Mini-review

Clinical and Translational Science

Progress Toward a Universal Biomedical Data Translator

Karamarie Fecho[1*], Anne Thessen[2*], Sergio Baranzini[3], Chris Bizon[4], Jennifer Hadlock[5], Sui Huang[6], Ryan Roper[7], Noel Southall[8], Casey Ta[9], Paul Watkins[10], Mark Williams[11], Hao Xu[12], William Byrd[13], Vlado Dancik[14], Marc Duby[15], Michel Dumontier[16], Gustavo Glusman[17], Nomi Harris[18], Eugene Hinderer[19], Greg Hyde[20], Adam Johs[21], Andrew Su[22], Guangrong Qin[23], Qian Zhu[24], and the Biomedical Data Translator Consortium

[1]Renaissance Computing Institute, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA
[2]Center for Health AI, University of Colorado Anschutz Medical Campus, Aurora, CO 80045

**These authors contributed equally to the manuscript and should be considered co-lead authors.*

+ Consortial authors [Consortium members who are not included in the primary author list]

Abstract: Clinical and translational science has reached an inflection point in the breadth and diversity of available data and the potential impact of such data to improve human health and well-being. Yet, the data are often siloed, disorganized, and not broadly accessible due to discipline-specific differences in terminology and representation. To address these challenges, the Biomedical Data Translator Consortium developed and tested a pilot knowledge graph–based 'Translator' system capable of integrating existing biomedical data sets and 'translating' those data into insights intended to augment human reasoning and accelerate translational science. Having demonstrated feasibility of the Translator system, the Translator program has since moved into development, and the Translator Consortium has made significant progress in the research, design, and implementation of an operational system. Herein, we describe the current system's architecture, performance, and quality of results. We apply Translator to several real-world use cases developed in collaboration with subject-matter experts. Finally, we discuss the scientific and technical features of Translator and compare those features to other state-of-the-art biomedical graph-based question-answering systems.

Introduction

The breadth and diversity of biomedical data available today hold great promise in the application of such data into actionable outcomes aimed at accelerating translational science and ultimately improving human health and well-being. Indeed, advancements in computing and storage capabilities have fostered a wealth of large datasets across clinical and translational

domains. Translational scientists now have unprecedented access to data and knowledge on genes, biological pathways, chemicals, metabolites, drugs, diseases, environmental exposures, clinical healthcare records, and more. Yet, the inherent power of the available data has not been fully harnessed due to long-recognized challenges related to the compartmentalization of data into separate domains, the lack of widely adopted standards or the adoption of standards that are domain-specific, and non-compliance with the principles of findability, accessibility, interoperability, and reusability (FAIR) [1].

The Biomedical Data Translator program ('Translator program') was launched in Fall 2016 by the National Center for Advancing Translational Sciences (NCATS) in an effort to overcome the many challenges that have long hindered translational science. The vision of the Translator program is to augment human reasoning and accelerate scientific discovery "through an informatics platform that enables interrogation of relationships across the full spectrum of data types" [2]. To achieve this goal, NCATS rapidly and adeptly established a diverse community of nearly 200 basic and clinical scientists, informaticians, ontologists, software developers, and practicing clinicians distributed over 11 teams and 28 institutions to form the Biomedical Data Translator Consortium ('Translator Consortium'). The Translator Consortium adheres to several core principles that have allowed the program to make considerable progress toward a shared vision: namely, team science; a bottom-up management approach; and open-source community-contributed software development.

The Translator Consortium last reported on the program in two 2019 publications [3,4]. The aim of this review is to provide an update on the Translator program. We first review approaches for knowledge representation in translational science. We then describe the technical solution that the Translator program has converged on. We demonstrate real-world use-case applications of the prototype Translator system ('Translator'). Finally, we end with a discussion of next steps and a comparison between Translator and similar systems.

Knowledge Representation in Translational Science

'Knowledge' vs 'Data'

The distinction between 'knowledge' and 'data' is most often captured as the data-to-information-to-knowledge-to-wisdom transformation or DIKW pyramid [5]. While the origins of this hierarchical representation model are uncertain, and other knowledge representations exist [6], the DIKW framework has been widely used in fields like information science, communications science, and library science. Within this hierarchical framework, data is viewed as abundant and characterized as discrete objective facts or observations; information is considered to be assertions derived from data and intended to provide interpretation of the data; knowledge is viewed as generally accepted, universal assertions derived from the accumulation of information; and wisdom is considered to be the most abstract layer of understanding derived from assertions and insights into acquired knowledge [7].

