

# Formalizing knowledge and evidence about potential drug-drug interactions

Jodi Schneider<sup>1</sup>, Mathias Brochhausen<sup>2</sup>, Samuel Rosko<sup>1</sup>, Paolo Ciccarese<sup>3,6</sup>,  
William R. Hogan<sup>4</sup>, Daniel Malone<sup>5</sup>, Yifan Ning<sup>1</sup>, Tim Clark<sup>6</sup>, and  
Richard D. Boyce<sup>1</sup>

<sup>1</sup> Department of Biomedical Informatics, University of Pittsburgh  
jos188, scr25, yin2, rdb20@pitt.edu

<sup>2</sup> Department of Biomedical Informatics, University of Arkansas for Medical Sciences  
mbrochhausen@uams.edu

<sup>3</sup> Innovation Lab at PerkinElmer, paolo.ciccarese@gmail.com

<sup>4</sup> University of Florida, hoganwr@ufl.edu

<sup>5</sup> College of Pharmacy, The University of Arizona, malone@pharmacy.arizona.edu

<sup>6</sup> Massachusetts General Hospital and Harvard Medical School  
tim.clark@harvard.edu

**Abstract.** Potential drug-drug interactions (PDDI) are a significant source of preventable drug-related harm. One contributing factor is that there is no standard way to represent PDDI knowledge claims and associated evidence in a computable form. The research we present in this paper addresses this problem by creating a new version of the Drug Interaction Knowledge Base, with scalable, interlinkable repositories for PDDI evidence and PDDI knowledge claims.

**Keywords:** Linked Data, drug-drug interactions, evidence bases, Micropublications, Nanopublications, knowledge bases

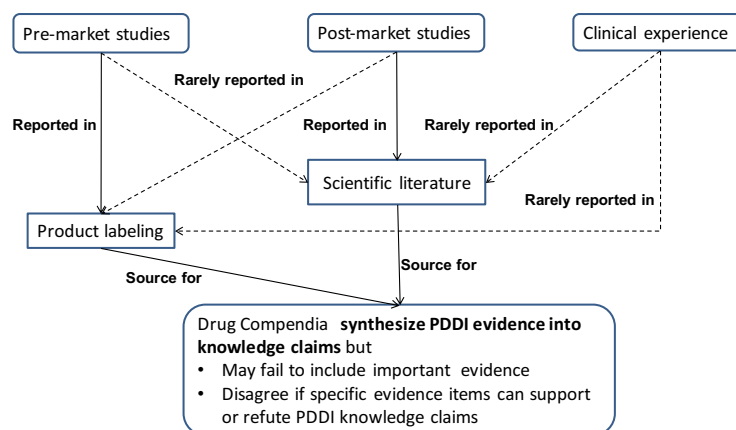
## 1 Introduction

A challenging area of focus for patient safety is the management of potential drug-drug interactions (PDDIs). These are defined as co-prescription or co-administration of two drugs known to interact, which potentially exposes the patient to adverse drug events [9]. PDDIs are a significant source of preventable drug-related harm: according to a recent review, clinically important events attributable to PDDI exposure occur in 5.3% to 14.3% of inpatients, and are responsible for 0.02% to 0.17% of the 129 million emergency department visits that occur each year [12].<sup>1</sup> Unfortunately, most drug information sources disagree substantially in their guidance about specific PDDIs [1, 16, 13, 2]. Addressing this is urgent as United States healthcare organizations consider PDDI screening in their strategies to achieve the effective use of electronic health records.

---

<sup>1</sup> <http://www.cdc.gov/nchs/fastats/ervisits.htm>

There are both technical and social factors underlying the disagreement that exists across drug information sources [14]. As Figure 1 shows, evidence that might be relevant for establishing PDDI knowledge claims is distributed across several sources including product labeling, the scientific literature, and case reports. Each source provides complementary evidence that editors of drug information resources (public sources [2] or proprietary sources such as Micromedex, Epocrates, and Medscape) must synthesize. A major social factor underlying disagreement is that drug information editors have different criteria for assessing evidence. Fortunately, two different conference series have brought leading drug information editors to discuss a standard set of methods for assessing evidence [14, 10].



