MEDINFO 2010
C. Safran et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-588-4-1095

Querying the National Drug File Reference Terminology (NDFRT) to Assign Drugs to Decision Support Categories

Linas Simonaitis[a,b], Gunther Schadow[a,b]

[a] *Regenstrief Institute, Inc., Indianapolis, IN,*
[b] *Indiana University, Indianapolis, IN*

Abstract

Introduction: The accurate categorization of drugs is a prerequisite for decision support rules. The manual process of creating drug classes can be laborious and error-prone. Methods: All 142 drug classes currently used at Regenstrief Institute for drug interaction alerts were extracted. These drug classes were replicated as fully-defined concepts in our local instance of the NDFRT knowledge base. The performance of these two strategies (manual classification vs. NDFRT-based queries) was compared, and the sensitivity and specificity of each was calculated. Results: Compared to existing manual classifications, NDFRT-based queries made a greater number of correct class-drug assignments: 1528 vs. 1266. NDFRT queries have greater sensitivity (74.9% vs. 62.1%) to classify drugs. However, they have less specificity (85.6% vs. 99.8%). Conclusion: The NDFRT knowledge base shows promise for use in an automated strategy to improve the creation and update of drug classes. The chief disadvantage of our NDFRT-based approach was a greater number of false positive assignments due to the inclusion of non-systemic doseforms.

Keywords

National Drug File Reference Terminology, Drug classification, Knowledge bases, Computer-assisted drug therapy

Introduction

Clinical decision support for computerized provider order entry (CPOE) depends on the assignment of drugs to classes used to express generalized medical knowledge. For example, a CPOE system could detect a risk of interaction between Macrolides and Statins (for a patient who is given both). As a prerequisite, the computer must know which of the drugs on the patient's profile belong to these classes.

The Gopher CPOE system was developed by Regenstrief Institute over the past two decades, and has provided decision support to thousands of physicians caring for patients in central Indiana. Regenstrief informaticians have created drug classes (e.g., "HMG-CoA reductase inhibitors"), and have manually populated these with individual drugs (e.g., "Lovastatin" and "Simvastatin"). Such drug classes facilitate writing decision support rules to detect drug-drug interactions or to suggest laboratory monitoring.

As the years pass, old medications are retired, and new ones enter the market. We have found that the manually-created Gopher drug classes accumulate older medications, but lack newer ones. Therefore, we are searching for a solution – an automated strategy – to improve and update the Gopher drug classes.

The Food and Drug Administration (FDA) is driving the Structured Product Labeling (SPL) initiative. SPL is an HL7 version 3 standard for drug knowledge representation based on the HL7 Reference Information Model (RIM). The FDA has already published over 5000 drug labels in this format. As of June 2009, all drug manufacturers must use the SPL format to register all of their new products with the FDA. An "indexing initiative" is underway, which will annotate these products using the Veteran Administration's National Drug File Reference Terminology (NDFRT). [1]

Already today, the NDFRT contains knowledge annotations for a large number of drugs, and NDFRT drug concepts are linked to NDC, RxNorm and other terminologies in the Unified Medical Language System (UMLS). Can NDFRT knowledge improve the way that drugs are categorized for clinical decision support? Carter and Brown have delivered an initial analysis to answer whether the drug classes used in real systems might be encoded using the available NDFRT categories. [2] We have previously reported on the use of this knowledge to organize basic drug terminology, [3] improve detection of drug intolerances, [4] and to create links between drugs and a patient's problem list. [5] In this paper, we investigate if another area of decision support – the detection of drug interactions – can be supported by NDFRT knowledge content.

Principles

NDFRT is an ontology of medication-related concepts that uses a highly restricted description logic formalism to define drugs in the form:

$$D \sqsubseteq B_1 \sqcap \ldots \sqcap B_n \sqcap \exists\ R_1.\ C_1 \sqcap \ldots \sqcap \exists\ R_m.\ C_m,$$

i.e. a drug concept D is described as the conjunction of base classes (a D is a B_1 and a D is a B_n) and existentially quantified role restrictions.