Approaches for Knowledge Representation

Application of the conceptual DIKW framework has focused primarily on knowledge discovery, or the systematic process whereby observations or data are organized and interpreted into information that is then scrutinized or tested in the context of existing knowledge, with any subsequent assertions disseminated for peer consensus and adjudication before being accepted as new knowledge. Approaches for knowledge discovery date back to ancient times

and form the foundation of the scientific method [8]. Approaches for knowledge representation likewise date back to ancient times [8]. Early forms of modern peer-reviewed publication represent one approach to knowledge representation that remains in use today.

Knowledge Graphs

In recent years, 'knowledge graphs' (KGs) have become a common approach for knowledge representation in a variety of fields [9,10]. In a KG, entities or data types are represented as nodes and connected to each other by way of edges with predicates that describe the relationship between entities. A 'schema' is used to constrain the KG by specifying how knowledge can be represented; as such, it provides a framework for validating specific instances of knowledge representation through rules that dictate the syntax and semantics. KGs allow users to pose questions that can then be translated into query graphs and applied to identify subgraphs within the KG that match the general structure of the query graph, thereby producing answers to user queries and generating new knowledge [11]. KGs have had successful applications, with Google's KG [10] perhaps the most widely known.

The Translator Solution

The Translator Consortium has adopted a federated KG-based approach for biomedical knowledge representation and discovery (**Figure 1**).

