

Digital objects to make computable biomedical knowledge FAIR: an infrastructural approach to knowledge representation, dissemination and implementation

Marisa Conte[‡], Allen J Flynn[§], Philip Barrison[‡], Peter Boisvert[‡], Zach Landis-Lewis[‡], Charles Friedman[‡]

[‡] Department of Learning Health Sciences, University of Michigan, Ann Arbor, United States of America

[§] Department of Learning Health Sciences and School of Information, University of Michigan, Ann Arbor, United States of America

Corresponding author: Marisa Conte (meese@umich.edu)

Reviewed v 1

Academic editor: Francisco Andres Rivera Quiroz

Received: 10 Jul 2023 | Accepted: 05 Dec 2023 | Published: 18 Dec 2023

Citation: Conte M, Flynn AJ, Barrison P, Boisvert P, Landis-Lewis Z, Friedman C (2023) Digital objects to make computable biomedical knowledge FAIR: an infrastructural approach to knowledge representation, dissemination and implementation. Research Ideas and Outcomes 9: e109307.

<https://doi.org/10.3897/rio.9.e109307>

Abstract

We present our work to develop digital objects to represent and convey a specific category of scientific knowledge: computable biomedical knowledge (CBK). Properly developed, validated, implemented and stewarded, CBK has the potential to accelerate the translation of actionable knowledge from scientific discovery to clinical application.

Our research takes an infrastructural approach to CBK, initially by focusing on the creation of a conceptual model for packaging computable biomedical knowledge - the Knowledge Object (KO) - and on corresponding efforts to create an architecture for KO management and implementation. Additionally, our work is grounded in the FAIR principles, such that KO artefacts should be findable, accessible, interoperable and reusable and we are exploring aligning KOs with emerging best practices for FAIR Digital Objects (FDO).

The outcomes of this work resonate in clinical contexts, health professions education, healthcare quality improvement, biomedical and translational research and population care. Our KO model is also of interest to researchers and practitioners interested in knowledge science, including those working with semantic technologies and other forms of digital objects.

Keywords

computable biomedical knowledge, digital objects, knowledge objects, metadata, FAIR principles

Introduction

We present our work to develop digital objects to represent and convey a specific category of scientific knowledge: computable biomedical knowledge (CBK). Computable biomedical knowledge (CBK) is defined as: “the result of an analytic and/or deliberative process about human health or affecting human health, that is explicit and, therefore, can be represented and reasoned upon using logic, formal standards and mathematical approaches” (Mobilizing Computable Biomedical Knowledge 2018). Other attributes of CBK include: validation and acceptance by a community; implementability, for example, in the form of an intervention (Friedman and Flynn 2019) and dynamism (Adler-Milstein et al. 2019). CBK is used in a variety of contexts, including clinical care, population health and biomedical research (Adler-Milstein et al. 2019) and consists of diverse types, including algorithms, clinical calculators, practice guidelines, message tailoring models, order sets and predictive models (Flynn et al. 2016).

Traditional methods of publishing and sharing biomedical knowledge – for example, scholarly scientific papers published in medical journals or knowledge contained in core medical texts – can neither keep pace with the accelerating speed of new knowledge generation (Guise et al. 2018) nor readily facilitate stewardship of dynamic knowledge. Additionally, knowledge in human-readable formats requires individual access, analysis and interpretation, leading to idiosyncratic application. In contrast, machine-actionable representations of computable biomedical knowledge enable rapid integration, mass application and management of knowledge at scale (Friedman and Flynn 2019).

When properly developed, validated, implemented and stewarded, CBK has the potential to close the gap between scientific discovery of actionable knowledge and its clinical application, a gap often measured in years (Balas and Boren 2000). CBK is also an essential element of a learning health system (LHS), which learns from a cycle of practice to data > data to knowledge > knowledge to practice (Friedman and Flynn 2019) (where data = patient or performance data, knowledge = clinical or biomedical knowledge and practice = clinical practice). This paper presents our work to develop and demonstrate a conceptual model for packaging computable biomedical knowledge, as well as initial efforts to create an architecture for package management and implementation.

Related work

This work is directly related to previous and current efforts to model knowledge in machine-accessible and machine-actionable formats, including digital objects, research objects and semantic publications. With respect to epistemology, in contrast to other classes of data or information, this work emphasises the modelling, symbolic representation and packaging of scientific knowledge. This knowledge has the character of being “accounted semantic information” (Floridi 2010), which is akin to saying it is evidence-based information that communities find meaningful.

