

TECHNICAL REPORT

The Knowledge Object Reference Ontology (KORO): A formalism to support management and sharing of computable biomedical knowledge for learning health systems

Allen J. Flynn^{1,2}  | Charles P. Friedman^{1,2,3} | Peter Boisvert² | Zachary Landis-Lewis² | Carl Lagoze¹

¹School of Information, University of Michigan, Ann Arbor, Michigan

²School of Medicine, University of Michigan, Ann Arbor, Michigan

³School of Public Health, University of Michigan, Ann Arbor, Michigan

Correspondence

Allen J. Flynn, Division of Learning and Knowledge Systems, Department of Learning Health Sciences 300 North Ingalls Street, Ann Arbor, MI48109.

Email: ajflynn@med.umich.edu

Abstract

Introduction: Health systems are challenged by care underutilization, overutilization, disparities, and related harms. One problem is a multiyear latency between discovery of new best practice knowledge and its widespread adoption. Decreasing this latency requires new capabilities to better manage and more rapidly share biomedical knowledge in computable forms. Knowledge objects package machine-executable knowledge resources in a way that easily enables knowledge as a service. To help improve knowledge management and accelerate knowledge sharing, the Knowledge Object Reference Ontology (KORO) defines what knowledge objects are in a formal way.

Methods: Development of KORO began with identification of terms for classes of entities and for properties. Next, we established a taxonomical hierarchy of classes for knowledge objects and their parts. Development continued by relating these parts via formally defined properties. We evaluated the logical consistency of KORO and used it to answer several competency questions about parthood. We also applied it to guide knowledge object implementation.

Results: As a realist ontology, KORO defines what knowledge objects are and provides details about the parts they have and the roles they play. KORO provides sufficient logic to answer several basic but important questions about knowledge objects competently. KORO directly supports creators of knowledge objects by providing a formal model for these objects.

Conclusion: KORO provides a formal, logically consistent ontology about knowledge objects and their parts. It exists to help make computable biomedical knowledge findable, accessible, interoperable, and reusable. KORO is currently being used to further develop and improve computable knowledge infrastructure for learning health systems.

KEYWORDS

BFO, IAO, knowledge management, knowledge object, KORO, ontology

1 | INTRODUCTION

A movement to create learning health systems (LHSs) is ongoing.¹⁻³ One important goal of an LHS is to greatly decrease the current latency of getting knowledge from bench to bedside, accelerating its movement from discovery into practice. This paper describes a formal,

logically consistent Knowledge Object Reference Ontology (KORO) to improve management, enable archiving, and accelerate sharing of actionable, machine-executable biomedical knowledge.

We postulate that the LHS “learns” through a cyclical process that engages a community in empirical analysis of data relevant to a biomedical problem, which leads to the discovery of new knowledge.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2018 The Authors. Learning Health Systems published by Wiley Periodicals, Inc. on behalf of the University of Michigan

The learning cycle is completed by direct application of that knowledge to change practice.^{3,4} Changed practice generates new data, driving the next iteration of the cycle, with improvement occurring via successive iterations. To support learning cycles at scale, the LHS requires a shared-technology infrastructure platform to support 3 information flows: (1) data to knowledge, (2) knowledge to practice, and (3) practice to data (Figure 1).

We are building an LHS infrastructure platform called the Knowledge Grid (KGrid) (www.kgrid.org).⁵ It specifically supports the knowledge-to-practice flow spanning steps 5, 6, and 7 in Figure 1. The KGrid is built upon mature and widely used technologies.⁶ It is for managing and rapidly sharing machine-executable (ie, “computable”) biomedical knowledge. KGrid allows its users to create, manage, and steward digital knowledge objects (KOs),⁷ which are structured packages holding instances of computable biomedical knowledge. KGrid includes a digital library component to store and manage KOs⁵ and a separate activator component to enable remote invocation of the computable knowledge they hold. In this way, KGrid’s activator provides knowledge as a service. Specifically, instances of the activator “activate” computable knowledge by (a) loading KOs stored as digital files, (b) providing a means to execute the computable knowledge held in those files, (c) making those means of execution available to external systems via web services, and (d) tracking the utilization of those web services.

This paper describes the design, development, and initial use of the KORO in support of KGrid. KORO extends the Information Artifact Ontology (IAO)⁸ to formally define what KOs are. Hence, although motivated by the need to more readily share biomedical knowledge, KORO is a general formalism for a digital “package,” called a KO, that holds functional, computable knowledge from any domain. KORO illuminates the required and optional parts that KOs have and the relationships among those parts.