**Fig. 1.** Editors of drug information resources might seek evidence for or against potential drug-drug interactions from numerous sources. Different information is reported in each type of source, making synthesis necessary.

A major technical factor yet to be addressed is that there currently does not exist a standard way to represent PDDI knowledge claims and associated evidence in a computable form. As a result, drug information editors resort to *ad hoc* information retrieval methods that can yield different sets of evidence to assess [10]. The research we present in this paper addresses this problem by creating scalable, interlinkable repositories for both PDDI evidence and PDDI knowledge claims. This paper describes our new approach. In Section 2, we outline requirements. In Section 3, we discuss the technical details. In Section 4, we present a benchmark analysis that tests the ability of the new approach to scale. After discussion, we conclude the paper.

## 2 Background and requirements

In prior work, we created the original Drug Interaction Knowledge Base (DIKB-old) [3, 4]. The DIKB<sup>2</sup> is an evidence-focused knowledge base designed to support pharmacoepidemiology and clinical decision support. It contains quantitative and qualitative knowledge claims about drug mechanisms and pharmacokinetic drug-drug interactions for over 60 drugs.

Prior work on the DIKB-old focused on development of an evidential approach representing the evidence associated with a scientific claim. The system considers the evidence board as a socio-technical reasoning system that manages both a *knowledge base* and an *evidence base*. The knowledge base holds PDDI knowledge claims while the evidence base stores information artifacts that can be used to support or challenge those claims. PDDI knowledge claims may be direct (e.g., “drug X interacts with drug Y”); or inferred from pharmacological properties (e.g., “drug X inhibits enzyme Q which is important for the clearance of drug Y from the body”).

In prior work [3, 4], all evidence was collected and entered by an evidence board consisting of an informaticist and a minimum of two drug-experts. The board used the following process to manage the evidence and knowledge base components:

1. All members of the board select drugs of interest. This determines the set of PDDI knowledge claims to be investigated.
2. The informaticist conducts a systematic search for evidence that might support or refute the pre-determined PDDI knowledge claims.
3. Retrieved items are filtered by applying study inclusion criteria.
4. Evidence items that meet inclusion criteria are entered into the system where they are linked to specific PDDI knowledge claims and any evidence use assumptions (knowledge claims that must be true for the evidence to hold).
5. A truth value for each knowledge claim is determined based on belief criteria.

Experience with the DIKB-old revealed a great need for improvements to the system that would make this process more efficient. First, a substantial amount of time was spent on reconciling and integrating information from various sources (Figure 1). Decision rationales were not recorded in a computable form and the evidence board did not have a process in place to keep up with relevant new evidence. Furthermore, the DIKB-old was ontologically informal, failed to adopt common biomedical ontology terms<sup>3</sup>, and did not distinguish drug and enzyme classes from individuals. This hindered automated reasoning that integrated external knowledge sources and resulted in treating PDDIs the same as observed

---

<sup>2</sup> When we do not need to distinguish between the old (‘DIKB-old’) and new (‘new DIKB’) versions, we simply mention ‘DIKB’.

<sup>3</sup> For example, the DIKB-old used the predicate ‘substrate of’ to represent the metabolic process of xenobiotic catalysis. However, this predicate was defined without reference to the formally defined biological process (e.g, such as that provided by the Gene Ontology).

drug-drug interactions. In the new system we wished to resolve these issues. We also wished to retain the ability to compute with a logical representation of drug mechanism knowledge claims, using a rule-based theory of how to infer PDDIs from metabolic mechanistic knowledge of how drugs interact [4].

We summarize these as three requirements for the new DIKB:

- R1 Create a maintainable structure that supports evidence entry of data, methods, and materials from multiple sources on an ongoing basis.
- R2 Create computable, logical representations of drug mechanism knowledge claims.
- R3 Link to biological processes while also carefully distinguishing between a drug drug interaction (an actual occurrence in a patient) and a potential drug drug interaction (an information content entity that may exist because of an observation or inference).