Some of the important roles in the NDFRT include the following: (a) has ingredient, (b) has mechanism of action, (c) has physiologic effect, (d) may treat, (e) may prevent, or (f) is contraindicated with a disease. Using such roles, we can define Statins as "drugs that have some mechanism of action which is a hydroxylmethylglutaryl-CoA reductase inhibitor"; or Tricyclic antidepressants as "drugs that have some ingredient which is a tricyclic ring structure derivative, and which may treat depression".

The NDFRT does not include fully-defined concepts to represent drug classes. (A fully-defined concept is one defined so that the constraints of the definition are sufficient criteria to declare a concept subsumed.) Instead, the NDFRT distribution file marks all concepts as primitive, where the criteria are descriptive (and necessarily true) for all subsumed concepts, but are not sufficient. In this paper, we attempt to show how we can describe conventional drug classes as fully-defined concepts in the NDFRT ontology.

Methods

NDFRT Knowledge Adapted to a Relational Database

The NDFRT knowledge base is made available for public use in a proprietary description logic XML format. [6] We downloaded the 2008.11.11 version of this file, and applied an XSLT transform to load it into our relational database.

As a matter of routine, we reason with such proprietary file formats in a relational database schema, which we have repeatedly described elsewhere. [3,4,7] In this schema, all NDFRT concept relationships are represented in one table with a relationship type, source concept id, and target concept id (see Figure 1). If the relationship type is transitive and reflexive (like the "is_a" relationship), then we compute the materialized transitive and reflexive closure and distance metric:

- reflexive: distance = 0
- direct: distance = 1
- transitive distance = 2,3,4…

We have found that this approach is fast even for large terminologies and instance databases.

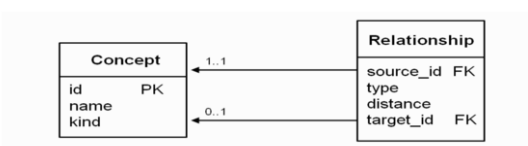

Figure 1 – NDFRT simple relational schema

Our relational database contained 136,054 NDFRT concepts. Of these, there were 120,877 drug concepts: NDC-level packages; drug products (with strength); and abstract "drug preparations" (without strength). There were also 8433 chemicals (derived from the MeSH chemical classification), 1815 physiologic effects (PE), 438 mechanisms of action (MoA), and 4258 diseases. Our relational database contained NDFRT-derived relationships as well: 158,717 is_a relationships, 4720 has_ingredient, 3167 has_PE, 2150 has_MoA, 5670 may_treat, and 793 may_prevent relationships.

Regenstrief Gopher Drug Classes

The Regenstrief Terminology Dictionary stores drug classes relevant to clinicians. Each drug class has a numerical identifier, a name, and a brief description. Physicians can create prescriptions for some of these drugs – in this paper, we will call these "Leaf-Level Drug Classes" (LLDC). Lisinopril is an example. There are also other drug classes, not at the leaf level, and not orderable by physicians. ACE Inhibitor is an example.

The Regenstrief Gopher CPOE system uses 808 non-leaf-level drug classes for various decision support purposes. Each such drug class has been assigned – manually – a set of LLDC identifiers representing orderable drugs. In order to focus our analysis on patient safety, we extracted only those 142 drug classes used by Gopher decision support to identify drug interactions. These drug interactions have been compiled over many years by Regenstrief pharmacists. Although there are several drug knowledge bases in the United States which list drug interactions, there is no single source recognized as a standard.

Defining Drug Classes as NDFRT Queries

The 142 Regenstrief drug classes were encoded using a few description logic templates (see Figure 2) and implemented as relational database queries (written in SQL). The authors, both physician-informaticians, captured the definition of each of the 142 drug classes taking into consideration: (a) the name of the class, (b) the LLDC members of the class, and if in doubt (c) the purpose of the class as it is used in the interaction rules. For example, "Macrolides" and "HMG-CoA reductase inhibitors" are understood from the name. For "Antipsychotics", the name was not sufficient, and the LLDC members of the class had to be examined to see that "Neuroleptics" are meant – excluding Antidepressants. For "Non-sedating antihistamines", the use of the interaction rule had to be examined to see that only those Antihistamines with the risk of QT interval prolongation (e.g., Terfenadine) were intended.