Unlike biomedical application ontologies, whose purposes are to support specific applications or to annotate existing records,⁹ KORO

is a reference ontology that defines KOs as information artifacts.¹⁰ However, since KORO supports practical reasoning about KOs, it has the character of an application ontology to some degree.¹⁰ KORO’s purpose is to serve as a formal specification for the design of KOs, thereby supporting computable knowledge stewardship¹¹ by helping to make computable knowledge findable, accessible, interoperable, and reusable.¹²

2 | BACKGROUND AND SIGNIFICANCE

This section begins by describing the LHS context for which KORO was developed. Then, it explains how KGrid and KOs support routine learning. Next, perspectives on knowledge and knowledge management systems are highlighted, and KGrid is differentiated from similar platforms. Then, learning objects (LOs) and research objects are briefly reviewed to differentiate them from KOs. Finally, the significance of KORO is stated plainly.

2.1 | Context of providing support for the LHS

In LHSs, after a community decides to study a problem, discrete steps can be repeated sequentially to learn (Figures 1 and 2). An example of learning cycle is illustrated in Figure 2. The cycle’s goal is to learn: “How can those at higher risk for lung cancer best be identified for screening?” To begin, data are collected (step 1), assembled (step 2), and analyzed (step 3), resulting in a potentially useful predictive model for estimating individuals’ lung cancer risk¹³ over the next 6 years. This predictive model is actionable and can be represented as machine-executable knowledge (step 5). An implementation of the predictive model can be held in a KO, which in turn can be managed in a KGrid library (step 6). From there, the predictive model can be deployed to support a prediction calculating service using a KGrid activator. When the lung cancer risk predictive model is remotely invoked by an electronic health record (EHR) or other health information technology (IT) system, it is applied as knowledge (step 7). This happens by combining the computable predictive model with facts about a person to generate a prediction for that person’s lung cancer risk that may be useful in practice.

The KO plays 2 important roles in the example above. As shown in Figure 2, at step 6, a KO is used to manage a working implementation of the lung cancer risk predictive model as a persistent resource. At step 7, the computable predictive model held in a KO is accessed and made serviceable via an activator. As previously described above, an activator activates computable knowledge in way that brings about and tracks the utilization of a Web service, which can be called on by other systems. Early efforts to build KGrid indicate that KOs are workable and potentially useful. To date, we have (a) implemented a conceptual model of a KO,⁷ (b) developed and described a digital library to manage KOs,⁵ and (c) built an activator to apply the knowledge held in KOs via knowledge services. Our work anticipates a future where computable biomedical knowledge is created, used in practice, and archived as an integral part of the scientific record.¹¹ These capabilities provide LHSs with the means to recall and use what they learn.

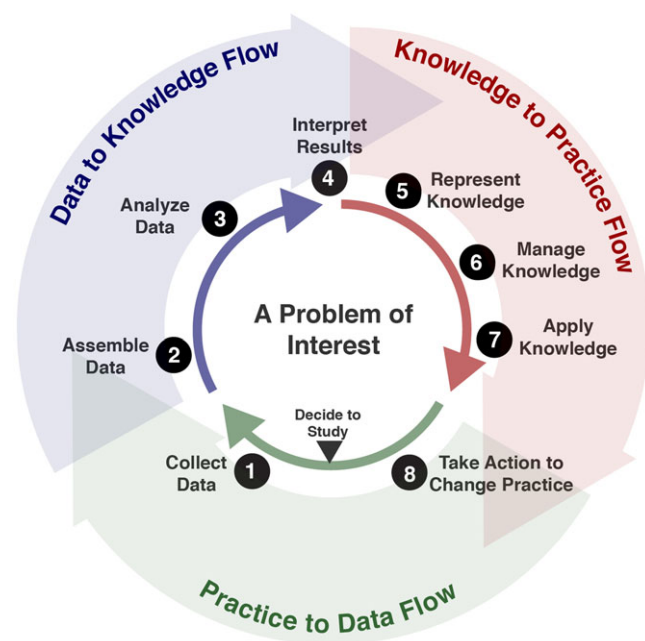


FIGURE 1 The learning health cycle of the learning health system with 3 information flows and 8 steps