Our approach to addressing these criteria are as follows:

- Addressing R1 We adopt the emerging Micropublications (MP) [6] model for literature integration using ‘argument graphs’ to represent published claims as formal assertions linked to primary data and resources.
- Addressing R2 We extended the MP ontology to add two new properties, `MP:formalizedAs`/`MP:formalizes`, to enable natural language claims to be linked to useful logical formalizations.
- Addressing R3 To stress that potential drug drug interactions are information artifacts, we use a new ontology called DIDEO [5] which has several advantages. DIDEO:
  - (a) Reuses identifiers from existing ontologies (e.g., CHEBI, PRO) that represent biological entities and processes;
  - (b) Differentiates between the representation (statements about drugs and drug-drug interactions) and the represented (actual drug-drug interactions);
  - (c) Prevents unwanted existential import (further explained in Section 3.3 below); and
  - (d) Distinguishes between the type of a drug or enzyme and portions of a specific drug or enzyme, by using punning.

### 3 Technical implementation

#### 3.1 Create a maintainable structure that supports evidence entry of data, methods, and materials from multiple sources

We used micropublications to create a structure that supports evidence entry of data, methods, and materials from multiple sources [15]. We now represent PDDI knowledge claims and supporting evidence as queryable RDF statements<sup>4</sup> constructed using the Micropublication ontology (MP) [6]. PDDI knowledge claims

---

<sup>4</sup> Queryable at <http://purl.org/net/nlprepository/swat-4-med-safety-sparql-endpoint>

and evidence were transformed from the DIKB-old model into the new one using Python scripts. Drug identifiers were converted to ChEBI identifiers to enable the use of DIDEO. So far, the mapping has been completed for 70% of the drugs that had data from clinical studies or mechanistic experiments. We envision that additional DIKB micropublications could be created by multiple parties, including evidence boards, and potentially the original authors, as we describe in Section 5.

Figure 2 shows the generic form of a DIKB micropublication graph using the example erythromycin - simvastatin interaction. Notice that MP has rigorously defined ontology classes that support the DIKB evidence curation process discussed above. In MP, the primary object of interest is the claim. A claim is supported by methods, materials, and data:

- MP:Claim, a text string representing a scientific claim.
- MP:Method, representing a scientific method.
- MP:Materials, for materials, such as study participants and drugs.
- MP:Data, such as the area under the concentration curve (AUC).

These are used for entering evidence and later, the evidence is used to determine truth values for the claims that the evidence supports.

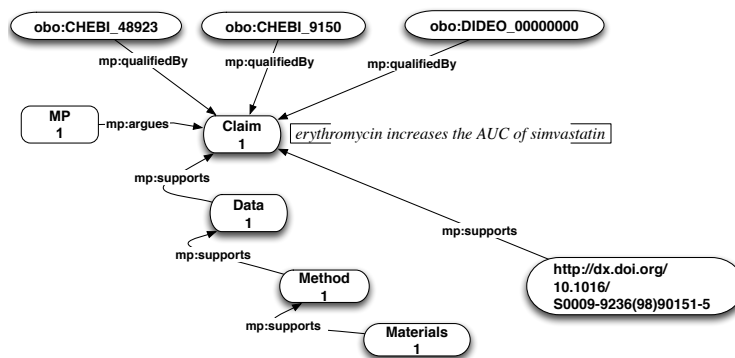


Fig. 2. DIKB micropublication graph for the erythromycin - simvastatin interaction.

The process for managing the evidence base and knowledge base described in Section 2 includes assessing the truth value of each PDDI knowledge claim using belief criteria. Operationally, the evidence board uses labels from a taxonomy of evidence types<sup>5</sup> to tag each evidence item as it is entered into the evidence base. The board then decides which evidence types are credible for specific types of PDDI knowledge claims: this specifies a belief criterion.

As an example, the evidence board might decide that, to support a claim that a drug is a substrate of an enzyme, only clinical drug-drug interaction studies

<sup>5</sup> <http://purl.org/net/drug-interaction-knowledge-base/evidence-types-and-inclusion-criteria>

are admissible. This would become a belief criterion for all ‘substrate of’ claims. To implement a belief criterion in the new DIKB, the evidence base is queried to find all PDDI knowledge claims that have at least one supporting evidence item meeting the criterion. The resulting claims are assigned the value of ‘True’.

