Enabling Trust in Clinical Decision Support Recommendations through Semantics

Oshani Seneviratne[1], Amar K. Das[2], Shruthi Chari[1], Nkechinyere N. Agu[1], Sabbir M. Rashid[1], Ching-Hua Chen[2], James P. McCusker[1], James A. Hendler[1], and Deborah L. McGuinness[1]

[1] Rensselaer Polytechnic Institute, Troy, NY 12180, USA
[2] IBM Research, USA

Abstract. In an ideal world, the evidence presented in a clinical guideline would cover all aspects of patient care and would apply to all types of patients; however, in practice, this rarely is the case. Existing medical decision support systems are often simplistic, rule-based, and not easily-adaptable to changing literature or medical guidelines. We are exploring ways that we can enable clinical decision support systems with Semantic Web technologies that have the potential to support representation and linking to details in the related items in the scientific literature, and that can quickly adapt to changing information from the guidelines. In this paper, we present the ontologies and our semantic web-based tools aimed at trustworthy clinical decision support in three distinct areas: guideline representation and reasoning, guideline provenance, and study cohort modeling.

Keywords: Health Data Management · Knowledge Representation · Guideline Modeling · Data Integration · Reasoning

1 Introduction

Clinical Practice Guidelines (CPGs) consist of diagnostic and therapeutic recommendations for specific health conditions and are commonly considered a community standard for patient management. The source of these recommendations often comes from published medical literature, which is rigorously reviewed and synthesized. Domain experts with clinical authority will then develop a plan of care, integrating these recommendations into diagnostic or treatment steps, pathways, and algorithms. Since the evidence changes over time, given new tests, interventions and studies, the plans within CPGs are regularly updated. Provider acceptance and use of CPG recommendations depend on several factors: (i) A provider must view the recommendation as relevant to his or her patient's clinical situation. (ii) A provider must accept that the published study or studies supporting the recommendation are rigorous. (iii) A provider must understand that the study population is similar to his or her patient or patient population.

CPGs are initially published in text form and are later translated by domain experts and IT specialists into computer-based clinical decision support rules

that can be embedded within Electronic Health Record (EHR) Systems. Through this dissemination process, however, the clinical relevance, study provenance, and evidence specificity of a CPG recommendation are often not captured or conveyed when a rule is triggered. Providers, as a result, may be less inclined to trust new recommendations that are surfaced without an understanding of their source or applicability [22,21,3]. Furthermore, a list of factors that impact the trustworthiness of a guideline can be found in [13].

The ontologies, rules, and the special-purpose reasoners we are developing can be used within a clinical decision support system to address the challenge of having a provider trust when to use a recommendation in a clinical decision support system, based on how that technology ensures transparency, explainability, and specificity. We focus on three technical aspects enabled through semantics: (i) the plausible reasons for an observed intervention, (ii) connecting a study to a guideline recommendation, and (iii) the people studied in the studies that support a guideline.

In this paper, we focus on the semantic modeling of recommendations from pharmacological treatment guidelines published by the American Diabetes Association[3]. Our work strives to answer the following key questions:

1. Can we represent guideline provenance in a way that enables the tracking of revisions in guidelines that leads to a better understanding of the evolution of the guideline as new medical evidence comes to light?
2. Can we represent study cohorts reported in the medical literature in a way that enables effective querying to pinpoint studies that may be applicable for a patient?
3. Can we understand which guideline recommendation is applicable in the context of discrepancies between past interventions applied to an individual patient and what the guideline would have recommended at those decision points?

We organize our paper as follows: In section 2, we describe how we have represented the provenance of clinical guidelines. In section 3, we describe how we represent the study cohorts reported in the medical literature that is referenced by the ADA CPGs. In section 4, we describe how abductive reasoning can be used to explain discrepancies between the interventions noted in the patient's EHR relevant guideline recommendations. Finally, we present related work in section 5.