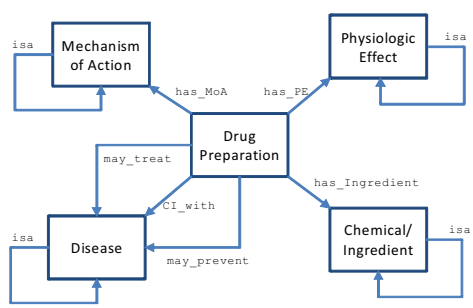

Figure 2 – Some of the Concepts and Relations in the NDFRT. The ones shown were used in this study to define drug classes.

Most (137 of 142) class definitions could be written as one of the following two schemata:

1. For 133 cases: $D \sqsubseteq \exists\ R.\ C$

Role R indicated these relationships: has_ingredient (94 times), has_MoA (21 times), has_PE (7 times), may_treat (7 times), may_prevent (3 times), and contraindicated with (1 time). Oxidizing meds were defined as those contraindicated with glucose-6-phosphate dehydrogenase insufficiency.

2. For 4 cases: $D \sqsubseteq \exists R_1.\, C_1 \sqcap \exists R_2.\, C_2$

In two of these cases, the roles R_1 and R_2 indicated the "has_ingredient" relationship (Cotrimoxazole and Advicor had been defined as sets, and each represents two ingredients). In the two other cases, one of the roles indicated "has_ingredient" and the other role indicated "may_treat" (e.g., Azole antifungals and Tricyclic antidepressants)

These NDFRT-based class definitions can be implemented as SQL queries. For the definition $D \sqsubseteq \exists R.\, C$, the role R is implemented as a database table linked to two other tables: the source concept of the role, and the target concept of the role. We join the source concept via a transitive and reflexive "is_a" relationship to the drug concept D. We join the target concept via another transitive and reflexive "is_a" relationship to the concept C. For the definition $D \sqsubseteq \exists R_1.\, C_1 \sqcap \exists R_2.\, C_2$, the conjunction is implemented as an SQL intersection.

Finally, 5 of the 142 drug classes could not be defined based on the knowledge in the NDFRT ("CYP3A4 inhibitors", "Class Ia antiarrhythmics", "Class III antiarrhythmics", "Non-sedating antihistamines", and "Kayexelates"). The minor effect (e.g., increased repolarization time for "class III antiarrhythmics") was not annotated in the NDFRT knowledge base. In order not to introduce a bias into the comparison of the manual approach with the NDFRT-based approach, these classes were not excluded from analysis, but were left to count against the NDFRT method.

Mapping NDFRT Drug Preparations to Regenstrief LLDC

In order to compare the NDFRT-based drug classes with the manually-created Regenstrief drug classes, we needed an element common to both. NDFRT-based drug classes consist of NDFRT Drug Preparations. Regenstrief drug classes consist of Regenstrief LLDC. We needed to link the two terminologies together. (See Figure 3 for a summary of the links required.)

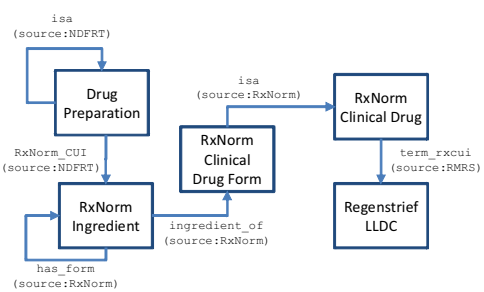

Figure 3 – Relations linking NDFRT Drug Preparations to Regenstrief LLDC

Comparing NDFRT Classes with Regenstrief Drug Classes

The analysis of the comparison is based on a full outer join table that connects NDFRT class members with Regenstrief class members, if either exists. The unit for comparison was the Regenstrief LLDC. In other words, the definition of Regenstrief classes as sets of Regenstrief LLDCs was not disturbed. However, the definition of NDFRT classes was expressed as a set of Regenstrief LLDCs, relying on the mapping linkages described above.