### 3.2 Create computable, logical representations

PDDI knowledge claims mention specific entities such as drugs, drug metabolites, enzymes, and biological pathways whose relationships with each other are more generally modeled in a rule-based theory that infers PDDIs [4]. Sources external to the DIKB provide additional formalized knowledge about these entities. For example, the Gene Ontology provides cellular location and molecular function for the enzyme CYP3A4; this is relevant when the evidence board seeks information about gene expression and about enzyme metabolization.

The spans of unstructured text in `MP:Claim` are not inherently computable entities, and the semantic qualifiers (`MP:qualifiedBy`) cannot specify the order (i.e. separate the object drug from the precipitant drug). Therefore, we extended the MP ontology to add two new properties, `MP:formalizedAs/MP:formalizes`, that enable natural language claims represented as `MP:Claim` to be linked to their logical representation.

RDF is also the language chosen for the formalization of `MP:Claim` resources, so that a single query language (i.e., SPARQL) can be used to retrieve information from the whole evidence base. We chose to represent the logical form of knowledge claims using OWL for two reasons. First, OWL provides classes and properties that enable the representation of logical statements in RDF. Second, logical statements written in OWL can be checked for logical consistency and new inferences by a reasoner such as Hermit [7].

We chose to formalize claims using the Nanopublication (NP) [8] ontology because:

1. NP provides a class called `NP:Assertion` that can hold any RDF graph, including logical statements written in OWL.
2. OWL logical statements stored as an `NP:Assertion` can be integrated into full nanopublications that combine the `NP:Assertion`, the provenance of the assertion, and the provenance of the nanopublication into a single publishable and citable entity.

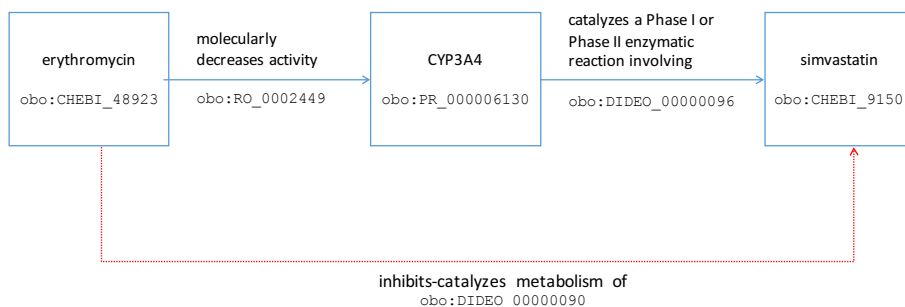
A nanopublication represents the logical structure of a claim as an RDF graph. Like micropublications, nanopublications are publishable and citable entities. Their citability and use of provenance enable us to make the evidence review process transparent and auditable. The uptake of nanopublications by the wider community suggests that nanopublication is a relevant publishing mechanism for reconsumption by others.<sup>6</sup> Unlike micropublications, nanopublications have no

<sup>6</sup> One measure of uptake is the variety of authors of papers using nanopublications; see the bibliography at [http://nanopub.org/wordpress/?page\\_id=638](http://nanopub.org/wordpress/?page_id=638). Another is the size and geographic distribution of current nanopublication datasets: see [11] Table 1 and Figure 3, respectively.

explicit evidence structure and do not support claim conflict. They are therefore complementary to micropublications, which provide these missing features.

### 3.3 Handling reasonable extrapolation

Reasonable extrapolation is an important way to infer a PDDI. In contrast to drug-drug interactions (DDIs) that are based on observing an actual drug-drug interaction in some patient, inferred PDDIs based on reasonable extrapolation might not actually occur in reality. Since we do not know whether a PDDI occurs, we cannot assume the existence of any instance of a drug interaction. To model this correctly, we differentiate between actual drug interactions and statements about PDDIs using the DIDEO ontology [5].

