

Provenance-Centered Dataset of Drug-Drug Interactions

Juan M. Banda¹(✉), Tobias Kuhn^{2,3}, Nigam H. Shah¹, and Michel Dumontier¹

¹ Stanford University - Center for Biomedical Informatics Research,
1265 Welch Road, Stanford, CA 94305, USA

{jmbanda,nigam,michel.dumontier}@stanford.edu

² Department of Humanities, Social and Political Sciences, ETH Zurich,
Zürich, Switzerland

tokuhn@ethz.ch

³ Department of Computer Science, VU University Amsterdam,
Amsterdam, Netherlands

Abstract. Over the years several studies have demonstrated the ability to identify potential drug-drug interactions via data mining from the literature (MEDLINE), electronic health records, public databases (Drugbank), etc. While each one of these approaches is properly statistically validated, they do not take into consideration the overlap between them as one of their decision making variables. In this paper we present LInked Drug-Drug Interactions (LIDDI), a public nanopublication-based RDF dataset with trusty URIs that encompasses some of the most cited prediction methods and sources to provide researchers a resource for leveraging the work of others into their prediction methods. As one of the main issues to overcome the usage of external resources is their mappings between drug names and identifiers used, we also provide the set of mappings we curated to be able to compare the multiple sources we aggregate in our dataset.

Keywords: Drug-drug interactions · Nanopublications · Data mining

1 Introduction

Studies analyzing costs over time have shown that adverse drug reactions (ADRs) cost over \$136 billion a year [13]. One significant cause of ADRs are drug-drug interactions (DDIs) which greatly affect older adults due to the multiple drugs they are taking [3]. A DDI occurs when the effect of any given drug is altered by another drug which results in an unpredictable effect. While new drugs, before market approval, are tested in both *in vivo* and *in vitro* methods [19], it is unfeasible to test their interactions with all other approved and experimental drugs. In the recent years, computational approaches have been trying to infer potential DDI signals using a wide variety of sources [6, 9, 16, 17], however these methods produce thousands of statistically plausible predictions [6, 8], thus making the task of testing them in an experimental setting impractical.

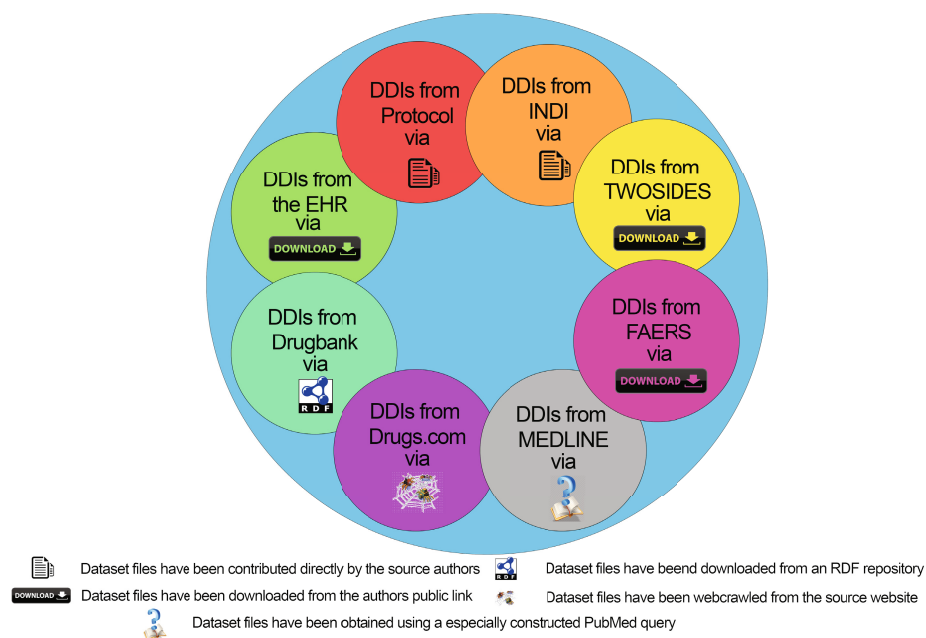


Fig. 1. Overview of the Linked Drug-Drug Interactions dataset and the data sources that it incorporates and their original formats

We are therefore facing the need of combining data from various disparate sources, but at the same time have to remain aware of their provenance due to differences in individual quality and overall confidence. Depending on their primary source and their provenance path, two data entries about the same DDI might substantially increase our confidence about a relation. By relying on Semantic Web technologies such as RDF – and more specifically nanopublications [7] – these provenance paths can be represented in an explicit and uniform manner. In this paper we introduce LInked Drug-Drug Interactions (LIDDI), a dataset that consolidates multiple data sources in one cohesive linked place that allows researchers to have immediate access to multiple collections of DDI predictions from public databases and reports, biomedical literature, and methods that take different and sometimes more comprehensive approaches.