2 Understanding the Provenance of Guideline Recommendations

Through the translation of text-based CPG recommendations into clinical decision-support rules, the source of a rule is often not made evident to a provider. As

[3] American Diabetes Association Standards of Care
`https://care.diabetesjournals.org/content/42/Supplement_1`

a result, the provider may not know the rationale and applicability of that recommendation to his or her patient. We address this challenge in our work on guideline provenance.

G-Prov Ontology

We developed an ontology framework called G-Prov [28] that captures provenance metadata at different granular levels. The ontology can be used to annotate rules with citation-backed evidence sentences and other sources of knowledge such as tables and figures. We demonstrate the reasoning capability of G-Prov for three different clinical questions. (i) Where does this treatment suggestion come from? (ii) Which studies support the recommendation? (iii) How recent is this recommendation?

G-Prov provides the physician with information about the rule that fired to generate the suggestion. We started our work using the ADA guidelines and annotation of the Semantic Web Rule Language (SWRL) rules from the Diabetes Mellitus Treatment Ontology (DMTO) [8]. The ontology also includes information from other guidelines, thus providing the physician with multiple sources/evidence for the suggested treatment.

For more information on G-Prov, please refer to `https://tetherless-world.github.io/GProv`.

3 Modeling Clinical Research Study Cohorts

Recommendations within CPGs are derived from the results of clinical trials and other types of clinical research studies. These studies are based on a recruited cohort of subjects, which are often not reflective of a provider's clinical population because of the inclusion and exclusion criteria used for the study as well as sites of recruitment. As a result, a provider will be interested in knowing which studies are the basis of a particular recommendation and whether his or her patient or patient population is similar to the study cohort(s). However, achieving these goals is no simple task because the published description about a cohort varies significantly across studies. Furthermore, a single patient may differ with multiple attributes from those in the study cohort, and comparing similarity or dissimilarity across these dimensions can be challenging. Therefore, we need a robust representation of studies and cohorts and semantic technology to evaluate and visualize cohort similarity.

Study Cohort Ontology (SCO)

The SCO developed by our team [4] defines classes and properties to describe content related to demographics, interventions, cohort statistics for each variable of a study cohort, as well as study's inclusion/exclusion criteria. As a proof-of-concept, we modeled eight cited research studies in the pharmacologic recommendations chapter in the ADA guidelines. In our modeling effort, we leverage

best practice ontologies, including: (i) Disease Ontology (DOID) (ii) Clinical Measurement Ontology (CMO) (iii) Unit Ontology (UO) (iv) Phenotypic Quality Ontology (PATO) (v) Semantic Science Integrated Ontology(SIO) We will continue to support interlinking and expansion as needed.

Representing aggregations in OWL and RDF has been a long-studied research problem, and there are multiple approaches to the modeling of aggregations. Having analyzed patterns across several cohort summary tables, we see that aggregations are manifested as descriptive statistics. The descriptive statistics are often measures of central tendency or dispersion like Mean, Median, Mode, Standard Deviation, Interquartile Range, and Rate. In the SCO we built a view inspired by the upper-level Ontology SemanticScience Integrated Ontology (SIO) [7] to model the descriptive statistics of characteristics on a set of patients administered a medical intervention and studied for an outcome in a research study. Since the terminology across the research study descriptions varies, we have begun to incorporate a medical meta-thesaurus such as UMLS and MeSH Check Tags for alignment.

We have converted several ADA guidelines to a Computer Interpretable Guideline (CIG) JSON format and extracted cited research studies. Further, we are automating the extraction of population descriptions from these studies using a PDF table extractor tool developed by IBM [27]. Additionally, for Medline citations, we have extracted additional study metadata from PubMed, and also plan on incorporating mappings to MeSH terms.

Applications may use SCO to support analyses that require a deep understanding of study populations. SCO annotated knowledge graphs support visualizations that provide the capability to view the fit of a patient as a whole with the various treatment arms to help the physician ascertain study applicability. The cohort similarity visualizations are powered by results of SPARQL queries, targeted to population descriptions stored in the study knowledge graphs built on SCO. Physicians could choose a subset of characteristics they wish to view as a part of the deep dive, or our visualizations could build off characteristics common to both patients and patient groups studied within the study. Furthermore, cohort similarity scores learned through *Semantically targeted Analysis* [18] techniques can assist in decision-making capabilities when picking the most relevant studies (in situations where more than one study is applicable for a patient).