Digital objects

Kahn and Wilensky pioneered the concept of a digital object, defining it as “an instance of an abstract data type that has two components, data and metadata”, in addition to a unique identifier or handle (Kahn and Wilensky 1995). This digital object model was foundational to work to build digital libraries and architectures (Payette and Lagoze 2000, Payette and Lagoze 2013). The concept of compound digital objects also directly provided information for our conceptual model for packaging and describing CBK, as described below.

Our work is also enhanced by the development and implementation of digital objects engineered for scientific applications. The Research Object (RO) model, which containerises “essential information relating to experiments and investigations”, is particularly relevant as it addresses packaging, publishing and preserving the RO for reuse (Bechhofer et al. 2013). Recently, the legacy RO concept has been updated and extended through the development of the RO-Crate, which utilises Linked Data principles to provide information about both resources and relationships (Soiland-Reyes et al. 2022). Current work by our lab to extend the metadata for KOs, which for us are digital objects containing CBK, has responded to these same information needs.

Other related work includes objects to represent experimental or analytic processes. This includes the development of an RDF-encoded, ontology-based model and templates for digital objects describing scientific workflows (Belhajjame et al. 2012) and, more recently, the BioCompute Object (BCO), a standardised JSON representation of a bioinformatics pipeline with a schema for domains that impact reproducibility, interpretation, reuse and verification (BioCompute Portal 2022). Taken as a whole, efforts to represent scientific workflows, processes and pipelines share with our work: a) the encapsulation and description of activities or tasks in machine-usable forms and b) Linked Data representations of these objects.

Semantic publications

Our work is also related to the development of semantic publication methods, for example, nanopublications or micropublications, which disseminate specific elements from scientific papers in machine-readable and -usable formats. The initial nanopublication model extracts elements from a scientific paper as a machine-readable combination of statement

(e.g. a finding or research result) and attribution (information about a statement, for example, author, date, instrumentation), expressed as RDF triples and resolved to a Uniform Resource Identifier (URI) (Groth et al. 2010), while later work stipulates that a nanopublication should also contain an assertion, provenance and publication information (Kuhn et al. 2018).

Scientific Knowledge Objects (SKO) and micropublications aim to represent the contents of a paper in more useful ways. SKO are a semantic representation of the content of scientific papers, including metadata, methods, results and discussion, with patterns corresponding to deductive, inductive and abductive reasoning (Giunchiglia et al. 2010). Micropublications start from the belief that scientific papers represent arguments rather than facts and that arguments rely on both supporting evidence for previous work on the problem or in the domain. As such, while the semantic model of a micropublication includes claims and attributions, it also presents a broader range of concepts and relationships, including the ability to represent relationships like supports/challenges, link to methods and data (Clark et al. 2014).

Our work to package and share CBK has several points of similarity to semantic publications. First, both initiatives recognise that knowledge and information need to be shared in formats that are accessible and actionable by both humans and machines in order to be discovered and used at scale. Semantic publications focus broadly on the extraction and dissemination of machine-usable elements from scientific papers, while our work extracts and makes computer-applicable empirical evidence from clinical practice, studies, guidelines etc. Second, both initiatives recognise the importance of providing information about context, provenance and relationships, in addition to the content itself.

FAIR principles

The FAIR principles (Wilkinson et al. 2016), which call for data to be machine- Findable, Accessible, Interoperable and Reusable, are also essential for the development, management and dissemination of CBK packaged in KOs. While they were developed with machine-readable data in mind, the FAIR principles have been extended to other research products, for example, computational workflows (Goble et al. 2020) and research software (Barker et al. 2022). Similarly, our work is grounded in the FAIR principles; both the conceptual model described below and our metadata encourage the FAIRness of CBK. Additionally, we are exploring aligning our conceptual model for packaging CBK with emerging best practices for FAIR Digital Objects (FDO) (Wittenburg et al. 2019). Here, our focus is on empirically-derived biomedical knowledge rather than individual or aggregated patient data. In that regard, it is related, but different to important efforts to standardise or FAIRify biomedical datasets, for example, recent work utilising ontological models to make observational patient data FAIR for clinical analysis (Queralt-Rosinach et al. 2022).

Knowledge Objects (KOs) and the Knowledge Grid

The Knowledge Systems Lab (<https://knowledge-systems.lab.medicine.umich.edu/>) is a health informatics research group based in the Department of Learning Health Sciences at the University of Michigan's Medical School. Our foundational work is the development of the Knowledge Object (KO), a packaged artefact conveying various representations of modular and extensible computable biomedical knowledge with corresponding metadata. The KO model is content- and language-agnostic and modular. This modularity increases interoperability and makes it possible to implement KOs as single units of knowledge or combine them for more complex operations.