With great potential for use in drug safety studies, LIDDI provides the linking components to branch from DDIs into individual properties of each drug via Drugbank and UMLS, as well as mappings to other types of biomedical ontologies which are used to describe drug-drug interactions. In terms of direct applications, the resources in our dataset allow researchers to quickly determine if there is support by other DDI sources for their own predictions, allowing them to evaluate them via consensus rather than only by independent statistical measures (AUROC values, odds ratios, etc). LIDDI offers all the mappings needed to bridge the incorporated resources into a single comparable entity, providing extra value

for researchers looking to bridge the data sources we have connected, for example Drugbank drug identifiers to UMLS Concept Unique Identifiers (CUI) to Medical Subject Headings (MeSH) codes, opening endless possibilities for reuse.

To our knowledge, this is the first public dataset exemplifying the use of nanopublications for the integration of knowledge from multiple diverse data sources in the field of drug safety surveillance. We believe it will facilitate cross-referencing of results between different domains not previously available, as well as the enrichment of drugs involved in the DDIs thanks to the linkages provided.

2 Methods

2.1 Nanopublications with Trusty URIs

Nanopublications [7] are a concept to use Semantic Web techniques (most importantly RDF and named graphs) to closely link data to their provenance and meta-data in a uniform manner. The vision is that small data packages should become the primary format for research outputs instead of narrative articles [15]. Technically, a nanopublication consists of an assertion graph with triples expressing an atomic statement (about drug-drug interactions in our case), a provenance graph that reports how this assertion came about (e.g. where it was extracted from or what mechanism was used to derive it), and a publication information graph that provides meta-data for the nanopublication (such as its creators and a timestamp).

Trusty URIs [11,12] are a recent proposal to make URI links for digital artifacts verifiable, immutable, and permanent. Such URI identifiers contain a cryptographic hash value calculated on the content of the represented digital artifact — i.e. RDF content of nanopublications in our case — in a format-independent way. With trusty URIs, any reference to an artifact thereby comes with the possibility to verify with 100% confidence that a retrieved file really represents the correct and original state of that resource.

2.2 Data Schema

We represent DDIs in the following way: We instantiate the DDI class from the Semantic science Integrated Ontology (SIO) [5]. We define a vocabulary in our own namespace, as is normally done with Bio2RDF [4], to assign domain-specific predicates to point from the DDI to each drug, and to the clinically-relevant event resulting from the drug interaction. This assertion is stored in the assertion graph of the nanopublication. We also create nanopublications to capture the mappings between each drug and our seven sources. We use BioPortal PURL identifiers for UMLS URIs. We use the PROV ontology [14] to capture the provenance of each DDI by linking the assertion graph URI to a PROV activity. This activity is linked to the software used to generate the drug and event mappings, as well as to the citation for the method, the generation time, and the dataset where the DDIs were extracted from. In this last case we also provide a direct link to the original dataset used. Figure 2 highlights our general schema with a particular example for a DDI that contains a drug-drug set.

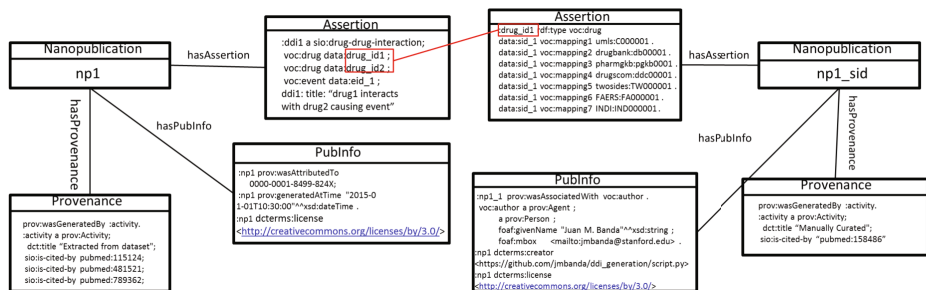


Fig. 2. Dataset Schema shows two types of nanopublications included in the dataset. The first captures the DDI, and the second captures drug mappings. A third one not show captures the event mappings.