For more information on SCO, please refer to `https://tetherless-world.github.io/study-cohort-ontology`.

4 Understanding Rationales of Past Treatment Options

Even though CPGs enable evidence-based clinical decision making using the technologies outlined in Sections 2 and 3, observed clinical actions could deviate from recommended actions. Several factors may explain this discrepancy, such as no clear recommendation for a particular clinical decision, a contraindi-

cation that the patient has to a recommendation, and patient or provider disagreement with the recommendation [22,21,3]. When such a discrepancy occurs, understanding which guideline recommendation applies based on the observed interventions can be challenging. We are developing a semantically-enabled abductive reasoning component to assist in discovering the potential rationales of past treatment options.

4.1 Type 2 Diabetes Example

For the recommendation *"Metformin, if not contraindicated and if tolerated, is the preferred initial pharmacologic agent for the treatment of type 2 diabetes,"* we can construct the following rule given in Listing 1.1 for deductive inference to learn new facts about the patients.

```
{
  ?patient rdf:type sio:Human ;
    sio:hasRole sio:PatientRole ;
    sio:isParticipantIn
      [ rdf:type sio:Diagnosis ;
        sio:hasValue "Diabetes" ] ,
      [ rdf:type efo:0003785 ;
        sio:hasValue ?allergy ] .
  NOT EXIST
    {dmto:Metformin gl:hasContraindication ?allergy . }
}
=>
{
?patient sio:isParticipantIn
    [ rdf:type dmto:TreatmentPlan ;
      sio:hasPart [ rdf:type ncit:C28180 ;
                    sio:hasAgent dmto:Metformin ] ]
}
```

Listing 1.1. Diabetes Treatment Rule Example

The rule given in Listing 1.1 indicates that only the patients that do not have an allergy, indicated by the concept *efo:0003785*[4] from the Experimental Factor Ontology (EFO), to the drug Metformin, indicated by the concept *dmto:Metformin* from the Diabetes Mellitus Treatment Ontology (DMTO), should be administered Metformin in their prescription, identified by the concept *ncit:C28180*[5], from the National Cancer Institute Thesaurus (NCIt). However, this rule may not be a perfect match to the patient's record. Reasons for such partial matches could vary including: (i) Substitution of medication is acceptable, but the guideline does not specify which medication (e.g., the physician did not use Metformin for initial monotherapy), (ii) Patient has a contraindication to the drug (e.g., the

[4] efo:0003785 ⇒ `http://www.ebi.ac.uk/efo/EFO_0003785`
[5] ncit:C28180 ⇒ `http://purl.obolibrary.org/obo/NCIT_C28180`

patient has a comorbidity that is potentially worsened by the drug), (iii) Insufficient treatment period for drug regimen (e.g., treatment given for more than three months), (iv) Treatment intensification does not match guideline (e.g., the patient is switched from one dual therapy to another dual therapy without going to triple therapy),.