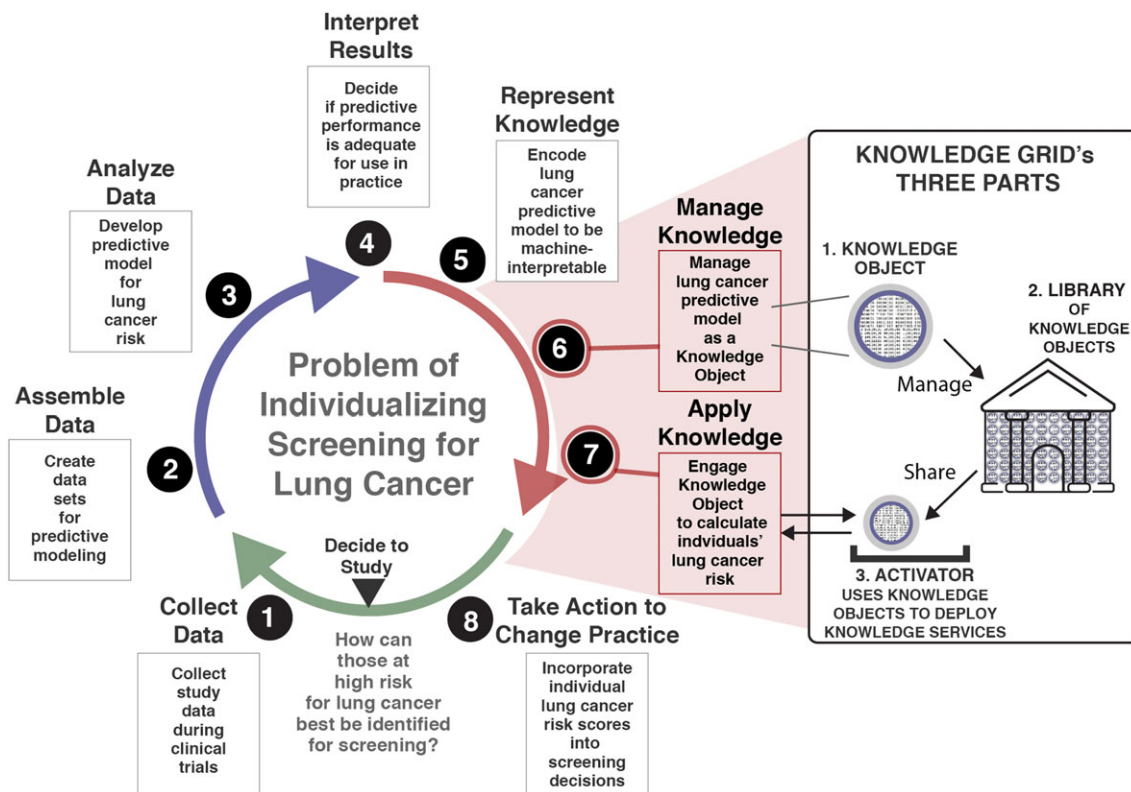


FIGURE 2 Learning cycle example portraying support from knowledge objects and Knowledge Grid

2.2 | Perspectives on knowledge and knowledge management systems

In our work on knowledge infrastructure for LHSs, we adopt 2 perspectives of the computer-supported cooperative work (CSCW) community. Since the 1980s, CSCW researchers have examined how computer systems can be used to support and coordinate collaborative group activities.¹⁴ An annual Association for Computing Machinery conference (ACM) on CSCW showcases this research (<http://cscw.acm.org>).

The first CSCW perspective we adopt is that knowledge can be a result of insights gained from empirical analytic or deliberative processes that are situated in a social context occurring at a particular time and place.¹⁵ This epistemological perspective is very relevant to groups of researchers and practitioners collaborating to generate new health knowledge within LHSs.¹⁶ To support collaborative learning for health, KORO anticipates annotating KOs with a wide variety of descriptive metadata to document KO provenance. This includes metadata about the social context within which KOs are created.

The second CSCW perspective we adopt is that knowledge repositories are necessary but not sufficient to achieve effective knowledge sharing.¹⁵ In published examples,¹⁷⁻¹⁹ we find that knowledge repositories are augmented with capabilities for communication, annotation, and search. KGrid's library is also augmented with capabilities that help users to annotate and deploy computable biomedical knowledge. Furthermore, KGrid may differ some from these previous efforts by assigning to the KO the dual roles of resource and service enabler. These dual roles are defined in KORO.

2.3 | Distinguishing learning objects from KOs

Digital objects for managing knowledge within information systems are not new. In support of e-learning and instructional systems, LOs have a significant history. Yet it is hard to understand precisely what an LO is. An overview in 2003 noted numerous “confusing and arbitrary” definitions of LOs.²⁰ Learning object content includes text, images, audio, video files, and software applications. Similar to websites,²¹ the scope of LOs is so broad that it is difficult to formally define them in a way that can be widely accepted.²⁰ Even so, for the purpose of achieving interoperability, the technical Shareable Content Object Reference Model standard has been adopted and used in many e-learning systems to date.²² A “next-generation” version of Shareable Content Object Reference Model, based on service-oriented architectures, is planned.²³ Taking note of this history, our approach to formalizing the definition of a KO is to develop the KORO ontology and to share it as widely as possible. As indicated in KORO, what distinguishes KOs from LOs is that KOs exist specifically to help manage discrete instances of computable knowledge in ways that facilitate their rapid and widespread deployment and use via knowledge services.

2.4 | Distinguishing research objects from KOs

We have also explored, in a preliminary way, what makes KOs different from research objects. There exists a research object ontology.²⁴ (It should not be confused with “RO,” the relation ontology.²⁵) Per the research object ontology, a research object is an evolving aggregation of heterogeneous resources for describing and reproducing research workflows, along with annotations about those resources.²⁶

In contrast, a KO packages a discrete instance of functional computable knowledge so that the computable knowledge it holds can be shared as a static resource, implemented as a reliable service, and archived as part of the global scientific record. As noted above, KOs play dual roles as resources and service enablers and include both resource and service metadata.