3 Data Sources Aggregated

In this section we provide details about the transformations we performed on each of the data sources to be able to normalize them and use them together. Due to space constraints we provide a very high-level view of how each source procures and generates its data as this lies out of the scope of this dataset paper. All the mappings provided within our dataset for UMLS concept unique identifiers (CUI), RxNORM codes, MeSH codes and Drugbank identifiers, were done as an initial step once we acquired all our sources.

Electronic Health Records (EHR). In [9] the authors used Stanford Translational Research Integrated Database Environment (STRIDE) dataset comprising 9 million unstructured clinical notes corresponding to 1 million patients that span a period of 18 years (1994–2011) to mine DDIs from the EHR. Published as a comma delimited file in the supplemental materials, we use the highest confidence predictions as the initial basis of our dataset in terms of drugs (345) and events (10). This dataset provides drug-drug-event sets of UMLS concept unique identifiers (CUI), drug/event names in string literals.

FDA Adverse Event Reporting System (FAERS). We analyzed over 3.2 million reports found on the FDA Adverse Event Reporting System (FAERS), since there are no established algorithmic ways of determining the statistical significance of a drug-drug interaction from FAERS reports without any additional external information, we set a threshold of each DDI to have at least ten appearances in different FAERS reports for it to appear on our dataset. In this data source, the drugs are given in string literals that normalized to RxNorm drug names and the events normalized to MeDRA.

DrugBank. We used the RDF versions of Drugbank [18] made publicly available through Bio2RDF. The used version was dated as 2013-07-25. We proceeded to strip the Drugbank identifiers and their labels into string literals mapped

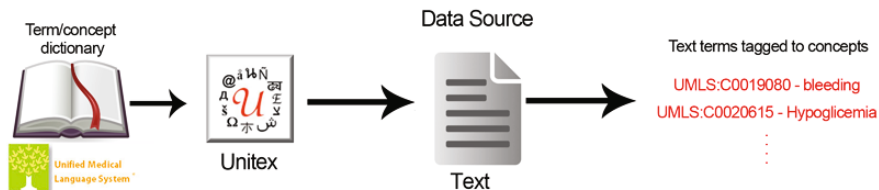


Fig. 3. Text-tagging pipeline. We build a term/concept dictionary from ontologies in UMLS. We then proceed to tag concepts found in any piece of text from any source via Unitex, resulting in a bag of UMLS CUIs for each input.

into RxNorm/SNOMED string literals. In order to extract DDIs we used the `drugbank_vocabulary:Drug-Drug-Interaction` type to match interacting drugs, we then used the `rdfs:label` to extract the event by concept tagging the indication section with a modified version of the NCBO annotator used in [9]. This concept tagging process is described on Figure 3.

Drugs.com. In order to mine DDIs from this resource, we used a web-crawling agent developed in Python. We programmatically checked the individual `drugs.com` page for each of the 345 drugs in our set and extracted all their interacting drugs. For each of the interacting drugs we build the event part of our DDI and by annotating the complementary free-text section, the same way as on Figure 3, which describes the interaction (similar to the Drugbank indication).

MEDLINE. In order to find potential DDIs signals in the biomedical literature, we have adapted a clever method developed by Avilach et al. [1], that allows for the use of a query with certain MeSH terms and qualifiers to find adverse events in the literature. We modified this query by adding a MeSH term for “Drug interactions” and modified the search to allow for two drugs instead of one from the original approach. These modifications of the method with some manipulation produces drug-drug-event sets.

TWOSIDES. Provided as a download by Tatonetti et al. [16], TWOSIDES contains polypharmacy side effects for pairs of drugs. Mined from FDAs FAERS, this method is designed to complement modern signal detection approaches, providing stratification, dampening or removing the effect of co-variables, without needing to divide drug-exposed reports into strata. This resource is available in a CSV format and features drugs in RxNorm normalized string literal form and events in UMLS CUI form.

Inferring Drug Interactions (INDI). Currently available on a website for checking drug-drug interactions and provided to us by the authors [6] in a comma separated file. This prediction method infers both pharmacodynamic and pharmacokinetic interactions based on their similarity to existing known adverse events that result from a common metabolizing enzyme (CYP). The provided resources consisted of drug-drug sets (no events) in UMLS CUI form.