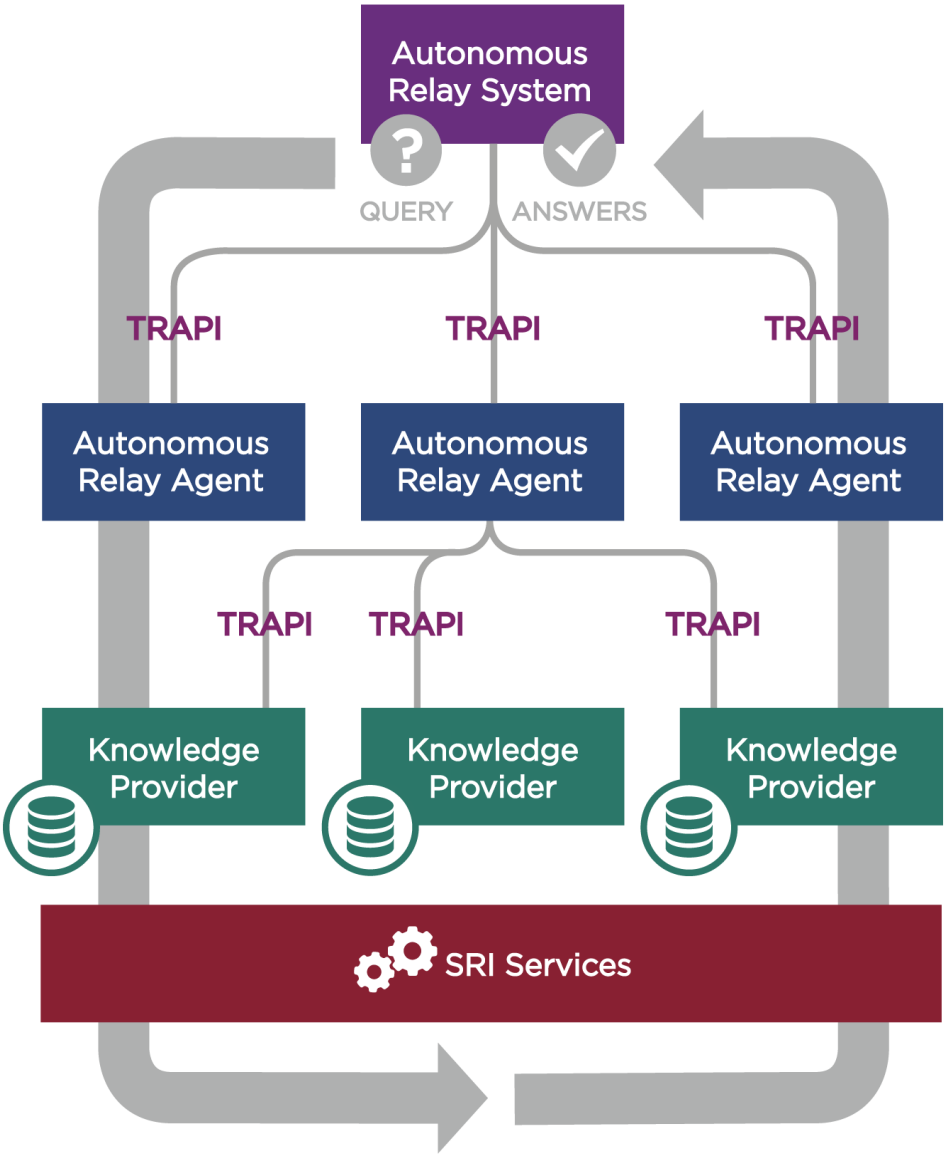

Figure 1. *Overview of the Translator Architecture. Note that while the high-level architecture depicted in the figure is accurate, certain components may deviate slightly from the architecture in their approach to implementation. SRI = Standards and Reference Implementation; TRAPI = Translator Reasoner Application Programming Interface. (Graphic prepared by Kelsey Urgo.)*

Translator comprises four main components: Knowledge Providers (KPs); Autonomous Relay Agents (ARAs); an Autonomous Relay System (ARS); and a Standards and Reference Implementation Component (SRI).

The objective of KPs is to contribute domain-specific, high-value information abstracted from one or more underlying 'knowledge sources', which may be raw data as defined by the DIKW framework or information that has been abstracted from the data. ARAs build upon the knowledge contributed by KPs by way of reasoning and inference and in response to user-defined queries. In addition, ARAs may independently expose information abstracted from data. The ARS functions as a central relay station between ARAs and broadcasts user queries to the

ARAs. The SRI Services are responsible for the development, adoption, and implementation of the standards needed to achieve the overall goals of the Translator Consortium.

Translator leverages integrated data from over 250 knowledge sources, each exposed via open application programming interfaces (APIs). The knowledge sources include, among others, highly curated biomedical databases such as Comparative Toxicogenomics Database [12] and ontologies such as Monarch Disease Ontology [13].

In addition, Translator openly exposes data derived from several electronic health record (EHR) systems, clinical registries, and clinical studies, from which future medical knowledge can be generated: Columbia University Irving Medical Center; UNC Health; the non-profit Providence Health System; the drug-induced liver injury (DILI) Network; the Personalized Environment and Genes Study within the National Institute of Environmental Health Sciences; the Institute for System Biology's Wellness cohort; and select cancer cohorts from within The Cancer Genome Atlas. Of importance, the Translator clinical KPs do not expose raw clinical data, but rather aggregated or semi-aggregated data and statistical associations or machine learning predictions derived from clinical data, in full compliance with all federal and institutional regulations [14].

The Translator Consortium has adopted several tools and approaches to support standardization, harmonization, and interoperability across the diverse Translator system. First, all Translator services are accessible via APIs. The APIs are standardized in their metadata, structure, and operations using the Translator Reasoner API standard (TRAPI) [15], which defines a standard HTTP protocol for transmitting queries and receiving answers, with both structured as graphs. Second, all Translator services are registered in the SmartAPI registry [16], thus adhering to FAIR principles. Third, the open-source Biolink Model [17–20] provides an upper-level graph-oriented data model universal schema that facilitates semantic harmonization and reasoning across disparate knowledge sources.

With these standards in place, users can query across the numerous data sources that are accessible via the federated Translator system. To demonstrate, we provide a simple example. Suppose a user asks *what chemical entities treat chronic pain?* The user is thus asking about approved drugs and other chemicals that may treat chronic pain. To answer this question, the user question must first be translated into a TRAPI-compliant directed query graph, structured in JSON format, with Biolink Model node and edge types specified and a compact unique resource identifiers (CURIE) used to constrain one node (**Figure 2**).