Via personal correspondence with the editors of the research object ontology, we are aware of the potential to embed KOs inside research objects. From the point of view of the research object ontology, embedded KOs coexist within research objects as “resources” that package computable workflow-related “scripts” in a formal way. Embedding KOs inside research objects could make it easier for researchers to put each other’s scripts to use, thereby facilitating evaluation of scripts. Doing this may also facilitate publication of computable analytic results, eg, predictive models.

2.5 | Related ontologies

We reviewed the ontologies from the National Center for Biomedical Ontology online Bioportal (<http://bioportal.bioontology.org/>) and the OBO Foundry (<http://www.obofoundry.org/>). We were unable to identify an existing ontology for digital objects intended to help manage and share computable biomedical knowledge. We did find the Software Ontology, which supports annotation of software tools by type, manufacturer, inputs, outputs, and uses. Software Ontology does not provide enough detail to model a KO and its parts.²⁷

This work raised the question of which, if any, upper-level ontology to use for KORO. We first selected the Basic Formal Ontology (BFO) as an upper-level ontology for KORO. BFO was chosen at first because it enables ontologies to coordinate and interoperate.^{28,29} However, we later discovered the IAO. IAO incorporates upper-level classes from BFO and then extends it. This extension of BFO by IAO includes classes for information entities that are useful for KORO, especially the classes information content entity and directive information entity.⁸ For this reason, we ultimately based KORO directly on IAO, which in turn is based on BFO.

We recognize that IAO, which is subject to ongoing refinement, will eventually need to be embedded into a wider framework of domain ontologies, including the Mental Functioning Ontology.³⁰ However, for KORO, we attempt only to slightly extend the scope of IAO to cover a new type of information artifact, the KO.

2.6 | Significance of this work

Learning health systems surface a complex set of knowledge management and archiving requirements. These include a requirement to directly connect knowledge to practice within the context of ongoing learning. In LHSs, communities of interest need to continuously learn and then share what they learn easily and widely via existing electronic health records and other health IT systems.^{3,4,16} In the KGrid approach, KOs play dual roles as archival knowledge resources, for knowledge management, and as service enablers, for knowledge sharing. By so doing, KOs have the potential to help people meet the knowledge management, archiving, and sharing needs of LHSs.^{4,5,7}

3 | RESEARCH QUESTIONS

The following 4 research questions, labeled RQ1 to RQ4, were investigated during the development of KORO.

- RQ1. What are required and optional parts of a Knowledge Object Content Package, and how may those parts and their relations be represented and described, in a logically consistent way, as entities in an ontology?
- RQ2. What classes, relations and constraints pertain among the whole and the parts of a Knowledge Object, when concretized and materialized, and how may they be represented and made logically consistent in an ontology?
- RQ3. What questions specifically about Knowledge Object parthood can KORO answer competently to demonstrate its logic and potential utility?
- RQ4. What are necessary implementation decisions that have to be made to successfully implement a Knowledge Object that bears a concretization of computable knowledge within its required instance of a ‘Knowledge Object Payload Item’ part?

4 | METHODS

4.1 | Ontology development

A multidisciplinary team at our site, comprised of faculty, graduate students, and developers, collaborated to create KORO. At the outset, we determined that KORO needed to serve as a formal way to describe all of the parts of KOs, and how those parts relate. KORO had to reflect the dual roles of KOs as information resources and knowledge service enablers. We started with the top-level BFO (BFO 2.0) and the IAO (IAO 1.0) merged with BFO,³¹ each represented in Web Ontology Language (OWL) format.^{29,31} We followed the steps to develop ontologies using BFO given by Arp et al.²⁹ In addition, when incorporating terms from BFO and IAO into KORO, we used the minimum information to reference an external ontology term method.³² However, we also included some labels, definitions, and examples to help clarify the meanings of BFO and IAO terms used for KORO. When other external sources inspired KORO terms, these sources were explicitly cited in the ontology.

To address KO parthood, we focused mostly on BFO continuants. We modeled a KO Content Package and its information content entity parts. We also modeled material KOs and their parts. KORO provides Aristotelian definitions for all KORO classes and properties.²⁹ After incorporating BFO and IAO classes and their relations (Figure 3), we added KORO-specific content through a process of term identification. We gathered established terms for the parts of KOs from the digital library and IT worlds. For example, the terms “fact sheet,” for a document comprised of statements that describe key facts and “log,” for a document resulting from automatic recording of activity by a computer. After assigning existing terms to most parts of KOs, we formalized what is meant by the novel terms in the ontology, eg, the term “knowledge payload item.”