Table 1. LIDDI DDI event distribution between sources

Event Name	EHR	MEDLINE	Drugbank	Drugs.com	FAERS	TWOSIDES	INDI*	Protocol*
Arrhythmia	700	68	286	3,148	100	4,632	NA*	NA*
Bradycardia	254	88	408	4,896	194	4,824	NA*	NA*
Hyperkalaemia	1,888	42	422	4,248	146	3,840	NA*	NA*
Hypoglycaemia	1,460	386	796	6,214	104	5,150	NA*	NA*
Long QT syndrome	14	270	334	3,510	2	0	NA*	NA*
Neutropenia	4,608	192	402	4,218	616	3,702	NA*	NA*
Pancytopenia	1,880	4	270	3,146	148	5,440	NA*	NA*
Parkinsonism	144	0	566	5,978	70	884	NA*	NA*
Rhabdomyolysis	122	198	392	3,842	214	3,264	NA*	NA*
Serotonin syndrome	896	0	384	3,960	122	1,094	NA*	NA*
Total:	11,966	1,248	4,260	43,160	1,716	32,830	8,370	224

Similarity-Based Modeling Protocol (Protocol). Provided to us by the authors, Vilar et al. [17], this protocol integrates a gold standard of DDIs with drug similarity information extracted from sources like: 2D and 3D molecular structure, interaction profile, target similarities, and side-effect similarities. This method generates drug interaction candidates that are traceable to pharmacological or clinical effects. We were provided with the resulting drug-drug (no events) sets with Drugbank identifiers.

4 Statistics and Access

Table 1 shows the total number of DDIs found with support in each of the multiple sources we integrated. Note that these are not unique DDIs in the sense that we have a drug1-drug2-event set that is equivalent to drug2-drug1-event set as the directionality does not matter, thus counting them twice. We left this reflexive drug drug interactions in our dataset to conform to the standard employed by bio2RDF’s DrugBank repository, in order to not limit the discovery of any potential interaction when a query is performed on any given single drug.

The dataset in its entirety contains a total of 98,085 nanopublications out of which 345 are used for drug mappings, 10 for event mappings and the remainder for DDIs extracted from the data sources we used. The dataset has a total of 392,340 graphs (four per nanopublications) and 2,051,959 triples, taking 723 MB in nquads representation.

LIDDI can be accessed as a bulk download via figShare [2] and SPARQL endpoint at: <http://http://liddi.stanford.edu:8890/sparql>. The dataset is also made available via the recently installed nanopublication server network [10]. The command ‘np’ from the Java nanopublication library¹ can be used to download the dataset [20]:

```
$ np get -c -o out.trig RA7SuQ0e661LJdKpt5E0S2DKykf1ht9LFmNaZtFSDMrXg
```

¹ <https://github.com/Nanopublication/nanopub-java>

5 Potential for Applications and Future Work

The first application that this dataset has been used for involves ranking DDI predictions from EHR data. By taking the EHR data source published by [9] and using this resource, researchers are ranking predictions in order to determine which ones are more feasible for experimental evaluation. By leveraging such diverse resources, the researchers have been able to use the data sources as voting mechanisms for certain DDIs to be prioritized over others. Similar studies can be performed to determine the priority of the other data sources and their predictions just using LIDDI and selecting a different resource to evaluate. The authors of each data source could enhance their own predictions by using their method and consider other sources we compiled in this dataset as gold-standards or at least 'silver-standards' as there is no de facto community-wide gold standard for drug safety when it comes to drug-drug interactions. Multiple applications can revolve around enhancing the DDIs listed here by enriching the drugs and events with properties available in the Drugbank, RxNorm and MeSH linked graphs that are not easily accessible via traditional means. This might lead to even further linkage of other resources tying all things back to a described DDI and helping to further explain its reason and impact. As more researchers become aware of this resource they should be encouraged to contribute to it by adding their own DDI predictions, thus enhancing the overall availability of DDI predictions sets for scientific comparisons. We will continue to contact authors and map their data sets into LIDDI to have a more diverse and rich set of resources. We also plan at some point incorporate more drugs and more events into the resource as they become available for it to have a wider coverage of the current drug space.

References

1. Avillach, P., Dufour, J.-C., Diallo, G., Salvo, F., Joubert, M., Thiessard, F., Mougín, F., Trifirò, G., Fourrier-Réglat, A., Pariente, A., Fieschi, M.: Design and validation of an automated method to detect known adverse drug reactions in medline: a contribution from the eu-adr project. *Journal of the American Medical Informatics Association* **20**(3), 446–452 (2013)
2. Banda, J.M., Kuhn, T., Shah, N.H., Dumontier, M.: Liddi: Provenance-centered dataset of drug-drug interactions. figshare July 17, 2015. <http://dx.doi.org/10.6084/m9.figshare.1486478>
3. Bushardt, R.L., Massey, E.B., Simpson, T.W., Ariail, J.C., Simpson, K.N.: Polypharmacy: Misleading, but manageable. *Clinical Interventions in Aging*. **3**(2), 383–389 (2008). 18686760[pmid] *Clin Interv Aging*
4. Callahan, A., Cruz-Toledo, J., Ansell, P., Dumontier, M.: Bio2RDF release 2: improved coverage, interoperability and provenance of life science linked data. In: Cimiano, P., Corcho, O., Presutti, V., Hollink, L., Rudolph, S. (eds.) *ESWC 2013. LNCS*, vol. 7882, pp. 200–212. Springer, Heidelberg (2013)
5. Dumontier, M., Baker, C.J., Baran, J., Callahan, A., Chepelev, L.L., Cruz-Toledo, J., Nicholas, R., Rio, D., Duck, G., Furlong, L.I., et al.: The semanticscience integrated ontology (sio) for biomedical research and knowledge discovery. *J. Biomedical Semantics* **5**, 14 (2014)