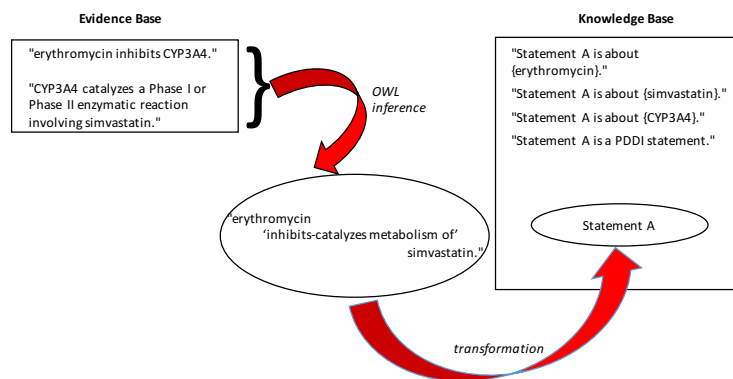


**Fig. 3.** OWL inference (dashed red arrow) of a potential drug drug interaction between erythromycin and simvastatin based on assertions in the evidence base.

Let us again consider erythromycin and simvastatin. Previously (Figure 2) we described a clinical study establishing this PDDI. We now reconsider this drug pair (e.g. in the absence of that clinical evidence), to show how a computable structure and OWL inferencing allow the new DIKB to compute a PDDI by reasonable extrapolation, as shown in Figure 3.

Assume that credible evidence also establishes that erythromycin inhibits the catalysis of CYP3A4 which is important to the metabolic clearance of simvastatin. We represent this relationship in the evidence base by creating OWL statements representing the following two knowledge claims (Figure 3): “erythromycin molecularly decreases activity of CYP3A4” and “CYP3A4 catalyzes a Phase I or Phase II enzymatic reaction involving simvastatin”. These claims are linked to the relevant evidence support using MP as discussed above, and formalized by two different `NP:Assertion` resources using RDF/OWL. If the evidence support meets the belief criteria specified by the evidence board, full nanopublications are created for the formalized claims. These are then passed to an OWL reasoner that would infer a relation between erythromycin and simvastatin called *inhibits-catalyses metabolism*. This inference generates a new entity,

namely an individual PDDI, based on the inferred statement, as shown on the left side of Figure 4. This individual resides in the knowledge base. The knowledge base contains other individuals based on inferred statements; it also contains entities that represent individual PDDIs based on data items created by clinical studies, clinical observation or physiological experiments. The key point here is that the inferred individuals do not refer to any class of specific drug-drug interactions, and so they do not imply the existence of at least one member of the class (known as existential import).



**Fig. 4.** Each OWL-inferred assertion in the evidence base generates a new individual in the knowledge base. Braces indicate punned entities, referring to types, not classes.

We also distinguish between references to a type and references to a class of collection of things. As an example, consider the two statements: “simvastatin treats dyslipidemia.” and “simvastatin was used in the treatment of patient John Doe.” The latter sentence refers to a collection of specific portions of simvastatin, whereas the former sentence refers to the entire type. In the context of PDDIs this allows us to differentiate specific portions of a drug used in for example a case study from the drug as a type. Traditionally, OWL represents classes and individuals, but not types. However, OWL 2 enables the representation of types using ‘punning’.<sup>7</sup>

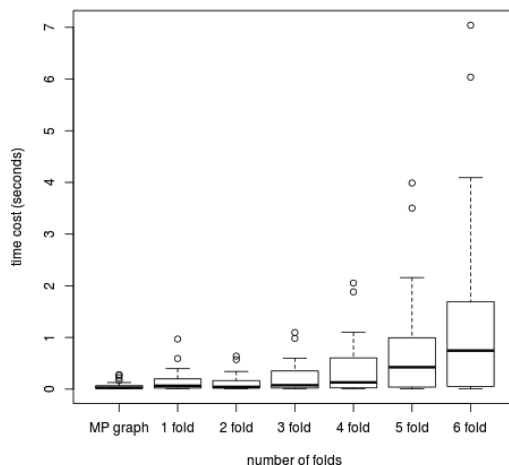
Punning means that we use a URI that is assigned to a class (for instance the class ‘simvastatin’) and also assign it to an individual. For instance, the ChEBI and PRO URIs are used to pun individuals that are intended to represent the type of drug or enzyme. To represent the result of the inference shown in Figure 3, we create a new individual which we will be able to link to both the drugs and the enzyme involved. As is shown in Figure 4, we use the punned URI for the ingredients, drug products and enzymes involved, to express that we are talking about the types, not about actual portions of them.