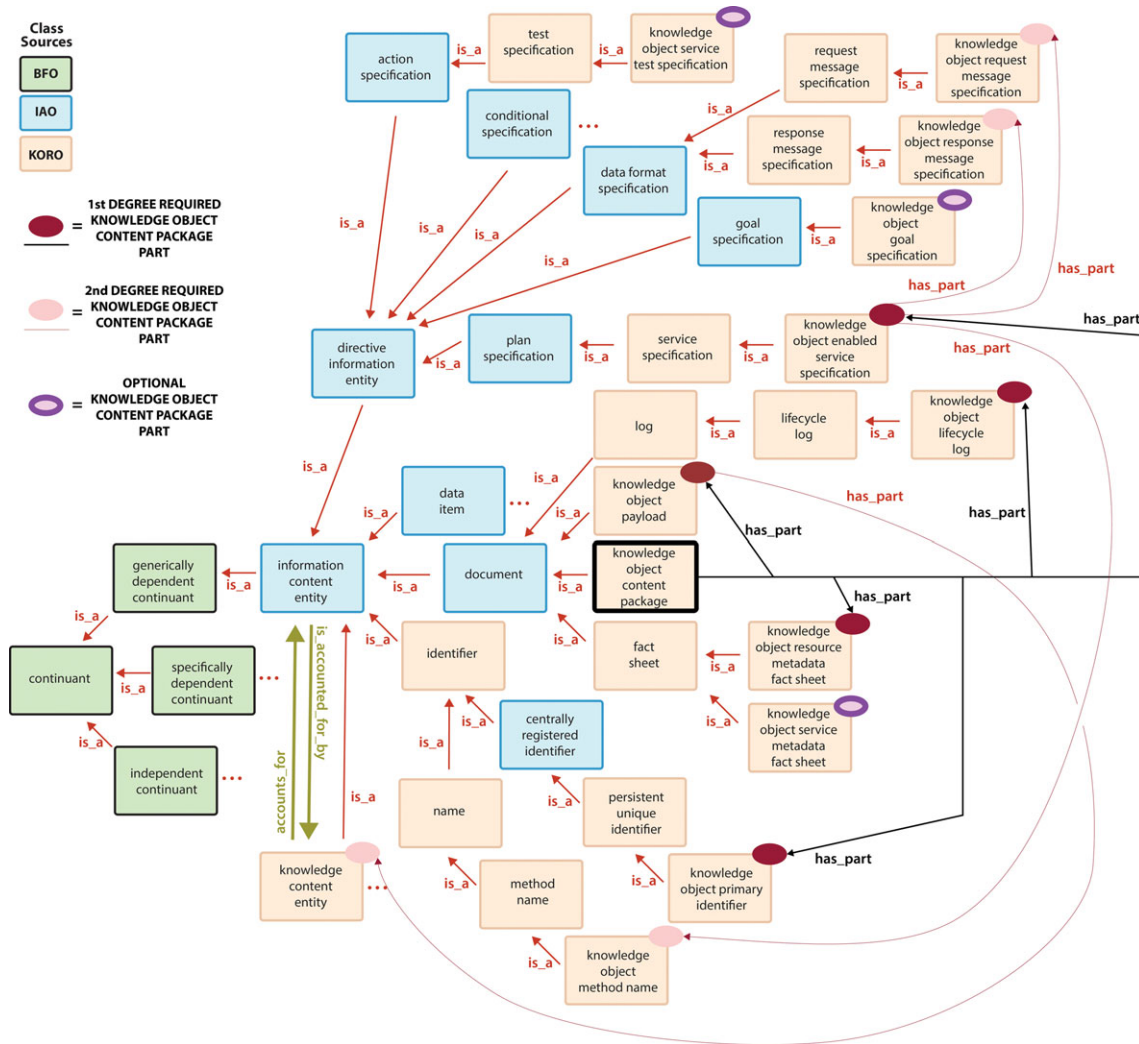


FIGURE 3 Portion of Knowledge Object Reference Ontology (KORO) indicating the required and optional parts of a knowledge object content package

Next, we ordered the terms for KORO's entities in a taxonomical hierarchy using the "is_a" relation to denote subtypes. Then, we added to KORO the minimum number of additional relations needed to represent parthood. We then debated and reordered KORO's entities multiple times through an iterative, collaborative process over 10 months. We finally arrived at a logically consistent version 1.3 of KORO, serialized in Web Ontology Language. After refining KORO using the feedback of experts, we uploaded our original and 2 improved versions of KORO to the National Center for Biomedical Ontology Bioportal at the following location on the World Wide Web: <http://purl.bioontology.org/ontology/KORO>.

4.2 | Evaluation of KORO using competency questions

We developed 3 competency questions (CQs) about KOs and their parts to help illuminate the scope and use of KORO.³³ KORO enables inferences to be made to answer the following 3 CQs:

CQ1. According to its parts, is a given instance of an arbitrary entity a Knowledge Object? (Answer: Yes or No)

CQ2. If an arbitrary entity IS an instance of Knowledge Object, what Knowledge Object part items does it have?