6. Gottlieb, A., Stein, G.Y., Oron, Y., Ruppin, E., Sharan, R.: Indi: a computational framework for inferring drug interactions and their associated recommendations. *Molecular Systems Biology* **8**, 592–592 (2012). 22806140[pmid] *Mol. Syst. Biol*
7. Groth, P., Gibson, A., Velterop, J.: The anatomy of a nano-publication. *Information Services and Use* **30**(1), 51–56 (2010)
8. Haerian, K., Varn, D., Vaidya, S., Ena, L., Chase, H.S., Friedman, C.: Detection of pharmacovigilance-related adverse events using electronic health records and automated methods. *Clinical pharmacology and therapeutics* **92**(2), 228–234 (2012)
9. Iyer, S.V., Harpaz, R., LePendu, P., Bauer-Mehren, A., Shah, N.H.: Mining clinical text for signals of adverse drug-drug interactions. *J. Am. Med. Inform. Assoc.* **21**(2), 353–362 (2014)
10. T. Kuhn, C. Chichester, M. Krauthammer, and M. Dumontier. Publishing without publishers: a decentralized approach to dissemination, retrieval, and archiving of data. In: *Proceedings of ISWC 2015. Lecture Notes in Computer Science*. Springer (2015)
11. Kuhn, T., Dumontier, M.: Trusty URIs: verifiable, immutable, and permanent digital artifacts for linked data. In: Presutti, V., d'Amato, C., Gandon, F., d'Aquin, M., Staab, S., Tordai, A. (eds.) *ESWC 2014. LNCS*, vol. 8465, pp. 395–410. Springer, Heidelberg (2014)
12. Kuhn, T., Dumontier, M.: Making digital artifacts on the web verifiable and reliable. *IEEE Transactions on Knowledge and Data Engineering* (2015)
13. Lazarou, J., Pomeranz, B.H., Corey, P.N.: Incidence of adverse drug reactions in hospitalized patients: A meta-analysis of prospective studies. *JAMA* **279**(15), 1200–1205 (1998). doi:10.1001/jama.279.15.1200
14. Lebo, T., et al.: PROV-O: The PROV ontology. Recommendation, W3C (2013)
15. Mons, B., van Haagen, H., Chichester, C., den Dunnen, J.T., van Ommen, G., van Mulligen, E., Singh, B., Hooft, R., Roos, M., Hammond, J., et al.: The value of data. *Nature genetics* **43**(4), 281–283 (2011)
16. Tatonetti, N.P., Ye, P.P., Daneshjou, R., Altman, R.B.: Data-driven prediction of drug effects and interactions. *Science Translational Medicine* **4**(125), 125ra31 (2012). doi:10.1126/scitranslmed.3003377
17. Vilar, S., Uriarte, E., Santana, L., Lorberbaum, T., Hripcsak, G., Friedman, C., Tatonetti, N.P.: Similarity-based modeling in large-scale prediction of drug-drug interactions. *Nat. Protocols* **9**(9), 2147–2163 (2014)
18. Wishart, D.S., Knox, C., Guo, A.C., Cheng, D., Shrivastava, S., Tzur, D., Gautam, B., Hassanali, M.: Drugbank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic Acids Research* **36**, D901–D906 (2008). 18048412[pmid] *Nucleic Acids Res*
19. Zhang, L., Zhang, Y., Zhao, P., Huang, S.-M.: Predicting drugdrug interactions: An fda perspective. *The AAPS Journal* **11**(2), 300–306 (2009)
20. Linked Drug-Drug Interactions (LDDI) dataset. Nanopublication index, July 17, 2015. <http://np.inn.ac/RA7SuQ0e661LJdKpt5EOS2DKykf1ht9LFmNaZtFSDMrXg>