Since 2016, our lab has developed hundreds of KOs (Flynn et al. 2022), as well as the Knowledge Grid (KGrid), a coordinated set of CBK artifact models and prototype infrastructural components for end-to-end creation, management, deployment and widespread application of CBK artefacts.

The following section describes our KO conceptual model, relevant ontologies and KGrid platform architecture. KGrid materials are freely available under a GPLv3 licence and include purpose-built collections of KOs (<https://kgrid-objects.github.io/>), demo projects (<https://demo.kgrid.org/>) and applications (<https://kgrid.org/guides/download/>), including the Activator, Library and command line interface.

KO conceptual model

We understand KOs as having a dual nature: KOs are both resources to be managed and services to be implemented. As a resource, KOs can be curated, stewarded and disseminated. As a service, KOs can be implemented in specific contexts and applied to case data automatically and, therefore, at scale. Both aspects of this nature are represented in our conceptual model (Fig. 1).

The KO conceptual model includes:

- a payload: a bit sequence encoding machine-actionable biomedical knowledge, for example, a computable clinical guideline or statistical prediction model. The knowledge payload takes as input patient data and provides as output the relevant computed results, for example, risk calculation, treatment recommendation etc.;
- a service description – currently an OpenAPI document - that specifies the inputs and outputs and the webservice operation required to implement the payload;
- a deployment description that provides instructions for running the specified payload.

These elements are packaged in a wrapper containing administrative, descriptive and technical metadata and the package is assigned a unique persistent identifier.

The Figure below (Fig. 2) presents a simple example of a functional KO as viewed in the KGrid Library. The payload is a risk score interpreter, which takes as its input a patient's

calculated 3-year risk score for developing hepatocellular cancer and provides a corresponding treatment recommendation. In this case, the executable CBK payload is formatted and concretised as a Javascript file. Metadata for managing this KO as a resource includes title, contributors, keywords and citations to the source knowledge (here, the paper describing the development and validation of the scoring algorithm). Additional technical information detailing how to deploy the payload so that it runs and to engage the running payload as a service is provided in the YAML files for service and deployment descriptions. The entire KO (ark:/99999/fk4474n87d/v1.0) and its files may be implemented in a hosted runtime environment or downloaded for other uses. Clicking the Play button packages the KO and deploys it to a hosted runtime environment or a local server.

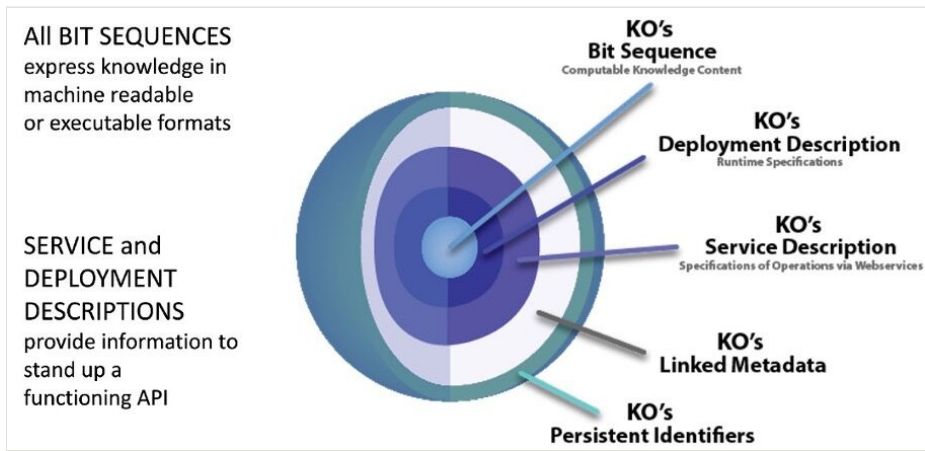


Figure 1. [doi](#)

Conceptual model of a Knowledge Object (KO) containing a payload, machine-actionable service and deployment specifications, metadata and a unique persistent identifier. We are exploring aligning our conceptual model with emerging best practices for FAIR Digital Objects. Derived from Wittenburg et al's *Digital Objects as Drivers towards Convergence in Data Infrastructures* (Wittenburg et al. 2019).