CQ3. If an arbitrary entity IS NOT an instance of Knowledge Object, does it partially fulfill the requirements of Knowledge Object, and, if so, what specific Knowledge Object Part Items does it have?

Using Protégé³⁴ and the Pellet reasoner,³⁵ we created instances of "Knowledge Object" and "Knowledge Object Part Item" suitable to demonstrate that KORO properly infers answers to these 3 CQs.

4.3 | Creation of exemplar KOs to illuminate implementation decisions

To create KOs, a variety of implementation-specific decisions have to be made. One goal we have is to future-proof KORO, as much as possible, so that as technologies and capabilities evolve, it will still provide a useful formalism for a package of computable knowledge. For this reason, KORO sets only a general pattern for all KOs. It does not constrain the format(s) of KO parts, leaving these to be determined by their creators. Hence, to help identify and illuminate the implementation

decisions that have to be made when creating KOs, we have created exemplar KOs of various types. The process of creating actual KOs allowed us to list an initial set of implementation questions that need to be answered about each part of a KO when creating new KOs.

5 | RESULTS

Table 1 includes counts of KORO's classes and properties. To begin this results section, the definitions for key terms from BFO, IAO, and KORO are provided next as background to assist in reporting the remainder of the results.

Fundamentally, BFO separates all "entities" in the world into the following 2 classes:

CONTINUANT	def. an <i>entity</i> that continues or persists through time
OCCURRANT	def. an <i>entity</i> that occurs or happens, such as an event or process with a beginning and an end

Further, BFO separates "continuants" into 3 classes, all 3 of which are included in KORO:

GENERALLY DEPENDENT CONTINUANT	def. a <i>continuant</i> that is dependent for its existence on one or other <i>independent continuants</i> that can serve as its bearer
INDEPENDENT CONTINUANT	def. a <i>continuant</i> that is the bearer of <i>qualities</i> such that <i>qualities</i> inhere in it
SPECIFICALLY DEPENDENT CONTINUANT	def. a <i>continuant</i> that is dependent for its existence on one or more specific <i>independent continuants</i> that serve as its bearer

The difference between generic and specific dependence is that generic dependence accounts for exact copies or clones of an entity. Some entities can migrate from bearer to bearer, as when copies of the same digital file migrate from one hard drive to another hard drive. A "specifically dependent continuant" is an entity that cannot migrate in this manner and instead depends on a specific "independent continuant" bearer for its existence.

IAO extends BFO, starting with the term for this key subclass of "generically dependent continuant":

INFORMATION CONTENT ENTITY (ICE)	def. an <i>entity</i> which is generically dependent on some <i>material entity</i> and which stands in a relation of aboutness to some <i>entity</i>
----------------------------------	---

TABLE 1 Knowledge Object Reference Ontology (KORO) summary statistics

	Number of classes	Number of properties
Total	110	19
From BFO	16	12
From IAO	16	3
Unique to KORO	78	4

Abbreviations: BFO, Basic Formal Ontology; IAO, Information Artifact Ontology.

The term "material entity," used in the definition of ICE above, is a subclass of independent continuant from BFO that is defined in this way:

MATERIAL ENTITY	def. an <i>entity</i> that has some portion of matter as part
-----------------	---

KORO incorporates the term ICE and several terms for ICE subclasses from IAO, with 1 noted modification:

DATA ITEM	def. an <i>ICE</i> that is intended to be a truthful statement about some thing
DIRECTIVE INFORMATION ENTITY	def. an <i>ICE</i> whose concretizations indicate to their bearer how to realize <u>some entity</u> in a process*

*Note: The underlined portion of this definition differs from the current definition given in IAO, which reads, "how to realize *them* in a process." This change is made in KORO because the referent for "them" in the IAO definition of a "directive information entity" was not clear to us.

DOCUMENT	def. an <i>ICE</i> that is a collection of other <i>ICEs</i> intended to be understood together as a whole
----------	--

KORO then adds a number of new terms. It uses them to help define a key new subclass of "document," which has the term "knowledge object content package." That subclass is comprised of instances of 5 other ICE subclasses. Its definition and the definition of one of its most important parts, "knowledge object payload," are given next:

KNOWLEDGE OBJECT CONTENT PACKAGE	def. a <i>document</i> including or containing, at a minimum:
----------------------------------	---

- some knowledge object primary identifier*
- some knowledge object resource metadata fact sheet*
- some knowledge object service specification*
- some knowledge object lifecycle log*
- some knowledge object payload*

KNOWLEDGE OBJECT PAYLOAD	def. a <i>document</i> comprised of one or more <i>knowledge content entities</i>
--------------------------	---

This brings us to a core term in KORO, which is the term "knowledge content entity." It is defined as follows:

KNOWLEDGE CONTENT ENTITY	def. an <i>ICE</i> that describes a result, found to be meaningful to one or more person(s) or a community, such that it can be interpreted by them in ways that they value, and which arises from a systematic analytic and/or deliberative process of investigation and study of other <i>entities</i>
--------------------------	--