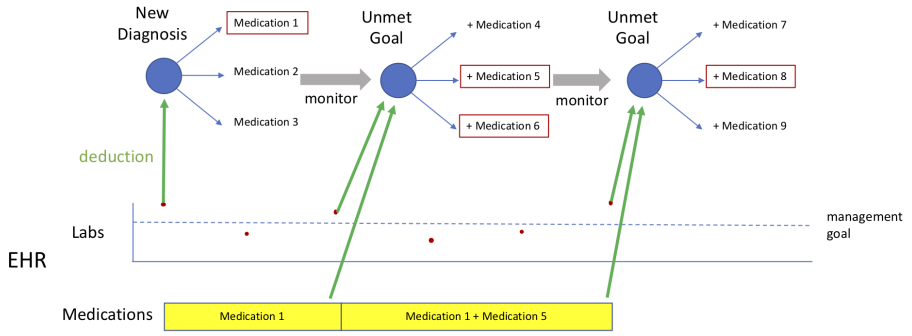

Fig. 1. Traditional decision support for an existing patient

To illustrate this further, consider the scenario outlined in Fig 1, where the data shows an initial' treatment decision that does not conform to the guideline. But not all past actions can be used to classify the degree of conformance to guideline recommendations. The physician may also need to examine prescribed drug(s) and match to recommended drug or drug combinations in the guidelines. She may also need to determine if the antecedents of a previously matched treatment recommendation still hold, and also determine what possible recommendation(s) would follow any part of the observed treatment that is conforming to the guideline. Therefore, we are investigating a new approach to applying guidelines to the patient's EHR data.

There are many decision points a physician must consider for a given pharmacotherapy for a specific disease. A domain expert can model the decision path for the treatment as specified in an official guideline for a patient. The reasoner should classify past actions to the degree of conformance to guideline recommendations, and examine prescribed drug(s) and match to recommended drug or drug combinations in the guidelines. Then, finally, it can determine if the antecedents of the matched treatment recommendation hold, and what possible recommendation(s) would follow any part of the observed treatment that is guideline conforming.

4.2 Abductive Reasoning

The goals for the reasoning process are as follows. First, we wish to learn possible explanations for why a physician decides whether or not to follow a guideline

recommendation. Second, we wish to determine explanations for adverse events associated with patients, i.e., whether an allergy (a protein-protein interaction, or a protein-drug interaction) caused it.

Example Reasoning Process Outline for Type 2 Diabetes Pharmacotherapy Consider the scenario mentioned in the rule outlined in Listing 1.1. Our deductive inference activity resulted in the Metformin prescription for a patient. However, we instead find that the physician had administered treatment differing from the CPG recommendation in the EHR data. With a semantic representation for Metformin that includes its contraindicators, we can abductively infer that a potential reason for the physician's action was that the *patient had a creatine clearance of 58 mL/min, which is below the threshold of 60 mL/min for recommended use of Metformin*. Intuitively, we can achieve this abductive goal by checking if knowledge about the patient corresponds to the 'consequent' of any rule in our ontology, by checking if the patient facts match the 'antecedent' of any matching rule. Such an approach would be to allow for incomplete matches of the patient's facts to the antecedent of a matching rule. For example, if the patient's facts include 3/4 triples in the antecedent of a given rule, we still consider that rule as a possible explanation.

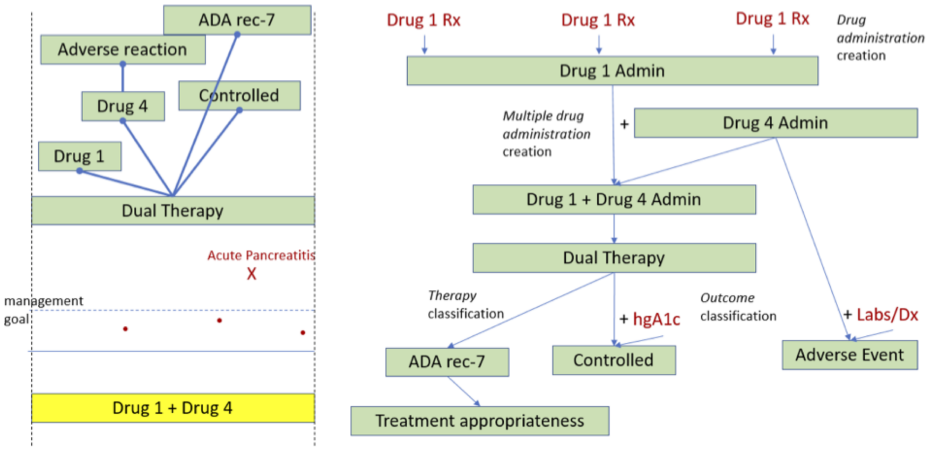

Fig. 2. Proposed inferential process to identify the best matching CPG recommendation for a physician's action.