For each of the 142 drug classes: we examined the Regenstrief LLDCs assigned to that class during the manual creation of the class; and we examined the Regenstrief LLDCs assigned to that class by the automated NDFRT-based queries. If a LLDC was assigned to the class by both strategies (manual and NDFRT-based), it was declared to be correctly assigned. If a LLDC was assigned by only one strategy, and not the other, then it required review. A reason for the discrepancy was determined. This review established a consensus set against which both methods were measured.

Results

Over the years, Regenstrief knowledge engineers had manually populated the 142 classes with a combined total of 1271 LLDCs. Our NDFRT-based definitions subsume a combined total of 1905 LLDCs. The two strategies overlap for 754 LLDCs. (See Figure 4.)

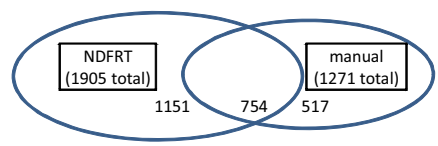

Figure 4 – Overlap of LLDCs collected by the two strategies

The combined totals are disproportionately affected by classes with many LLDCs; therefore, we broke down the counts by class. NDFRT definitions subsume more drug LLDCs in 64 cases; Regenstrief classes subsume more in 54 cases; and the two strategies collect equal numbers in 24 cases.

In order to study performance metrics, we needed to assign some "gold standard" of true class membership. We declared that the 754 LLDCs returned by both the NDFRT-based queries and the Regenstrief classes were true members of their classes. We reviewed the 517 LLDCs in the Regenstrief classes (which had not been included in NDFRT queries), and could only find 5 cases of inappropriate class membership. For example, Digoxin Fab antibody fragments do not share the proarrhythmic effects of the other cardiac glycosides, and thus were considered inappropriately placed in that class. Finally, we reviewed the 1151 in the NDFRT queries (which had not been manually included in the Regenstrief classes). We designated 774 as truly members of the class – in accordance with the original Gopher definition of the drug class. As discussed below, making this designation was a matter of judgment. In sum, we found 2040 class-member assignments (754 + 512 + 774 = 2040) which we considered true.

Having made this designation of true class membership, we calculated sensitivity and specificity. (See Figure 5.) Note that the 2616 LLDCs lacking true class membership are calculated as the sum of the 5 LLDCs inappropriately placed in Gopher classes, the 377 LLDCs inappropriately placed in the NDFRT-based classes, and the 2234 remaining Regenstrief Dictionary Drug LLDCs which do not participate in either strategy.

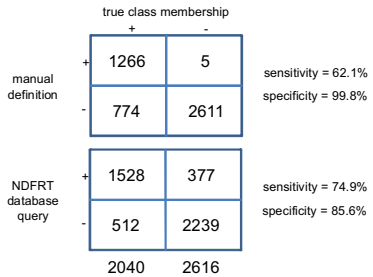

Figure 5 – Performance metrics of the two strategies for defining classes: manual (top) and NDFRT-based (bottom)

The low specificity calculated for the NDFRT strategy derives from 377 false positives. We discovered that 44 of these were due to a medication of incorrect formulation for a drug class where the formulation was explicitly stated. For example, topical Erythromycin gel had been assigned to the class of "Macrolides Systemic". An additional 192 were due to incorrect formulation, where the formulation was implied. For example, Atropine ophthalmic drops had been assigned to the class of "Antiarrhythmics". The systemic route is implied, though not explicitly stated.

The sensitivity calculated for the NDFRT strategy depends on the number of NDFRT Preparations successfully mapped to a Regenstrief LLDC. But 387 of the LLDCs in the manually defined classes are not mapped to RxNorm clinical drugs, and thus cannot be linked to NDFRT Preparations. The greater part of these unmapped LLDCs are outdated medications which are no longer marketed (e.g., Oxytriphylline). These unmapped LLDCs contribute to the count for the manual strategy, but cannot contribute to the count for the NDFRT strategy. Simply excluding the unmapped LLDCs would improve sensitivity of the NDF-RT approach to (1528/1653 =) 92.4%.

