboards;
- Definition of ontology development work required including integration and harmonization of existing ontologies;
- Documentation of ontology development methods to be pursued;
- Elaboration of scientific and business cases for the ontology;
- Definition of resources required for subsequent phases.

Piloting (1 year)

In this phase the following activities will be carried out:

- Development of ontology for prioritized use cases;
- Implementation of ontology into prioritized use cases;
- User testing and evaluation of prioritized use cases;
- Elaboration of service and business models for prioritized use cases.

Ontology Development & Implementation (3 years)

In this phase the following activities will be carried out:

- Extensive development of ontology for all accepted use cases;
- Implementation of ontology into substantial infrastructure supporting use cases;
- User testing and evaluation of all use cases;
- Deployment and testing of service and business models for use cases.

Sustainable Development & Maintenance (ongoing)

In this phase the following activities will be carried out on an ongoing basis:

- Establishment of sustainable ecosystem of infrastructure, services, business models, and governance structure for ontology standard;
- Ongoing evaluation and acceptance of new cases;
- Implementation of ontology into new use cases;
- User testing and evaluation of new use cases;
- Support and maintenance of ontology.

5 Conclusions

Numerous existing ontology and standards initiatives can contribute to the creation of a toxicology ontology supporting the needs of predictive toxicology and risk assessment. Additionally, there is the need for the creation of new ontologies to satisfy practical use cases and scenarios, to harmonize ontologies and to implement them into infrastructures. A well-coordinated

and sustained effort across numerous stakeholders engaged in a public-private partnership supported by sufficient resources to achieve the goals of the creation and deployment of a toxicology ontology is required. The initiative will require scientific collaboration and consensus on concepts and vocabulary, business model development, establishment of a well-organized governance structure and a sustainable development and support structure for the open standard. Academia will benefit from improved access to scientific knowledge, industry will inherit a more competitive and effective infrastructure for innovation and safety, and society will benefit as a whole from the increased availability of knowledge related to toxicological risk and exposure. Products on the market will have undergone superior evaluation and management of their risks to consumers, supported by a stronger integrated knowledge ecosystem.

References

- Arvidson, K. B. (2008). FDA toxicity databases and real-time data entry. *Toxicol. Appl. Pharmacol.* 233, 17-19.
- Ashburner, M., Ball, C. A., Blake, J. A., et al. (2000). Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat. Genet.* 25, 25-9.
- Hardy, B., Douglas, N., Helma, C., et al. (2010). Collaborative development of predictive toxicology applications. *J. Cheminform.* 2, 7.
- Hardy, B., Apic, G., Carthew, P., et al. (2012). Toxicology ontology perspectives. *ALTEX* 29, 139-156.
- NRC - National Research Council (2007). *Toxicity Testing in the 21st Century: A Vision and a Strategy*. Washington DC, USA: National Academies Press.

Acknowledgements

OpenTox – An Open Source Predictive Toxicology Framework, www.opentox.org, was funded under the EU Seventh Framework Program: HEALTH-2007-1.3-3 Promotion, development, validation, acceptance and implementation of QSARs (Quantitative Structure-Activity Relationships) for toxicology, Project Reference Number Health-F5-2008-200787 (2008-2011). The research leading to these results has received support from the Innovative Medicines Initiative Joint Undertaking under grant agreement n° 115002, resources of which are composed of financial contribution from the European Union's Seventh Framework Programme (FP7/2007-2013) and EFPIA companies' in kind contribution.

Correspondence to

Barry Hardy, PhD
Douglas Connect
Baermeggenweg 14
4314 Zeiningen
Switzerland
Phone: +41 61 851 0170
e-mail: Barry.Hardy@douglasconnect.com