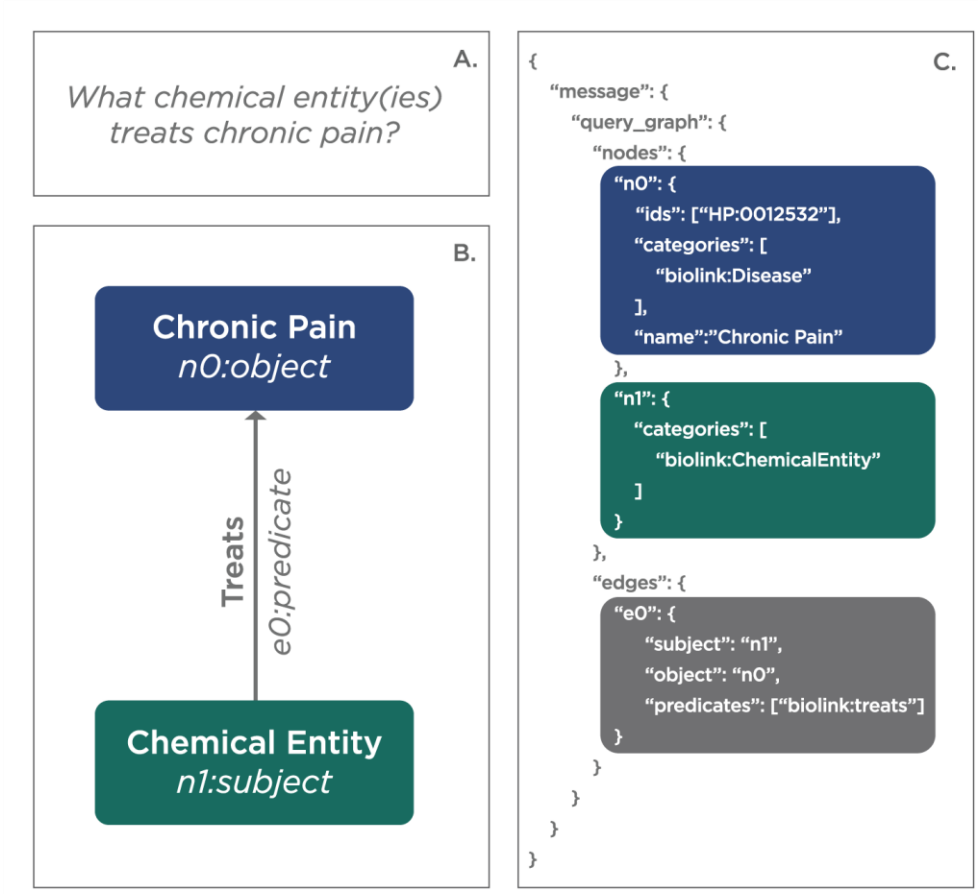

Figure 2. *An example of a natural language question translated into a TRAPI directed query graph in JSON format. (a) the natural language question:* what chemical entity(ies) treats chronic pain?. *(b) the natural language question represented as an object-predicate-subject 'triple'. (c) the TRAPI query that was executed by Translator.*

In this query, 'chronic pain' is specified as a *biolink:Disease* type node 'n0' with the CURIE 'HP:0012532', which is defined by the Human Phenotype Ontology as 'chronic pain'. A second node *n1* is specified only as a *biolink:ChemicalEntity* type. Nodes *n0* and *n1* are related by an edge with the relation defined by a predicate specified as *biolink:treats*. The query graph is thus structured to ask *what chemical entity(ies) treats chronic pain?* The query graph is then sent to the ARS, which parses the query and distributes it to the ARAs. The ARAs then distribute it to those KPs that have provided a meta-graph within the SmartAPI registry indicating that they are able to respond to queries of this type. The ARAs may apply a variety of sophisticated reasoning and inference algorithms to the answers returned by the KPs, including different approaches for ranking and scoring answers such as weighting by supporting publications or abstract co-occurrence of subject and object nodes. Finally, the ARS compiles the ARA results for the user.

A review of the answers to the query finds expected answers such as oxycodone, hydrocodone, codeine, lidocaine, and ibuprofen. There are also answers that are accurate but may not be responsive to the user's query such as methadone, which is used to treat opioid dependence [21], and caffeine, which is an adjuvant in certain pain medicine formulations [22]. In addition, the answer set includes perhaps unexpected answers such as naloxone and naltrexone, which are opioid antagonists. An examination of the evidence and provenance that Translator returns in

support of these answers identifies publications in the form of PubMed identifiers (PMIDs), with links to PubMed abstracts that suggest that these compounds may be effective in the treatment of chronic pain conditions such as fibromyalgia and inflammatory bowel conditions (**Figure 3**). While a pain specialist may not find these findings surprising, many users likely would be surprised to find that there are cases in which an opioid antagonist is beneficial in the treatment of pain, for which opioid agonists are often administered.