As noted above, this pragmatic definition of knowledge content entity is informed by prior work of the CSCW community.¹⁵ The definition of a knowledge content entity reflects the contended epistemological notion that knowledge is information that has somehow been upgraded.³⁶⁻³⁸ Instead of defining precisely how information is upgraded to knowledge, KORO only stipulates the involvement of a systematic process of study. Otherwise, KORO leaves the precise determination of what is knowledge to people and their communities. The following definitions of terms for subclasses of knowledge content entity serve as examples of these types of entities:

CONTENT ANALYSIS RESULT	def. a <i>knowledge content entity</i> that is the result of systematically coding and analyzing qualitative (non-numerical) <i>ICEs</i>
STATISTICAL ANALYSIS RESULT	def. a <i>knowledge content entity</i> that is the result of systematically coding and analyzing quantitative (numerical) <i>ICEs</i>
MODEL	def. a <i>knowledge content entity</i> comprised of a collection of two or more <i>data items</i> and their relationship(s) to each other
EMPIRICAL MODEL	def. a <i>model</i> that describes relationships, or correspondences, found to exist among <i>data items</i> that describe observations
THEORETICAL MODEL	def. a <i>model</i> that describes, in a hypothetical way based on previous observation, potential relationships, or possible correspondences, among <i>data items</i> that have yet to be observed or demonstrated

By definition, then, “models” and other analytic or deliberative results are knowledge content entities. These knowledge content entities are parts of “documents” of the class “knowledge object payload.”

In addition to having as one of its parts a document that is a “knowledge object payload,” a “knowledge object content package” also has a “primary identifier,” a “resource metadata fact sheet,” a “lifecycle log,” and a “knowledge object service specification” as its parts.

Next, turning the focus to subclasses of “material entity,” KORO adopts the definitions of terms for an “artifact” and an “information artifact” given by Smith and Ceusters³⁰:

ARTIFACT	def. a <i>material entity</i> created or modified or selected by some <i>agent</i> to realize a certain function or <i>role</i>
INFORMATION ARTIFACT	def. an <i>artifact</i> whose function is to bear an <i>information carrier</i>

KORO provides a definition for the above term “agent” (not shown), and IAO provides the following definition for the above term “information carrier”:

INFORMATION CARRIER	def. a <i>specifically dependent continuant</i> that is a quality of an information bearer that imparts the information content
---------------------	---

By defining and relating entities, KORO finally arrives at these formal definitions of a “knowledge object” and “knowledge object part item”:

KNOWLEDGE OBJECT	def. an information artifact with knowledge object part items that bears a concretization of a knowledge object content package
KNOWLEDGE OBJECT PART ITEM	def. an <i>information artifact</i> that bears a concretization of an <i>ICE</i> that exists within a <i>knowledge object content package</i>

Although the definitions of the classes “knowledge object” and “knowledge object part item” above do not require it, we anticipate that most actual KOs and KO part items in the world will take the form of digital files. When they take a digital form, instances of these 2 subclasses of “information artifact” better enable the rapid, widespread sharing, deployment, and use of computable knowledge that we intend to achieve.

Having reviewed this background material, we can now proceed to report the remainder of the results. What follows are results organized as responses to the 4 research questions, RQ1 to RQ4, given above.

RQ1. What are required and optional parts of a Knowledge Object Content Package, and how may those parts and their relations be represented and described in a logically consistent way as entities in an ontology?

According to KORO, a “knowledge object content package” is a document, which is an “information content entity,” which is a “generally dependent continuant.” Figure 3 below depicts the portion of KORO’s “*is_a*” relation hierarchy that shows this. It illustrates the parts of a “knowledge object content package,” thereby answering RQ1.

In Figure 3, the key class, “knowledge object content package,” appears toward the center. There are 5 “first-degree” parts that have to be included for a document to be an instance of the class “knowledge object content package.” Subclasses for these 5 first-degree parts are marked with a red oval and connected to the “knowledge object content package” class in Figure 3 by a thick black line portraying the “*has_part*” relation. (Note: Other, optional subclasses for parts of an instance of “knowledge object content package” are marked by purple-pink ovals. To avoid clutter, the “*has_part*” relation to these optional part subclasses is not portrayed in Figure 3.)