<sup>7</sup> [http://www.w3.org/TR/owl2-new-features/#F12:\\_Punning](http://www.w3.org/TR/owl2-new-features/#F12:_Punning)



## 4 Results and Evaluation

To evaluate the feasibility of using the model in practice, we investigated how long querying will take. As a benchmark, we execute a set of 31 queries<sup>8</sup> across 7 different sizes of MP graphs<sup>9</sup>, from a single Virtuoso endpoint. Results are shown in Figure 5.



**Fig. 5.** Range of performance: minimum and maximum run times for all 31 queries.

To generate  $G_0$ , the first MP graph, the original DDI data is extracted from the DIKB-old and reshaped into the MP model, along with its supporting evidence. We include DIKB-old data with assertions that are typed as ‘inhibits’, ‘substrate of’, ‘increase AUC’, making  $G_0$  2.0 MB with 16,670 triples and 3498 entities. We also generated 6 additional graphs, with  $G_i$  containing  $2^i$  ( $i = 1$  to 6) copies of the  $G_0$  graph. Claims, methods, materials, and data in these MP graphs have the same relationships as in  $G_0$  so that all items are evenly distributed in the testing graphs. One limitation is that, although this is a real-world graph, we scale by making multiple copies of the same entities and relationships.

## 5 Discussion

Drug information resources are created by groups, which may include an information specialist, drug specialists, etc. Currently, multiple different editorial

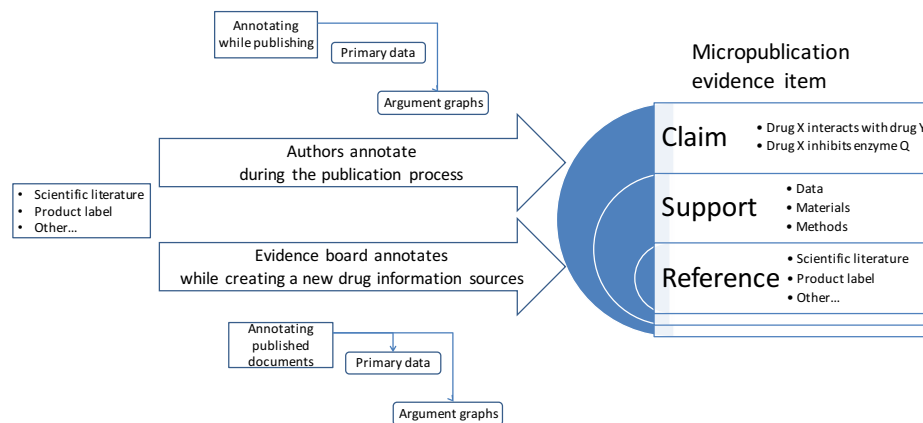
<sup>8</sup> <http://purl.org/net/drug-interaction-knowledge-base/micropublication-queries>

<sup>9</sup> <http://purl.org/net/drug-interaction-knowledge-base/published-NP-and-MP-graphs>

boards undertake this process entirely independently. Evidence from scientific studies, new drug applications, and product labeling is first searched for, and then evaluated. Complex cases are reviewed by multiple people in order to resolve issues, such as to determine whether inclusion criteria are met.

Currently, each group must independently spend a substantial amount of time on reconciling and integrating information from various sources. This slows the process, because multiple different systems, including literature alerts, spreadsheets, and external databases, must be combined in order to evaluate whether a PDDI should be recorded. One new potential is for some of the work of centralizing information to be shared, even across different proprietary sources. Even for commercial entities, shared data resources can be beneficial: For instance, the Open PHACTS public-private partnership has received commercial funding to pre-process and integrate certain aspects of pharmaceutical data into nanopublications [17], which saves individual drug companies time and resources in preparing proprietary data.

*Maintainability* Maintainability will also be a key factor in ensuring the success of the new DIKB. We have already transformed existing annotations into the new format. We are starting to construct a user-friendly annotation process, so that evidence board members and their delegates can annotate knowledge claims and primary data. Our ongoing research will suggest optimal approaches for publishing PDDI claims and evidence in a structured form.



**Fig. 6.** Structured publications such as micropublications could make the evidence search and synthesis process more efficient and scalable. Two routes for creating PDDI knowledge claims and evidence: authors could annotate micropublications for their own papers (pre-publication) or an evidence board could create them (post-publication).