The sensitivity of the manually-defined classes is greatly affected by the failure to include combination products. For example, among the medications detected by NDFRT queries, but not included in manually-defined classes, were 229 combinations of ingredients. For example, Fiorinal had been correctly included in the Gopher Barbiturates class; but it had not been included in the Gopher Aspirins class.

Discussion

In this study, knowledge derived from the NDFRT was used to categorize medications and enable a specific type of decision support: checking for drug interactions. However, the accurate categorization of drugs has broader applicability, and can enhance other types of decision support, such as treatment guidelines (e.g., Beta blockers for myocardial infarct).

We demonstrated that classes defined in the NDFRT ontology subsume a greater number of medications than the manually defined classes currently used by our institution. The manual classes are incomplete, and lack many medications which could rightly be assigned to them. This fact is not surprising. We know how difficult it is for a limited number of knowledge engineers to monitor the unrelenting arrival of new drugs on the market, and to manually update drug classes and other features of a complex decision support system. This is especially problematic for combinations of ingredients (e.g., Amlodipine/atorvastatin should be assigned to two classes: Statins and Calcium channel blockers). If we could harness an automatic process of class assignment based on NDFRT definitions, we could improve the completeness of our Gopher classes.

We measured a modest improvement (from 62% to 75%) in sensitivity with the use of NDFRT. But these numbers understate the potential benefit of the NDFRT strategy. Our analyses used Regenstrief LLDCs as the unit of measure. This had no negative impact on the Gopher classes. However, this put the NDFRT queries at a disadvantage. Not all of the NDFRT Preparations in the queries were successfully mapped to Regenstrief LLDCs. Despite the handicap of incomplete mapping, the NDFRT strategy outperformed the manual strategy.

Nevertheless, it would be premature to use NDFRT definitions "as is". This strategy could introduce a lack of specificity. This problem is especially important in the domain of drug interaction decision support. Physician users consider most drug interaction reminders unhelpful; in one study, they overrode 89.4% of such alerts. [8]

We determined that the majority of false positives were due to the wrong formulation of the right ingredient. Thus medications intended for non-systemic (e.g., topical or ophthalmic) use were placed with medications intended for systemic use. As noted by Carter et al, a reference hierarchy of formulated routes would greatly improve the NDFRT. [2] Another promising approach is the definition of drug classes based on the SPL model, which includes dose forms and routes, and will include NDFRT annotations. Such an approach would routinely take form and route into consideration. [3]

An important lesson learned is that to correctly assign drugs to a class, one must understand the purpose of the drug class. A drug class assembled for the purpose of allergy detection may include more members (be more sensitive at a cost of specificity). For example, Penicillamine might be included in a set of warnings for Penicillin allergy. A drug class assembled for treatment guidelines should have less members (be more specific at a cost of sensitivity). For example, not all Quinolones should be suggested for treatment of pneumonia. A drug class to monitor interactions might include only those formulations known to produce sufficient systemic concentrations. Unfortunately, at our institution, a drug class, once it is defined, is often reused for several different purposes – and may not suit them all.

The expressiveness of NDFRT has already been investigated: Rosenbloom et al determined that the Physiological Effect hierarchy is adequate for representing the effects of commonly prescribed medications. [9] Nevertheless, not all of the possible physiological effects of a medication have been instantiated as relationships. We found that the primary treatment effects were well represented; but not all the possible side effects were. For example, the class of anticholinergic medications collected by the NDFRT query did not include some medications (e.g., Diphenhydramine) which geriatricians at our institution have flagged for anticholinergic side effects.

One limitation of our study is the use of Regenstrief LLDCs as the unit of measure: other institutions do not use Regenstrief terminology. However, it must be noted that Regenstrief LLDCs have survived real-world testing over several decades. Furthermore, many Regenstrief LLDCs have a direct correspondence to RxNorm identifiers. Finally, we needed to use Regenstrief LLDCs to put the NDFRT strategy to a real-world challenge. The fact that the NDFRT strategy still outperformed in sensitivity adds credibility to our belief that this strategy should replace the manual maintenance of drug classes.