ark:/99999/fk4474n87d/v1.0

3-Year Liver Cancer Risk Score Interpreter

v1.0

3-Year Liver Cancer Risk Score Interpreter

Hepatocellular Cancer

KGID TEAM

service.yaml

deployment.yaml

http://www.ncbi.nlm.nih.gov/pubmed/21497551

Server response

Code Details

200

```

Response body
{
  "result": "<div class='\"'interpretation'\">p style='\"'font-weight: bold;'\">TREATMENT RECOMMENDATION: Consider a liver biopsy or a non-invasive liver fibrosis assessment and, if moderate or greater information on fibrosis are detected,then treat Hepatitis B with one of two first-line oral antiviral medications,entecavir (1 mg tablet by mouth on an empty stomach daily for adults) or Tenofovir (300mg tablet by mouth daily for adults)</p><p>ADMONITORING RECOMMENDATION: Every 3 to 6 months, monitor and reassess liver function (ALT), Hepatitis B e Antigen status (HBeAg), and Hepatitis B DNA virus copy counts (HBV-DNA) </p></div>";
  "metadata": {
    "id": "99999/fk4474n87d/v1.0",
    "type": "KnowledgeObject",
    "identifier": "ark:/99999/fk4474n87d/v1.0",
    "version": "v1.0",
    "title": "3-Year Liver Cancer Risk Score Interpreter",
    "contributors": "KGID TEAM",
    "description": "This object accepts a 3-Year Hepatocellular Cancer risk score as input and generates a recommendation depending on an individual's risk of developing liver cancer.",
    "citations": [
      "http://www.ncbi.nlm.nih.gov/pubmed/21497551"
    ]
  },
  "keywords": [
  ]
}

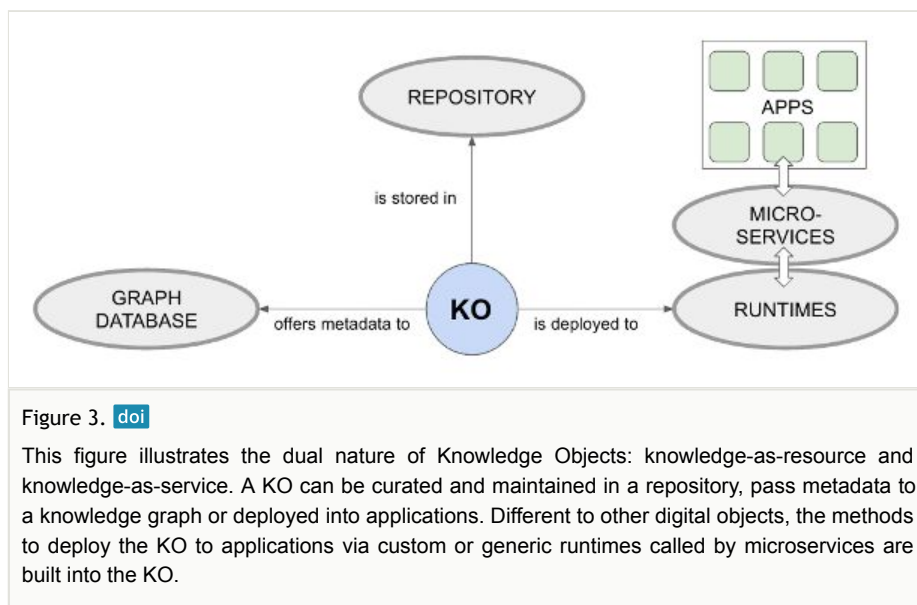
```

Download

Figure 2. [doi](#)

(L) Sample KO as viewed from the KGrid Library, from which the KO can be implemented in a hosted runtime environment or downloaded. (R) Sample output results from deploying the KO.

KOs containing computable biomedical knowledge are similar in many ways to the digital objects and semantic publishing models mentioned above. Similarities include a shared foundational understanding that the contents of a digital object should be machine-accessible, packaged with metadata and uniquely identified by a resolvable persistent resource identifier. However, there are also significant differences. The greatest difference stems from the tendency of the above-mentioned models to treat the object and its payload (whether data or knowledge) primarily as a resource, while Knowledge Objects containing CBK have essential properties of both a resource and a service. For example, where micropublications include both the representation and argumentation of statements and are meant to augment the scholarly publishing ecosystem, KOs are made to make knowledge machine-actionable, such that it can be implemented in existing infrastructure or stand-alone applications to perform specific tasks (Fig. 3). In other words, the potential for knowledge contained within a KO to be implemented by machines, with limited human involvement, is an essential component of the conceptual model and a KO must contain, link to, or describe, in both human- and machine-actionable ways, the methods by which the knowledge can be applied.



KO ontologies - KORO and KOIO

Representing KOs with an ontology facilitates machine interpretability. The Knowledge Object Reference Ontology (KORO) (<https://bioportal.bioontology.org/ontologies/KORO>) is a Basic Formal Ontology-based ontology, which extends the Information Architecture Ontology to formally specify a Knowledge Object. KORO's scope includes what is needed to build compound knowledge objects, implement them and make them FAIR. The ontology defines both the parts of a KO and the relationship between these parts utilising 110 classes and 19 properties, 78 and 4 of which, respectively, are unique to KORO (Flynn

et al. 2018a). The Knowledge Object Implementation Ontology (KOIO) is a subset of concepts taken from the larger KORO. Its purpose is to specify KOs and their constituent parts using common metadata elements, starting with `koio:KnowledgeObject` (Knowledge Grid 2021).