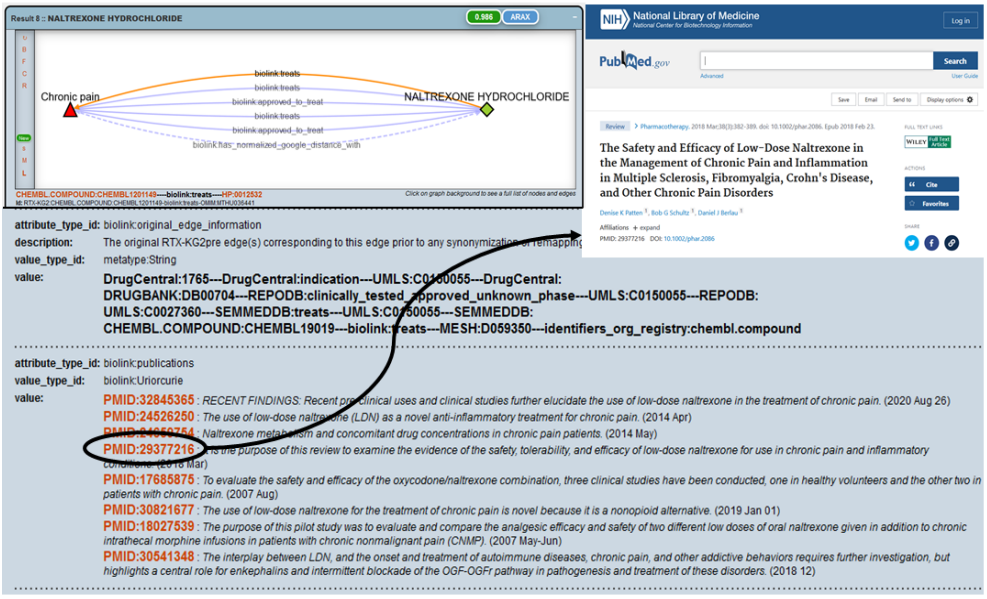

Figure 3. *An example of Translator evidence and provenance in support of naltrexone hydrochloride as an answer to the query in Figure 2.*

Application Use Cases

The chronic pain use case illustrates basic Translator functionalities in the context of a simple ‘one-hop’ Translator query (i.e., two nodes connected by one edge) and the types of insights and discoveries that the Translator Consortium intends to achieve. Here, we provide an overview of three additional use cases (**Figure 4**).

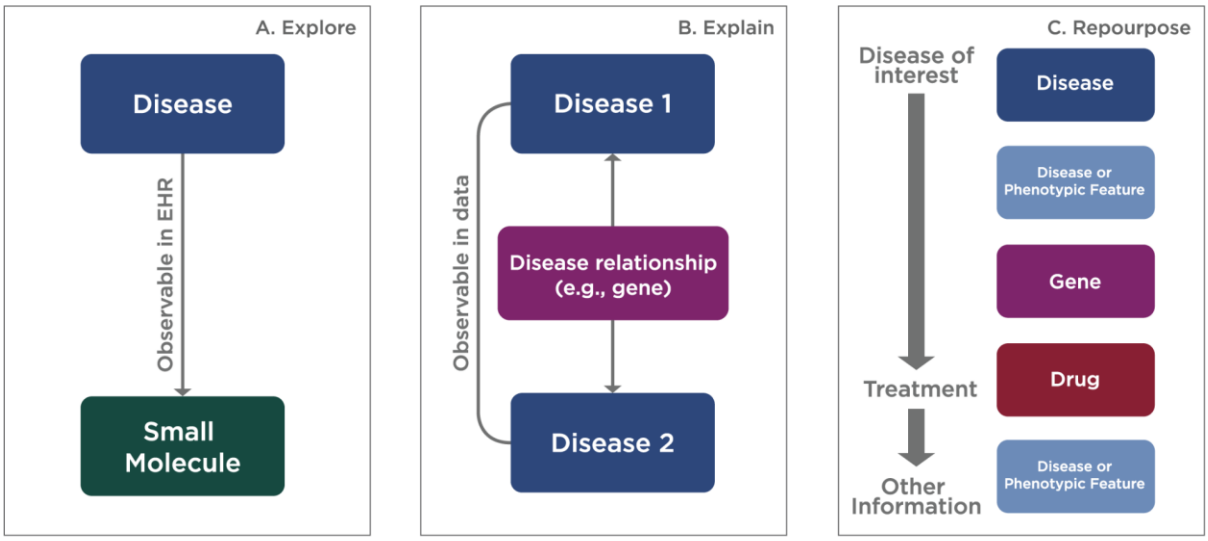

Figure 4. *Schematic of three generalizable Translator workflows applied to support specific use-case queries on (a) immune-mediated inflammatory disease, (b) Parkinson's-Crohn's Disease relationship, and (c) drug-induced liver injury. (Graphic prepared by Kelsey Urgo.)*

Explore: Immune-mediated Inflammatory Diseases (IMID)