Structured publishing is of increasing interest, and the two main sources of PDDI evidence (product labels and scientific papers) are already moving toward

more structured representations. All drug product labels in the United States are written by drug companies and submitted to the U.S. Food and Drug Administration using the Structured Product Labeling (SPLs) standard<sup>10</sup>. Scientific publishers may in the future deeply interlink scientific papers and data; this is a proposed service of Elsevier, the largest scientific publisher<sup>11</sup>. For authors to annotate their own work with DIKB-compatible micropublications, and for publishers to redistribute micropublications, thus seems conceivable.

If either authors or drug editorial boards share DIKB-compatible micropublications, other groups who need to process the same evidence would also benefit because they would not need to separately integrate and reconcile the information artifacts that support PDDI knowledge claims. This greatly improves maintainability.

*Shared Representation* In the contested area of drug compendium information, centralizing collective information has another benefit: it can help make visible disagreements and discrepancies between these sources, which can be difficult to resolve. Currently, there is no way to answer questions such as: did compendium A and compendium B consider the same evidence? Only the final assertions—lists and severity rankings of drug-drug interactions—are exposed; the underlying evidence is not retrievable or comparable, without directly approaching evidence boards, one by one.

We see the potential for deeper structures, that represent intermediary steps of the process in a shareable format. The evidence base of the new DIKB is a shared representation to centralize and unify the information, and to support tasks such as commenting on evidence, rejecting evidence, and asking for opinions about evidence. This would also allow, for instance, citing a particular statement, indicating a disagreement with it, or marking some evidence as more credible than others. We believe that a computable, shared representation could enable compendium editors to more easily integrate and cross-check different sources of information – as well as to compare different potential interpretations, amongst different editorial groups.

## 6 Conclusions & Future work

As we have discussed, the new DIKB has several advantages compared to the DIKB-old. It is designed to meet 3 key requirements: maintainability, computability, and the ability to link to biological processes while retaining logical consistency.

In future work we will test whether the DIKB does in fact simplify the editor’s task of searching and synthesizing clinically relevant PDDI information. In particular, we aim to show how clinicians and others interested in PDDIs

---

<sup>10</sup> <http://www.fda.gov/downloads/Drugs/GuidanceComplianceRegulatoryInformation/Guidances/ucm072317.pdf>

<sup>11</sup> <http://libraryconnect.elsevier.com/articles/best-practices/2013-02/research-data-driving-new-services>

can use the DIKB to quickly identify the rationale for any given assertion of knowledge. We plan to compare existing tools to new search tools derived from the new DIKB, in terms of the completeness of information retrieval by drug experts using the tools.

One major advantage of the DIKB is that its evidence base serves as a shared representation, recording the evidence assessed and the results taken into account. To further test the relevance of recording this information explicitly, in the future we will explore an architecture enabling multiple ‘possible worlds’ to be described: multiple knowledge bases, generated from different belief criteria.

We hope that in the future some of the bookkeeping work of compendium creation can be shared—ideally amongst authors annotating their own materials—and certainly amongst editors of different compendia. We will advocate for authors to annotate publications of various sorts with DIKB-compatible micropublications once appropriate tools can be made available. This route has great potential to reduce much of the bookkeeping work conducted by the evidence board because PDDI knowledge claims and the data supporting those claims would be annotated by the individuals who wrote the submission. It would also help evidence avoid becoming ‘stale’ because the new evidence would be linked directly to knowledge claims as they are published.

Since our approach is compatible with other ontologies in the Open Biomedical Ontologies (OBO) Foundry through the reuse of identifiers, the import of external information is simplified compared to the previous approach. With the new DIKB architecture, as gene and protein knowledge is formalized in external sources, that knowledge can be pulled in and queried, in ways that may meet future use cases for other audiences.

## Acknowledgments

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 246016 for the ERCIM “Alain Bensoussan” Fellowship Programme; from a grant from the National Library of Medicine (1R01LM011838-01); and training grant 5T15LM007059-29 from the National Library of Medicine/National Institute of Dental and Craniofacial Research. We thank Carol Collins, Amy Grizzle, Lisa Hines, Harry Hochheiser, and John R Horn.