KGrid Library and Activator

KGrid was initially envisioned as an infrastructural platform that not only specified and packaged a knowledge object, but made it findable and accessible and facilitated its application or implementation. The original architecture consisted of two technical infrastructural components: a Library (enabling KO as resource) and an Activator to facilitate deployment of CBK payloads held in KOs (enabling KO as service). The first prototype library included components for standardised metadata and Archival Resource Key (ARK) ID assignment and registration and a gateway to allow for resource discovery and access through APIs (Flynn et al. 2017).

The first Activator was built in Java, using the Spring Framework and enables the provision of Knowledge-as-a-Service, allowing KOs to be requested and retrieved via an application programming interface (API) call, serialised in JavaScript Object Notation and deployed via webservices (Flynn et al. 2018a). Using the Activator, KO artifacts can be called by a RESTful API and deployed in a variety of environments, including web apps, without the need for special tooling.

KOs in research and application: past and future work

Early projects demonstrated the ability of KOs containing CBK to activate knowledge as a service at scale. One project addressed medication safety, combining KOs to alert physicians to atypical electronic prescriptions in an effort to minimise prescribing errors (Flynn et al. 2018b). Another demonstration project utilised existing knowledge, in the form of Clinical Pharmacogenomic Implementation Consortium guidelines, to provide patient-specific dosing guidelines for 28 drugs, based on genotype-phenotype patterns.

In addition to demonstration projects, KGrid's conceptual model and technology has supported the translation of knowledge into practice. KGrid APIs were used to manage and deliver scoring model calculations as part of an app that recommends precision chemotherapy treatments for paediatric neuro-oncology patients (Ravi et al. 2022).

We continue to expand our understanding of the dual nature of KOs representing knowledge-as-resource and knowledge-as-service. The following section describes current work in four emerging areas of research:

1. Expanding the knowledge-as-service potential for KOs carrying CBK;
2. Developing a typology of CBK through the development of different classes of CBK KOs;

3. Exploring the modular nature of KOs and inherent challenges with respect to engineering, knowledgebase management and composition;
4. Expanding CBK metadata, including Linked Data.

Expanding knowledge-as-service

We have recently completed a retrospective analysis of the technology and projects developed during the first five years of the Knowledge Grid platform (2016 - 2021) and lessons learned. One important finding is that, rather than enabling knowledge-as-service, reliance on the original Activator may complicate KO implementation in certain situations. As a result, we have identified a need for technical work to update the original Activator model, as well as development of a specification and reference implementation for activation using different runtimes. The goal of this work is to provide end-users and developers with multiple ways to access and activate KOs, rather than restricting them to any single implementation or workflow. This work is intended to simplify KO integration into clinical workflows. The development of the specification and reference implementation will simplify the method of accessing and deploying KOs through an API call, while providing a software development kit will allow developers to integrate knowledge objects into existing workflows and runtime environments.subsection text.

Developing new classes of KOs and a typology of CBK

By engineering and packaging a variety of CBK as KOs, in accordance with the FAIR principles, we are developing a typology of CBK, understanding the unique characteristics of knowledge developed for different purposes, for example, risk prediction, clinical decision support, patient classification etc. This typology, together with the work described above to expand models of activation, will require updates to KORO and KOIO, so that KOs can be adequately described in a standardised way consistent with interoperability.

One example of current work includes defining a class of KO that facilitates patient cohort identification for clinical studies (Conte et al. 2022). Computable phenotypes are a machine-processable expression of a pattern of observable characteristics of interest, related to a disease or condition, derived from data in an electronic health record system. Computable phenotypes have a variety of use-cases in clinical and population health research and patient care, yet there are currently no standardised ways to represent and share this important knowledge and phenotypes are difficult to repurpose, reproduce or reuse. To address this need, the goal of our work is to develop computable phenotypes as executable KOs (Flynn et al. 2018c), including developing multiple implementations of the core classification knowledge.