The IMID use case was motivated by an interdisciplinary team that was interested in learning more about immunomodulatory drugs that are used to treat IMIDs, including systemic sclerosis, which is a spectrum of rare diseases involving excess collagen that can lead to fibrosis of the skin and/or internal organs. The team was interested in many classes of drugs, including Janus kinase inhibitors (JAK-Is), which have been suggested in the literature as a potential treatment for systemic sclerosis. The team thus approached the Translator Consortium with the following question: *what real-world evidence is there for use of JAK-Is in patients with systemic sclerosis?*

Structured EHR data do not track the condition for which a medication is prescribed to a given patient. An investigator can examine co-occurrence rates between diagnoses and medications, but those rates can be deceptive due to the prevalence of commonly prescribed drugs such as acetaminophen among the general population. Translator clinical KPs have overcome this limitation of EHR data by allowing users to openly explore both co-occurrence rates and relative frequencies of medications, as well as information on whether a medication is contemporaneously predictive for a given disease or phenotype, thus provisioning informative EHR data and assertions without regulatory hurdles.

In this case, the Consortium approached the user's question by executing a one-hop query that targeted Translator clinical KPs (**Figure 4A**). Translator identified drugs such as methotrexate, dexamethasone, and sulfasalazine, which are commonly used to treat a variety of IMIDs. When examining JAK-Is, Translator found evidence of co-occurrence in patients with IMIDs, including systemic sclerosis. Translator also found that JAK-Is were contemporaneously predictive of systemic sclerosis, thereby supporting the assertion that they are being prescribed for people who have systemic sclerosis. With this evidence in hand, the investigative team now plans to use Translator to explore mechanistic evidence that supports the effectiveness of JAK-Is in the treatment of systemic sclerosis.

Explain: Crohn's Disease and Parkinson's Disease

This use case was motivated by clinical observations that patients with Crohn's disease are at risk of Parkinson's disease – two apparently independent diseases. Specifically, the investigative team approached the Translator Consortium with the following question: *why do patients with Crohn's disease have a higher risk of developing Parkinson's disease?*

The Consortium addressed this question by constructing a two-hop query that sought biomedical entities that might be shared by both Crohn's disease and Parkinson's disease (**Figure 4B**). The query was structured with two specified *biolink:Disease* nodes, each connected to an unspecified *biolink:NamedThing* node (i.e., a root class for all things and informational relationships).

Due to the structure of the query, Translator returned a variety of biomedical entities, including *genes*, *diseases*, *chemicals*, and *drugs*. The genes included *LRRK2* (leucine rich repeat kinase 2), *PARK7* (Parkinsonism associated deglycase), and *NOD2* (nucleotide binding oligomerization domain containing 2). Moreover, Translator provided quantitative publication support for each gene's involvement in both Crohn's disease and Parkinson's disease.

The investigative team had expected *LRRK2* to be among the answers returned to the query, so the fact that this gene indeed was returned by Translator provided the team with confidence in the accuracy and sensitivity of Translator answers. The investigative team now plans to take a deeper dive into the supporting evidence and generate new queries to determine if there are common biological processes that might explain how these shared genes contribute to two presumably unrelated diseases.

Repurpose: DILI

[pending approval by Paul Watkins and the DILI Network Steering Committee - FYI, the DILI Network Pubs Committee requires a final draft before conducting a full review of the manuscript]

The DILI use case was motivated by a partnership between the Translator Consortium and the DILI Network. A high priority for the DILI Network, which is the longest running cohort-based study funded by the National Institutes of Health, is to support a DILI clinical trial. This priority is motivated by the fact that the only consensus treatment for DILI is to discontinue the causal agent, leaving patients with few therapeutic options until the drug injury resolves and leaving underlying diseases and conditions untreated. DILI Network investigators have been unable to identify a suitable therapeutic, namely, one that is generally safe, with sufficient biological justification to support a clinical trial.

Hence, the DILI Network approached the Translator Consortium with this goal in mind. The specific question that was asked was *what drug candidate(s) might be repurposed for the treatment of DILI, and is there sufficient biological plausibility to justify the use of those candidates in a clinical trial?*

The Consortium approached this question with a two-fold solution (**Figure 4C**): (1) implement a complex asynchronous three-hop query to identify candidate drugs, leveraging the knowledge provided by Translator clinical KPs; and then (2) implement a simple one-hop query to find additional support for any candidate drugs thus identified, leveraging the real-world and curated knowledge provided by all KPs.