As can be seen in Fig. 2, we attempt to infer the intent of clinical actions when the physician intensified treatment, or when the treatment changes are not based on guideline recommendations, or the physician stopped medication in response to an adverse event. In this scenario, we model the ADA guidelines for pharmacotherapy by extracting the guideline recommendations from the ADA website

and representing the content as RDF that preserves provenance information. Additionally, we create a representation model for the patient in terms of their attributes, diagnoses, and prescribed treatment plans. Given an observed patient treatment history in the EHR, which consists of a sequential series of prescribed medications, laboratory test results, and comorbid health conditions, the goal of our reasoning method is to identify the best matching CPG recommendation for each action. If an action differs from the CPG recommendation, our reasoner infers the clinical intent for the action. Intents include intensifying treatment to meet the therapeutic goal, using a particular medication given a comorbid condition, or providing an alternative treatment due to contraindications.

We plan to evaluate the success of the abductive inference activities by comparing human physician conclusions to that of the system and examining discrepancies between the sets of conclusions.

5 Related Work

Guideline Provenance

Provenance Context Entity (PaCE) [23], a Scalable Provenance Tracking for Scientific RDF Data, creates provenance-aware RDF triples using the provenance context notion. The approach was implemented at the US National Library of Medicine for Biomedical Knowledge Repository. Curcin et al. [5] provide recommendations based on their experience with several EU based biomedical research projects, providing real issues and a high-level recommendation which could be reused across the biomedical domain. Kifor et al. [15] investigated provenance in a decision support system that attempts to trace all the systems execution steps to explain, on the patient level, the final results generated by the system. There is also extensive research in provenance in distributed healthcare management such as as the work by Deora et al. [6] that aims to ensure efficient healthcare data exchange. The work by Alvarez et al [2] and Xu et al [35] also goes further into provenance for effective healthcare data management. Guideline-based decision support systems such as the above work aim to assist healthcare practitioners with patient diagnosis and treatment choices. However, our G-Prov is different from these in that we embody the information present in published CPGs encoded into computable knowledge, such as rules as well as the evidence sentences from the CPG directly.

Modeling Study Cohorts

The Ontology of Clinical Research (OCRe) [26] is a widely cited study design ontology used to model the study lifecycle, and addresses goals similar to our study applicability scenario in SCO. They adopt an Eligibility Rule Grammar and Ontology (ERGO) [32] annotation approach for modeling study eligibility criteria to enable matching of a study's phenotype against patient data. ProvCaRe [33]

integrates OCRe and also supports evaluation with the PICO model. Although their ontological model captures statistical measures, their modeling is not as intuitive and does not seem to leverage the power of OWL math constructs to the fullest. We also found that most clinical trial ontologies, e.g., CTO-NDD [36], are domain-specific and not directly reusable for a population modeling scenario. Other ontologies, such as the EPOCH ontology [25] that was developed to track patients through their clinical trial visits, had class hierarchies that were insufficient to represent the types of publication cited in the ADA guideline.
Clinical trial matching has been attempted multiple times, mainly as a Natural Language Processing problem, including a knowledge representation (KR) approach that utilizes from SNOMED-CT to improve the quality of the cohort selection process for clinical trials [19]. However, the focus of their effort was mainly on efficient KR of patient data, and study eligibility criteria were formulated as SPARQL queries on the patient schema. We tackle the converse problem of identifying studies that apply to a clinical population based on the study populations reported.
Furthermore, Liu et al [17] detail an approach to creating precision cohorts. Their emphasis is on learning a distance metric which best suits the patient population, but they are not providing a quantified score of similarity. Lenert et al [16] have developed a couple of compelling visualizations of cohort similarity to county populations, but the data analysis, knowledge representation, and machine learning methods used are not elaborated well. Study bias is common in scientific research that can be mitigated through appropriate algorithm choice [1]. Specifically, the recent research focus has been on methods that identify and reduce the influence of potential bias attributes such as ethnicity and gender of data, and the biased algorithms in the models by using various bias measurement metrics and bias mitigation algorithms. Our solution is to incorporate semantic technologies that leverage many of the existing biomedical ontologies, builds upon decades of reasoner work that leads to explainable results that many of the machine learning models cannot yet provide.