Exploring modularity and composition

We are also interested in the modular nature of KOs, which holds exciting potential for combining different types of CBK for specific purposes. One recent example utilises KOs to prioritise preventative interventions for primary care providers and population health

researchers. In the Composite Model for Individualized Precision Prevention (CM-IPP) project, 42 KO submodels were created, in a nested hierarchy, with each submodel representing one preventative medical service recommended by the United States Preventive Task Force (USPTF). At the top level, an executive submodel utilises conditional logic to determine which other submodels should be engaged and these submodels rank preventative service recommendations, based on the patient's unique characteristics (Flynn et al. 2021). In addition to demonstrating the utility of composite models, this project revealed the need for deeper understanding of how to modularise computable knowledge. Further work in this area will focus on both the engineering and management requirements of KOs and knowledge bases. Additionally, this project extends our understanding of knowledge-as-resource as we consider the requirements of knowledge management for submodels and composite models, for example, to account for versioning and updates to knowledge.

Expanding CBK metadata

Finally, metadata is a primary focus of our research exploring the knowledge-as-resource nature of KOs. Metadata is an essential component of making computable biomedical knowledge FAIR and our work with specific types of CBK includes developing sufficient metadata, such that the artifacts can be discovered, accessed and implemented. This work relates to ongoing efforts within the Mobilizing Computable Biomedical Knowledge (MCBK) community to describe thirteen categories of metadata for computable knowledge Alper et al. (2021) and to develop a minimal metadata specification for CBK artefacts. Additionally, recent work has focused on several metadata categories that are specific to KOs, including formalising relationships between KOs and between KOs and empirical evidence. KOs may relate to other KOs sequentially (e.g. diagnostic knowledge preceding treatment knowledge), dependently (e.g. stratification depends on measurement) or comparatively (where multiple models estimate the same factor) and more work is needed to formalise these relationships. Moreover, metadata about two different kinds of evidence are required: features and details about how the computable knowledge was generated and evidence of effects of applying the computable knowledge to simulated or real cases.

We are also exploring the use of semantic technologies, including linked data to extend our metadata model (Bizer et al. 2008) and implementing KOs in conjunction with knowledge graphs (Flynn et al. 2022) Our linked metadata conform to the Resource Description Format (RDF), where domains of interest are represented using a pattern of subject-predicate-object "triples". These RDF triples constitute machine-actionable metadata records that can be visualised as directed graphs. Working examples of these metadata have been successfully loaded into the JSON-LD Playground and an instance of the Blue Brain Nexus knowledge graph system. To demonstrate machine-actionability, we have used SPARQL queries to extract and filter elements from these linked metadata records. Next steps include studies of constructing and managing knowledge bases comprised of many interacting KOs.

Conclusions

This paper presents our work to develop and demonstrate a conceptual model for packaging computable biomedical knowledge aligned with the FAIR principles, as well as initial efforts to create infrastructural components that can be added to an architecture for CBK management and implementation. Within healthcare and health informatics, this work has immediate relevance to two communities: learning health systems and a growing international community dedicated to Mobilizing Computable Biomedical Knowledge (MCBK) (<https://mobilizecbk.med.umich.edu/>). The outcomes of this work resonate in clinical contexts, health professions education, healthcare quality improvement, biomedical and translational research and population care. Our model will also be of interest to researchers and practitioners interested in knowledge science, including those working with semantic technologies and other forms of digital knowledge objects.

Hosting institution

Department of Learning Health Sciences, University of Michigan Medical School, Ann Arbor MI.

Conflicts of interest

The authors have declared that no competing interests exist.

References

- Adler-Milstein J, Nong P, Friedman C (2019) Preparing healthcare delivery organizations for managing computable knowledge. *Learning Health Systems* 3 (2). <https://doi.org/10.1002/lrh2.10070>
- Alper B, Flynn A, Bray B, Conte M, Eldredge C, Gold S, Greenes R, Haug P, Jacoby K, Koru G, McClay J, Sainvil M, Sottara D, Tuttle M, Visweswaran S, Yurk RA (2021) Categorizing metadata to help mobilize computable biomedical knowledge. *Learning Health Systems* n/a (n/a). <https://doi.org/10.1002/lrh2.10271>
- Balas EA, Boren SA (2000) Managing Clinical Knowledge for Health Care Improvement. *Yearbook of Medical Informatics* 09 (01): 65-70. <https://doi.org/10.1055/s-0038-1637943>
- Barker M, Chue Hong N, Katz D, Lamprecht A, Martinez-Ortiz C, Psomopoulos F, Harrow J, Castro LJ, Gruenpeter M, Martinez PA, Honeyman T (2022) Introducing the FAIR Principles for research software. *Scientific Data* 9 (1). <https://doi.org/10.1038/s41597-022-01710-x>
- Bechhofer S, Buchan I, De Roure D, Missier P, Ainsworth J, Bhagat J, Couch P, Cruickshank D, Delderfield M, Dunlop I, Gamble M, Michaelides D, Owen S, Newman D, Sufi S, Goble C (2013) Why linked data is not enough for scientists. *Future Generation Computer Systems* 29 (2): 599-611. <https://doi.org/10.1016/j.future.2011.08.004>