Translator successfully executed both queries and identified two candidate drugs, both antioxidants that are available over-the-counter and in prescription formulation: resveratrol and quercetin. Translator provided additional evidence to justify the use of these candidates in a clinical trial, including: the identification of intermediary genes that suggest biological plausibility; evidence of effectiveness in rodent models of DILI; and clinical trial precedence in other diseases and conditions such as chronic obstructive pulmonary disease. Having met the criteria for viable drug candidates in clinical trials of DILI, members of the Translator Consortium now plan to prepare a formal report on Translator's findings for consideration by the DILI Network Steering Committee.

Discussion

The Translator program is soon to begin year three of development, having first demonstrated feasibility. While a prototype Translator system now exists, with demonstration of its success in returning valid answers to user questions, there are several areas of improvement required to truly achieve a production-level Translator system.

First, the scoring and ranking algorithms that are invoked by the ARAs are intentionally varied to provide breadth in answer sets and associated evidence. We acknowledge a need to refine the scoring and ranking algorithms in order to prioritize those answers with strong evidence, more complete provenance, and high confidence, thereby enriching for answers that are likely to provide the greatest insights to users.

Second, the TRAPI standard and Biolink Model are critical to standardize queries and answers across the federated Translator system. However, standardization often results in a lack of granularity and an inability to pose nuanced queries. For instance, workflow operations are only minimally supported in the current TRAPI standard. We are working to provision a variety of logical operations such as a graph overlay operation. We are also working to extend the Biolink Model to support nuanced statements by developing a core set of qualifiers that can be used to capture semantic richness.

Finally, while several Translator teams have developed user interfaces (UIs) that support TRAPI queries and answers, a uniform cross-component UI is not yet available, although NCATS recently funded a team to develop one. We recognize the urgent need for such an interface, which will allow us to more efficiently engage users and grow our user base, thus promoting long-term sustainability.

We note that the Translator system is one of several available biomedical KG–based question-answering systems. Others include Causaly [23], Elsevier's Biology Knowledge Graph [24] and related Pathway Studio [25], and Google's Knowledge Graph [10]. We emphasize a few differences between these systems. First, the Translator system is the only open-source, community-contributed system; Causaly and Elsevier's systems are commercial, and Google's Knowledge Graph is largely proprietary. Second, Elsevier's systems are highly specific to basic biology and do not span the translational spectrum. Causaly's system supports a broader set of translational questions, but only a subset of those supported by Translator. Third, Translator supports a more sophisticated set of queries than the other systems. For instance, Google's Knowledge Graph only supports simple 'lookup' operations, albeit with highly sophisticated natural language parsing of user questions. Causaly's system is currently limited to linear two-hop queries. Neither Causaly's nor Elsevier's systems support batch or asynchronous queries, in contrast to the Translator system. Finally, none of the other systems support clinical knowledge such as EHR data.

In conclusion, we have developed a biomedical KG–based Translator system capable of integrating a wide range of data sets and translating those data into insights intended to augment human reasoning and accelerate translational science. We are now working on refinements to the prototype Translator system.

Conflicts of Interest

JH has received grant funding from Pfizer and Novartis for research unrelated to the Biomedical Data Translator program. [To be completed]

Acknowledgements

The authors are grateful to the NCATS Translator Extramural Leadership Team and the Intramural Research Program for their ongoing leadership and support. They also thank Kelsey Urgo for assistance with graphics design.

Funding Support

Funding for graphics design was provided by Stanley C. Ahalt, Renaissance Computing Institute, University of North Carolina at Chapel Hill. [Pending approval as acknowledged or funded by Stan]

Other funding, where applicable

Translator Phase I and II Award Numbers:
OT2TR003434 [Clinical Data Services Provider]
OT2TR003436 [Connections Hypothesis Provider]
OT2TR003428 [Expander Agent]
OT2TR003448 [Explanatory Agent]
OT2TR003427 [Exploring Agent]
OT2TR003430 [Exposures Provider]
OT2TR003433 [Genetics Provider]
OT2TR003450 [(im)Prove Agent]
OT2TR003437 [Molecular Data Provider]
OT2TR003443 [Multiomics Provider]
OT2TR003441 [Ranking Agent]
OT2TR003449 [Standards and Reference Implementation]
OT2TR003445 [Service Provider]
OT2TR003422 [Text Mining Provider]
OT2TR003435 [Unsecret Agent]
OT3TR002026 [Blue]
OT3TR002020 [Green]
OT3TR002025 [Grey]
OT3TR002019 [Orange]
OT3TR002027 [Red]
OT2TR002517 [Alpha]
OT2TR002514 [Gamma]
OT2TR002515 [IR]
OT2TR002584 [UV]
OT2TR002520 [X-Ray]