Guideline Modeling and Reasoning

There have been over 30 years of research in medical informatics on guideline modeling. Many executable guideline models have been created (Proforma [10], EON [31], PRODIGY [14], Asbru [11], GLEE [34], GLARE [29], SAGE [30]), each with different reasoning capabilities, and GLIF [20] was created as an interchange language for guideline knowledge, while OpenClinical.org [9] was an online clearinghouse of models and tools. Ongoing projects include ATHENA (hypertension, pain management, and others) [12] and DESIREE (breast cancer) [24]. In our work, we have learned from these previous efforts but focus primarily on the pain points that physicians identified when using the ADA guidelines. Our work differs from prior guideline modeling work by connecting recommendations in guidelines to the source literature and their study populations. Our knowledge modeling approach thus provides transparency that may

increase healthcare provider trust in following a recommendation. In addition, we use our guideline model to support abductive reasoning, which prior approaches have not supported.

6 Conclusion

In this paper, we highlighted several challenges with building a guideline execution engine within todays EHR systems. In our work, the usage of semantic technologies is not limited to knowledge representation, the reuse, and the contribution to the expanding body of biomedical ontologies. Our work spans guideline provenance (i.e., G-Prov) to represent the evolution of guidelines in a highly evolving medical information space, scientific study cohorts (i.e., SCO), and novel reasoning strategies to address some of the lapses in reasoning systems deployed on EHR systems to guide the physicians to treat their patients better. These tools are tailored to the unique needs of health professionals to give personalized recommendations based on their patients' unique situations.

Because we use standards-based community-accepted vocabularies and practices, we achieve high interoperability in our work. For example, the G-Prov ontology can also provide the SCO information on the citations for the recommendation, and SCO, in turn, will provide information about the patient cohorts used within the study. Therefore, we can trace back the provenance for cited clinical studies and vice versa. Our special-purpose inference engine is being built to identify this association chain and point to the rating level of associated guideline recommendation, as well as for pinpointing the missing information and rules that led to a physician's treatment decision.

Therefore, our ultimate goal is to provide practitioners with evidence-based clinical knowledge that enables transparency when integrating guideline content into CDS systems. We are primarily working with guideline data on diabetes, and the current focus includes both the ADA guideline and the AACE guideline. However, for future purposes, we plan on incorporating guidelines for other diseases as well.

Acknowledgements

This work is partially supported by IBM Research AI through the AI Horizons Network. We thank our colleagues Kristen Bennett, John Erickson, Morgan Foreman and Dan Gruen, who provided insight and expertise that greatly assisted the research.

References

1. Althubaiti, A.: Information bias in health research: definition, pitfalls, and adjustment methods. Journal of multidisciplinary healthcare **9**, 211 (2016)