- Belhajjame K, Corcho O, Garijo D, Zhao J, Newman D, Palma R, Bechhofer S, Garcia E, Gomez-Perez JM, Klyne G, Page K, Ruiz JE, Soiland-Reyes S, Roure DD, Goble CA (2012) Workflow-Centric Research Objects: First Class Citizens in Scholarly Discourse. 903. CEUR Workshop Proceedings. URL: <http://ceur-ws.org/Vol-903/paper-01.pdf>
- BioCompute Portal (2022) Introduction - BCOeditor Wiki. URL: <https://wiki.biocomputeobject.org/index.php?title=Introduction>
- Bizer C, Heath T, Idehen K, Berners-Lee T (2008) Linked data on the web (LDOW2008). Proceedings of the 17th international conference on World Wide Web. [ISBN 978-1-60558-085-2]. <https://doi.org/10.1145/1367497.1367760>
- Clark T, Ciccarese PN, Goble CA (2014) Micropublications: a semantic model for claims, evidence, arguments and annotations in biomedical communications. Journal of Biomedical Semantics 5 (1). <https://doi.org/10.1186/2041-1480-5-28>
- Conte M, Flynn A, Boisvert P, Landis-Lewis Z, Richesson R, Friedman C (2022) Computable phenotypes for cohort identification: core content for a new class of FAIR Digital Objects. Research Ideas and Outcomes 8 <https://doi.org/10.3897/rio.8.e95856>
- Floridi L (2010) Semantic information and the network theory of account. Synthese 184 (3): 431-454. <https://doi.org/10.1007/s11229-010-9821-4>
- Flynn A, Friedman C, Boisvert P, Landis-Lewis Z, Lagoze C (2018a) The Knowledge Object Reference Ontology (KORO): A formalism to support management and sharing of computable biomedical knowledge for learning health systems. Learning Health Systems 2 (2). <https://doi.org/10.1002/lrh2.10054>
- Flynn A, Adler-Milstein J, Boisvert P, Gittlen N, Lagoze C, Meng G (2018b) ScriptNumerate: A Data-to-Advice Pipeline using Compound Digital Objects to Increase the Interoperability of Computable Biomedical Knowledge. AMIA ... Annual Symposium proceedings. AMIA Symposium 2018: 440-449. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6371244/>
- Flynn A, Boisvert P, Gittlen N, Gross C, Iott B, Lagoze C, Meng G, Friedman C (2018c) Architecture and Initial Development of a Knowledge-as-a-Service Activator for Computable Knowledge Objects for Health. Studies in Health Technology and Informatics 247: 401-405. <https://doi.org/10.3233/978-1-61499-852-5-401>
- Flynn A, Taksler G, Caverly T, Beck A, Boisvert P, Boonstra P, Gittlen N, Meng G, Rath B, Friedman C (2021) CBK model composition using paired web services and executable functions: A demonstration for individualizing preventive services. Learning Health Systems n/a (n/a). <https://doi.org/10.1002/lrh2.10325>
- Flynn A, Conte M, Boisvert P, Richesson R, Landis-Lewis Z, Friedman C (2022) Linked Metadata for FAIR Digital Objects Carrying Computable Knowledge. Research Ideas and Outcomes 8 <https://doi.org/10.3897/rio.8.e94438>
- Flynn AJ, Shi W, Fischer R, Friedman C (2016) Digital Knowledge Objects and Digital Knowledge Object Clusters: Unit Holdings in a Learning Health System Knowledge Repository. 2016 49th Hawaii International Conference on System Sciences (HICSS). <https://doi.org/10.1109/HICSS.2016.413>
- Flynn AJ, Bahulekar N, Boisvert P, Lagoze C, Meng G, Rampton J, Friedman CP (2017) Architecture and Initial Development of a Digital Library Platform for Computable Knowledge Objects for Health. Stud Health Technol Inform (235)496-500. <https://doi.org/10.3233/978-1-61499-753-5-496>
- Friedman C, Flynn A (2019) Computable knowledge: An imperative for Learning Health Systems. Learning Health Systems 3 (4). <https://doi.org/10.1002/lrh2.10203>