Another limitation is that these performance metrics depend on our own designation of which drug LLDCs are correctly assigned to a class. In many cases, the authors reached consensus promptly. However, in some cases, careful judgment was required. For example, an NDFRT relation states that Warfarin "may treat" Atrial Fibrillation. But Atrial Fibrillation is an Arrhythmia. Thus the "may treat" relation places Warfarin in the class of Antiarrhythmics. In our judgment, this was a false positive.

Categorization of medications is an important feature of commercial drug knowledge bases. Our experience with some of these products has been favorable. The advantage of the NDFRT is that it has been made freely available in the public domain and that it is part of the SPL system of drug product descriptions published by the pharmaceutical industry through the FDA. We believe there is future potential for decision support that uses a knowledge base maintained by these authoritative sources. An additional advantage of using our method to manipulate a medication terminology is the high degree of expressiveness allowing us to adjust the definition (especially the granularity) of a drug class provided the drugs are annotated with the NDFRT concepts at sufficient detail.

Conclusion

The manual process of creating and updating drug classifications could be automated by a strategy based on the NDFRT, with some human oversight. The level of detail and consistency of assignment of the NDFRT relationships should be improved to allow even more precise definition of classes. Knowledge engineers who write decision support rules could then be more specific about the drug classes they require for the purpose at hand.

An important improvement that our automated method still requires is the recognition of form and route. We believe that the most promising strategy to implement this may derive from the SPL indexing initiative, which will link detailed product descriptions with NDFRT role relationships.

Acknowledgments

This work was performed at the Regenstrief Institute and was funded in part by the Agency for Healthcare Research and Quality (AHRQ) grant R01 HS15377 and the Food and Drug Administration.

References

[1] US Food and Drug Administration. Guidance for Industry - Indexing Structured Product Labeling. Revision 1. June 2008. Available from: http://www.fda.gov/ForIndustry/DataStandards/StructuredProductLabeling/default.htm

[2] Carter JS, Brown SH, Bauer BA, Elkin PL, Erlbaum MS, Froehling DA, Lincoln MJ, Rosenbloom ST, Wahner-Roedler DL, Tuttle MS. Categorical information in pharmaceutical terminologies. AMIA Annu Symp Proc. 2006:116-20.

[3] Schadow G. Assessing the impact of HL7/FDA Structured Product Label (SPL) content for medication knowledge management. AMIA Annu Symp Proc. 2007:646-50.

[4] Schadow G. Structured product labeling improves detection of drug-intolerance issues. J Am Med Inform Assoc. 2009 Mar-Apr;16(2):211-9.

[5] Burton MM, Schadow G, Simonaitis L. Medication and indication linkage: a practical therapy for the problem list? AMIA Annu Symp Proc. 2008 Nov 6:86-90.

[6] National Cancer Institute website. NDF-RT Public Edition version 2008.11.11. Available for FTP download from: ftp://ftp1.nci.nih.gov/pub/cacore/EVS/NDF-RT

[7] Schadow G, Barnes MR, McDonald CJ. Representing and querying conceptual graphs with relational database management systems is possible. Proc AMIA Symp. 2001:598-602.

[8] Weingart SN, Toth M, Sands DZ, Aronson MD, Davis RB, Phillips RS. Physicians' decisions to override computerized drug alerts in primary care. Arch Intern Med. 2003 Nov 24;163(21):2625-31.

[9] Rosenbloom ST, Awad J, Speroff T, Elkin PL, Rothman R, Spickard A 3rd, Peterson J, Bauer BA, Wahner-Roedler DL, Lee M, Gregg WM, Johnson KB, Jirjis J, Erlbaum MS, Carter JS, Lincoln MJ, Brown SH. Adequacy of representation of the National Drug File Reference Terminology Physiologic Effects reference hierarchy for commonly prescribed medications. AMIA Annu Symp Proc. 2003:569-78.

Address for correspondence

Linas Simonaitis, MD, Regenstrief Institute, 410 West 10th St., Suite 2000, Indianapolis, IN 46202. EMail: lsimonaitis@regenstrief.org