Contract #75N95021P00636 [Biomedical Data Translator User Interface Development]
NCATS Intramural Research Funding (ZIA TR000276-05) [pending approval]

DILI Network Funding: [pending approval by DILI Network]
National Institute of Diabetes and Digestive and Kidney Diseases [1U01DK065201,1U01DK065193, 1U01DK065184, 1U01DK065211, 1U01DK065238 and 1U01DK065176]

References

1. Wilkinson MD, Dumontier M, Aalbersberg IJJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data*. 2016;3:160018.

2. Austin CP, Colvis CM, Southall NT. Deconstructing the Translational Tower of Babel. *Clin Transl Sci*. 2019;12(2):85.

3. Biomedical Data Translator Consortium. Toward A Universal Biomedical Data Translator. *Clin Transl Sci*. 2019;12(2):86-90.

4. Biomedical Data Translator Consortium. The Biomedical Data Translator Program: Conception, Culture, and Community. *Clin Transl Sci*. 2019;12(2):91-94.

5. Ackoff RL. From data to wisdom. *J Appl Syst Anal*. 1989;16(1):3-9.

6. Alavi M, Leidner DE. Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues. *Miss Q*. 2001;25(1):107-136.

7. Rowley J. The wisdom hierarchy: representations of the DIKW hierarchy. *J Inf Sci Eng*. 2007;33(2):163-180.

8. Schmitt CP, Cox S, Fecho K, et al. *Scientific Discovery in the Era of Big Data: More than the Scientific Method*. Vol 3. RENCI; 2015. Accessed December 13, 2021. https://renci.org/wp-content/uploads/2015/11/SCi-Discovery-BigData-FINAL-11.23.15.pdf

9. Wilcke X, Bloem P, de Boer V. The knowledge graph as the default data model for learning on heterogeneous knowledge. *Data sci*. 2017;1(1-2):39-57.

10. Singhal A. Introducing the Knowledge Graph: things, not strings. Google. Published May 16, 2012. Accessed January 18, 2022. https://blog.google/products/search/introducing-knowledge-graph-things-not/

11. Huang Z, Zheng Y, Cheng R, Sun Y, Mamoulis N, Li X. Meta Structure: Computing Relevance in Large Heterogeneous Information Networks. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '16. Association for Computing Machinery; 2016:1595-1604.

12. Davis AP, Grondin CJ, Johnson RJ, et al. The Comparative Toxicogenomics Database: update 2019. *Nucleic Acids Res*. 2019;47(D1):D948-D954.

13. Wg OT. Mondo Disease Ontology. Accessed December 13, 2021. http://www.obofoundry.org/ontology/mondo.html

14. Ahalt SC, Chute CG, Fecho K, et al. Clinical Data: Sources and Types, Regulatory Constraints, Applications. *Clin Transl Sci*. 2019;12(4):329-333.

15. *ReasonerAPI: NCATS Biomedical Translator Reasoners Standard API*. Github Accessed December 13, 2021. https://github.com/NCATSTranslator/ReasonerAPI

16. SmartAPI. SmartAPI. Accessed December 13, 2021. https://smart-api.info/registry?tags=translator

17. Unni, Moxon, Harris, Mungall. Biolink Model: a universal schema for knowledge graphs in biomedical and translational science. *Clin Transl Sci*.

18. Tree Viz of Biolink. Accessed February 4, 2022. http://tree-viz-biolink.herokuapp.com/

19. *Biolink-Model: Schema and Generated Objects for Biolink Data Model and Upper Ontology*. Github Accessed February 4, 2022. https://github.com/biolink/biolink-model

20. Biolink model. Biolink Model. Accessed January 18, 2022. https://biolink.github.io/biolink-model/

21. Bell J, Strang J. Medication Treatment of Opioid Use Disorder. *Biol Psychiatry*. 2020;87(1):82-88.

22. Derry CJ, Derry S, Moore RA. Caffeine as an analgesic adjuvant for acute pain in adults. *Cochrane Database Syst Rev*. 2014;(12):CD009281.

23. Causaly - Accelerate your research. Accessed December 13, 2021. https://www.causaly.com/

24. Elsevier. Biology Knowledge Graph. Accessed December 13, 2021. https://www.elsevier.com/solutions/biology-knowledge-graph

25. Elsevier. Pathway Studio. Accessed December 13, 2021. https://www.elsevier.com/solutions/pathway-studio-biological-research