2. Álvarez, S., Vázquez-Salceda, J., Kifor, T., Varga, L.Z., Willmott, S.: Applying provenance in distributed organ transplant management. In: International Provenance and Annotation Workshop. pp. 28–36. Springer (2006)
3. Bryn Nelson, P.: Why Arent Doctors Following Guidelines? (2016), https://www.the-hospitalist.org/hospitalist/article/121436/why-arent-doctors-following-guidelines
4. Chari, S., Qi, M., Agu, N.N., Seneviratne, O., McCusker, J.P., Bennett, K.P., Das, A.K., , McGuinness, D.L.: Making Study Populations Visible through Knowledge Graphs. International Semantic Web Conference (2019)
5. Curcin, V., Miles, S., Danger, R., Chen, Y., Bache, R., Taweel, A.: Implementing interoperable provenance in biomedical research. Future Generation Computer Systems **34**, 1–16 (2014)
6. Deora, V., Contes, A., Rana, O.F., Rajbhandari, S., Wootten, I., Tamas, K., Varga, L.Z.: Navigating provenance information for distributed healthcare management. In: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence. pp. 859–865. IEEE Computer Society (2006)
7. Dumontier, M., Baker, C.J., Baran, J., Callahan, A., Chepelev, L., Cruz-Toledo, J., Del Rio, N.R., Duck, G., Furlong, L.I., Keath, N., et al.: The Semanticscience Integrated Ontology (SIO) for biomedical research and knowledge discovery. Journal of biomedical semantics **5**(1), 14 (2014)
8. El-Sappagh, S., Kwak, D., Ali, F., Kwak, K.S.: DMTO: a realistic ontology for standard diabetes mellitus treatment. Journal of biomedical semantics **9**(1), 8 (2018)
9. Fox, J., Gutenstein, M., Khan, O., South, M., Thomson, R.: OpenClinical. net: A platform for creating and sharing knowledge and promoting best practice in healthcare. Computers in Industry **66**, 63–72 (2015)
10. Fox, J., Johns, N., Rahmanzadeh, A., Thomson, R.: PROforma: A method and language for specifying clinical guidelines and protocols. Studies in Health Technology and Informatics pp. 516–520 (1996)
11. Fuchsberger, C., Hunter, J., McCue, P.: Testing Asbru guidelines and protocols for neonatal intensive care. In: Conference on Artificial Intelligence in Medicine in Europe. pp. 101–110. Springer (2005)
12. Goldstein, M.K., Hoffman, B.B., Coleman, R.W., Tu, S.W., Shankar, R.D., O'Connor, M., Martins, S., Advani, A., Musen, M.A.: Patient safety in guideline-based decision support for hypertension management: ATHENA DSS. Journal of the American Medical Informatics Association **9**(Supplement_6), S11–S16 (2002)
13. Graham, R., Mancher, M., Wolman, D.M., Greenfield, S., Steinberg, E., et al.: Trustworthy clinical practice guidelines: Challenges and potential. In: Clinical Practice Guidelines We Can Trust. National Academies Press (US) (2011)
14. Johnson, P.D., Tu, S., Booth, N., Sugden, B., Purves, I.N.: Using scenarios in chronic disease management guidelines for primary care. In: Proceedings of the AMIA Symposium. p. 389. American Medical Informatics Association (2000)
15. Kifor, T., Varga, L.Z., Vazquez-Salceda, J., Alvarez, S., Willmott, S., Miles, S., Moreau, L.: Provenance in agent-mediated healthcare systems. IEEE Intelligent Systems **21**(6), 38–46 (2006)
16. Lenert, M.C., Mize, D.E., Walsh, C.G.: X Marks the Spot: Mapping Similarity Between Clinical Trial Cohorts and US Counties. In: AMIA Annual Symposium Proceedings. vol. 2017, p. 1110. American Medical Informatics Association (2017)
17. Liu, H., Li, X., Xie, G., Du, X., Zhang, P., Gu, C., Hu, J.: Precision Cohort Finding with Outcome-Driven Similarity Analytics: A Case Study of Patients with Atrial Fibrillation. In: MedInfo. pp. 491–495 (2017)