## References

1. Abarca, J., Malone, D.C., Armstrong, E.P., Grizzle, A.J., Hansten, P.D., Van Bergen, R.C., Lipton, R.B.: Concordance of severity ratings provided in four drug interaction compendia. *Journal of the American Pharmacists Association* 44(2), 136–141 (2003)
2. Ayvaz, S., Horn, J., Hassanzadeh, O., Zhu, Q., Stan, J., Tatonetti, N.P., Vilar, S., Brochhausen, M., Samwald, M., Rastegar-Mojarad, M., et al.: Toward a complete dataset of drug-drug interaction information from publicly available sources. *Journal of Biomedical Informatics* 55, 206–217 (2015)

3. Boyce, R., Collins, C., Horn, J., Kalet, I.: Computing with evidence: Part I. *Journal of Biomedical Informatics* 42(6), 979–989 (2009)
4. Boyce, R., Collins, C., Horn, J., Kalet, I.: Computing with evidence: Part II. *Journal of Biomedical Informatics* 42(6), 990–1003 (2009)
5. Brochhausen, M., Schneider, J., Malone, D., Empey, P., Hogan, W.R., Boyce, R.D.: Towards a foundational representation of potential drug-drug interaction knowledge. In: *Drug Interaction Knowledge Representation at ICBO* (2014)
6. Clark, T., Ciccarese, P., Goble, C.: Micropublications: a semantic model for claims, evidence, arguments and annotations in biomedical communications. *Journal of Biomedical Semantics* 5(1), 28 (2014)
7. Glimm, B., Horrocks, I., Motik, B., Stoilos, G., Wang, Z.: Hermit. *Journal of Automated Reasoning* 53(3), 245–269 (2014)
8. Groth, P., Gibson, A., Velterop, J.: The anatomy of a nanopublication. *Information Services and Use* 30(1-2), 51–56 (2010)
9. Hines, L.E., Murphy, J.E.: Potentially harmful drug-drug interactions in the elderly: a review. *The American Journal of Geriatric Pharmacotherapy* 9(6), 364–377 (2011)
10. Hines, L.E., Murphy, J.E., Grizzle, A.J., Malone, D.C.: Critical issues associated with drug-drug interactions. *American Journal of Health-system Pharmacy* 68(10), 941–946 (2011)
11. Kuhn, T., Chichester, C., Dumontier, M., Krauthammer, M.: Publishing without publishers: a decentralized approach to dissemination, retrieval, and archiving of data. In: *ISWC* (2015)
12. Magro, L., Moretti, U., Leone, R.: Epidemiology and characteristics of adverse drug reactions caused by drug-drug interactions. *Expert Opinion on Drug Safety* 11(1), 83–94 (2012)
13. Saverno, K.R., Hines, L.E., Warholak, T.L., Grizzle, A.J., Babits, L., Clark, C., Taylor, A.M., Malone, D.C.: Ability of pharmacy clinical decision-support software to alert users about clinically important drug-drug interactions. *Journal of the American Medical Informatics Association* 18(1), 32–37 (2011)
14. Scheife, R.T., Hines, L.E., Boyce, R.D., Chung, S.P., Momper, J.D., Sommer, C.D., Abernethy, D.R., Horn, J.R., Sklar, S.J., Wong, S.K., et al.: Consensus recommendations for systematic evaluation of drug-drug interaction evidence for clinical decision support. *Drug Safety* 38(2), 197–206 (2015)
15. Schneider, J., Ciccarese, P., Clark, T., Boyce, R.: Using the micropublications ontology and the open annotation data model to represent evidence within a drug-drug interaction knowledge base. In: *Linked Science at ISWC* (2014)
16. Wang, L.M., Wong, M., Lightwood, J.M., Cheng, C.M.: Black box warning contraindicated comedications: concordance among three major drug interaction screening programs. *Annals of Pharmacotherapy* 44(1), 28–34 (2010)
17. Williams, A.J., Harland, L., Groth, P., Pettifer, S., Chichester, C., Willighagen, E.L., Evelo, C.T., Blomberg, N., Ecker, G., Goble, C., et al.: Open phacts: semantic interoperability for drug discovery. *Drug discovery today* 17(21), 1188–1198 (2012)