- Giunchiglia F, Xu H, Birukou A, Chenu R (2010) Scientific knowledge object patterns. Proceedings of the 15th European Conference on Pattern Languages of Programs. [ISBN 978-1-4503-0259-3]. <https://doi.org/10.1145/2328909.2328928>
- Goble C, Cohen-Boulakia S, Soiland-Reyes S, Garijo D, Gil Y, Crusoe M, Peters K, Schober D (2020) FAIR Computational Workflows. *Data Intelligence* 2: 108-121. https://doi.org/10.1162/dint_a_00033
- Groth P, Gibson A, Velterop J (2010) The anatomy of a nanopublication. *Information Services & Use* 30 (1-2): 51-56. <https://doi.org/10.3233/ISU-2010-0613>
- Guise J, Savitz L, Friedman C (2018) Mind the Gap: Putting Evidence into Practice in the Era of Learning Health Systems. *Journal of General Internal Medicine* 33 (12): 2237-2239. <https://doi.org/10.1007/s11606-018-4633-1>
- Kahn R, Wilensky R (1995) A Framework for Distributed Digital Object Services. URL: <http://www.cnri.reston.va.us/k-w.html>
- Knowledge Grid (2021) Knowledge Grid Guides - KOIO. URL: <https://kgrid.org/guides/tutorial/ko/koio.html>
- Kuhn T, Banda J, Willighagen E, Ehrhart F, Evelo C, Malas T, Dumontier M, Merono-Penuela A, Malic A, Poelen J, Hurlbert A, Centeno Ortiz E, Furlong L, Queralt-Rosinach N, Chichester C (2018) Nanopublications: A Growing Resource of Provenance-Centric Scientific Linked Data. 2018 IEEE 14th International Conference on e-Science (e-Science). [ISBN 978-1-5386-9156-4]. <https://doi.org/10.1109/eScience.2018.00024>
- Mobilizing Computable Biomedical Knowledge (2018) MCBK Manifesto. URL: <https://mobilizecbk.med.umich.edu/about/manifesto>
- Payette S, Lagoze C (2000) Policy-Carrying, Policy-Enforcing Digital Objects. In: Borbinha J, et al. (Ed.) *Research and Advanced Technology for Digital Libraries*. 1923. [ISBN 978-3-540-41023-2 978-3-540-45268-3]. https://doi.org/10.1007/3-540-45268-0_14
- Payette S, Lagoze C (2013) Flexible and Extensible Digital Object and Repository Architecture (FEDORA). arXiv. arXiv:1312.1258 [cs]. URL: <http://arxiv.org/abs/1312.1258>
- Queralt-Rosinach N, Kaliyaperumal R, Bernabé C, Long Q, Joosten S, van der Wijk HJ, Flikkenschild EA, Burger K, Jacobsen A, Mons B, Roos M, BEAT-COVID Group, COVID-19 LUMC Group (2022) Applying the FAIR principles to data in a hospital: challenges and opportunities in a pandemic. *Journal of Biomedical Semantics* 13 (1). <https://doi.org/10.1186/s13326-022-00263-7>
- Ravi K, Meng G, Roberts H, Schepers A, Franson A, Koschmann C, Marini B, Flynn A (2022) Expanding Access to CNS-TAP: Design, Development, and Initial Use of a Complex Precision Health Specialty Web App for Neuro-Oncology. *MEDINFO 2021: One World, One Health – Global Partnership for Digital Innovation* 804-808. <https://doi.org/10.3233/SHTI220190>
- Soiland-Reyes S, Sefton P, Crosas M, Castro LJ, Coppens F, Fernández J, Garijo D, Grüning B, La Rosa M, Leo S, Ó Carragáin E, Portier M, Trisovic A, RO-Crate Community, Groth P, Goble C (2022) Packaging research artefacts with RO-Crate. *Data Science* 5 (2): 97-138. <https://doi.org/10.3233/ds-210053>
- Wilkinson M, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, Blomberg N, Boiten J, da Silva Santos LB, Bourne P, Bouwman J, Brookes A, Clark T, Crosas M, Dillo I, Dumon O, Edmunds S, Evelo C, Finkers R, Gonzalez-Beltran A, Gray AG, Groth P, Goble C, Grethe J, Heringa J, 't Hoen PC, Hooft R, Kuhn T, Kok R, Kok J, Lusher S,

Martone M, Mons A, Packer A, Persson B, Rocca-Serra P, Roos M, van Schaik R, Sansone S, Schultes E, Sengstag T, Slater T, Strawn G, Swertz M, Thompson M, van der Lei J, van Mulligen E, Velterop J, Waagmeester A, Wittenburg P, Wolstencroft K, Zhao J, Mons B (2016) The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* 3 (1). <https://doi.org/10.1038/sdata.2016.18>

- Wittenburg P, Strawn G, Mons B, Boninho L, Schultes E (2019) Digital Objects as Drivers towards Convergence in Data Infrastructures. [web publication] <https://doi.org/10.23728/B2SHARE.B605D85809CA45679B110719B6C6CB11>