18. New, A., Rashid, S.M., Erickson, J.S., McGuinness, D.L., Bennett, K.P.: Semantically-aware population health risk analyses. Machine Learning for Health (ML4H) Workshop at NeurIPS (2018)
19. Patel, C., Cimino, J., Dolby, J., Fokoue, A., Kalyanpur, A., Kershenbaum, A., Ma, L., Schonberg, E., Srinivas, K.: Matching patient records to clinical trials using ontologies. In: The Semantic Web, pp. 816–829. Springer, Busan, Korea (2007)
20. Peleg, M., Boxwala, A.A., Ogunyemi, O., Zeng, Q., Tu, S., Lacson, R., Bernstam, E., Ash, N., Mork, P., Ohno-Machado, L., et al.: GLIF3: the evolution of a guideline representation format. In: Proceedings of the AMIA Symposium. p. 645. American Medical Informatics Association (2000)
21. Peter Edelstein, M.: How Clinical Guidelines Can Fail Both Doctors and Patients (2016), `https://www.elsevier.com/connect/why-arent-all-physicians-using-clinical-practice-guidelines`
22. Robert McNutt, M., Nortin Hadler, M.: How Clinical Guidelines Can Fail Both Doctors and Patients (2014), `https://thehealthcareblog.com/blog/2014/01/30/how-clinical-guidelines-can-fail-both-doctors-and-patients`
23. Sahoo, S.S., Bodenreider, O., Hitzler, P., Sheth, A., Thirunarayan, K.: Provenance Context Entity (PaCE): Scalable provenance tracking for scientific RDF data. In: International Conference on Scientific and Statistical Database Management. pp. 461–470. Springer (2010)
24. Séroussi, B., Guézennec, G., Lamy, J.B., Muro, N., Larburu, N., Sekar, B.D., Prebet, C., Bouaud, J.: Reconciliation of multiple guidelines for decision support: a case study on the multidisciplinary management of breast cancer within the DESIREE project. In: AMIA Annual Symposium Proceedings. vol. 2017, p. 1527. American Medical Informatics Association (2017)
25. Shankar, R.D., Martins, S.B., O'Connor, M.J., Parrish, D.B., Das, A.K.: Epoch: an ontological framework to support clinical trials management. In: Proc. Int. Workshop on Healthcare Inf. and Knowl. Manage. pp. 25–32. ACM, Arlington, Virginia (2006)
26. Sim, I., Tu, S.W., Carini, S., Lehmann, H.P., Pollock, B.H., Peleg, M., Wittkowski, K.M.: The Ontology of Clinical Research (OCRe): an informatics foundation for the science of clinical research. Journal of biomedical informatics **52**, 78–91 (2014)
27. Staar, P.W., Dolfi, M., Auer, C., Bekas, C.: Corpus Conversion Service: A machine learning platform to ingest documents at scale. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 774–782. ACM (2018)
28. Stephen, S., Seneviratne, O., McGuinness, D.L., Chari, S., Das, A.: G-PROV: A Provenance Encoding Structure for Guideline Evidence. Knowledge Representation and Semantics Working Group Pre-Symposium at AMIA 2018 (2018)
29. Terenziani, P., Montani, S., Bottrighi, A., Torchio, M., Molino, G., Correndo, G.: The GLARE approach to clinical guidelines: main features. Studies in health technology and informatics pp. 162–166 (2004)
30. Tu, S.W., Campbell, J.R., Glasgow, J., Nyman, M.A., McClure, R., McClay, J., Parker, C., Hrabak, K.M., Berg, D., Weida, T., et al.: The SAGE Guideline Model: achievements and overview. Journal of the American Medical Informatics Association **14**(5), 589–598 (2007)
31. Tu, S.W., Musen, M.A.: The EON model of intervention protocols and guidelines. In: Proceedings of the AMIA Annual Fall Symposium. p. 587. American Medical Informatics Association (1996)

32. Tu, S.W., Peleg, M., Carini, S., Bobak, M., Ross, J., Rubin, D., Sim, I.: A practical method for transforming free-text eligibility criteria into computable criteria. Journal of biomedical informatics **44**(2), 239–250 (2011)
33. Valdez, J., Kim, M., Rueschman, M., Socrates, V., Redline, S., Sahoo, S.S.: ProvCaRe semantic provenance knowledgebase: evaluating scientific reproducibility of research studies. In: AMIA Annual Symposium Proceedings. vol. 2017, p. 1705. American Medical Informatics Association (2017)
34. Wang, D., Shortliffe, E.H.: GLEE–a model-driven execution system for computer-based implementation of clinical practice guidelines. In: Proceedings of the AMIA Symposium. p. 855. American Medical Informatics Association (2002)
35. Xu, S., Rogers, T., Fairweather, E., Glenn, A., Curran, J., Curcin, V.: Application of data provenance in healthcare analytics software: information visualisation of user activities. AMIA Summits on Translational Science Proceedings **2017**, 263 (2018)
36. Younesi, E.: A Knowledge-based Integrative Modeling Approach for In-Silico Identification of Mechanistic Targets in Neurodegeneration with Focus on Alzheimer's Disease. Ph.D. thesis, Department of Mathematics and Natural Sciences, Universitäts-und Landesbibliothek Bonn, Bonn, Germany (2014)