In addition, because some “knowledge object content package” parts have their own required parts, 4 more “second-degree” required parts of a “knowledge object content package” are indicated with pink ovals. These are connected to the classes of the parts that require them by thin, red, curved lines. Optional parts of a “knowledge object content package” include those parts indicated by the purple-pink ovals. Finally, in KORO inverse relationships exist between the classes “knowledge content entity” and “information content entity.” These are the “*accounts_for*” and “*is_accounted_for_by*” inverse relationships. They are defined in the following way:

ACCOUNTS FOR	def. a question-answer relationship between an <i>information content entity</i> and a <i>knowledge content entity</i> whereby some <i>knowledge content entity</i> accounts for, meaning explains, the answer to a question about the <i>information content entity</i>
--------------	--

IS ACCOUNTED FOR BY def. an answer-question relationship between a *knowledge content entity* and an *information content entity* whereby an answer to a question about some *information content entity* is accounted for by, meaning is explained by, some *knowledge content entity*

These relationships address a known issue with IAO, which is the need to refine or extend IAO in some manner to include the concepts of truth and/or knowledge.³⁰ In this regard, we assert that it is not the case that all ICEs are truth bearing. Indeed, the “directive information entity” class is defined in a way that instances of procedures may be members of this class. Yet, according to Floridi,³⁶ procedures, eg, “lock the door at 4 pm,” do not “qualify alethically,” meaning they cannot be correctly qualified as true or false. We have been able to create early versions of KORO while recognizing particular issues like this with IAO that remain unsettled.

RQ2. What classes, relations and constraints pertain among the whole and the parts of a Knowledge Object, when concretized and materialized, and how may they be represented and made logically consistent in an ontology?

By definition, a KO is a “material entity.” This means that all KOs have some “portion of matter.”²⁹ A potential ontological confusion arises here. In BFO and IAO, there are subclasses of “material entity” termed “object” and “object aggregate.” Yet in a more recent work, two of the creators of IAO have defined “artifact” and “information artifact” as subclasses of “material entity,” without making it entirely clear how these subclasses relate to the “object” and “object aggregate” subclasses.³⁰ For KORO, we opted to include all 4 of these subclasses of “material entity.” We did this because the definitions of “artifact” and “information artifact” are more recent and more specific to our purposes.³⁰ As indicated above, this means that a “knowledge object” is an “information artifact.” To limit confusion about the 2 similar terms, “object” and “artifact,” we reserve the term “object” in KORO for entities that do not bear concretizations of “information content entities.”

To answer RQ2 it is first necessary to recount how, in BFO, “generally dependent continuants” relate to “specifically dependent continuants.” Further, it is also necessary to recount how “specifically dependent continuants” in turn relate to “independent continuants.” Here are definitions from BFO for these relationships:

IS CONCRETIZED AS def. a relationship between a *generally dependent continuant* and a *specifically dependent continuant*, in which the *generally dependent continuant* depends on some *independent continuant* in virtue of the fact that the *specifically dependent continuant* also depends on that same *independent continuant*

INHES IN def. a relation between a *specifically dependent continuant* (the dependent) and an *independent continuant* (the bearer), in which the dependent specifically depends on the bearer for its existence

The 2 relationships defined immediately above have the following 2 corresponding inverse relationships that are also defined in BFO:

CONCRETIZES def. a relationship between a *specifically dependent continuant* and a *generally dependent continuant*, in which the *generally dependent continuant* depends on some *independent continuant* in virtue of the fact that the *specifically dependent continuant* also depends on that same *independent continuant*

BEARER OF def. a relation between an *independent continuant* (the bearer) and a *specifically dependent continuant* (the dependent), in which the dependent specifically depends on the bearer for its existence

At the top of Figure 4, the inverse relationships between the 3 subclasses of “continuant” are illustrated. Next, toward the middle of Figure 4, the “specifically dependent continuant” subclass of “quality” is shown. It has a key subclass from IAO, “information carrier,” which is also defined above. Instances of “information carrier” are essentially intermediating entities between ICEs and material “information artifacts.” As shown in Figure 4, through instances of “information carrier,” instances of “knowledge object content package” are concretized and then made to inhere in instances of material “knowledge objects.” Hence, all KOs are instances of “material entity,” which in turn are instances of “independent continuant.”

To give further detail, the following key terms for 6 of KORO’s classes of “independent continuant” are defined below and portrayed on the right in Figure 4:

ARTIFACT AGGREGATE def. a *material entity* consisting exactly of a collection of *artifacts* as its member parts at all times at which it exists

INFORMATION ARTIFACT AGGREGATE def. an *artifact aggregate* that is made up of a collection of *information artifacts* bearing concretizations of *information content entities*

KNOWLEDGE OBJECT PART ITEM def. an *information artifact* that is a bearer of a concretization of a required or optional part of a *knowledge object*

INFORMATION ARTIFACT MANIFESTATION def. an *information artifact aggregate* comprised of all instances in reality of some particular *information artifact*

KNOWLEDGE OBJECT MANIFESTATION def. an *information artifact aggregate* comprised of all instances in reality of a particular and uniquely identifiable *knowledge object*

KNOWLEDGE OBJECT PART ITEM MANIFESTATION def. an *information artifact aggregate* comprised of all instances in reality of a particular part item of a *knowledge object*

In addition, the following key terms for inverse relationships between “information artifact” and “information artifact aggregate” are also portrayed in Figure 4:

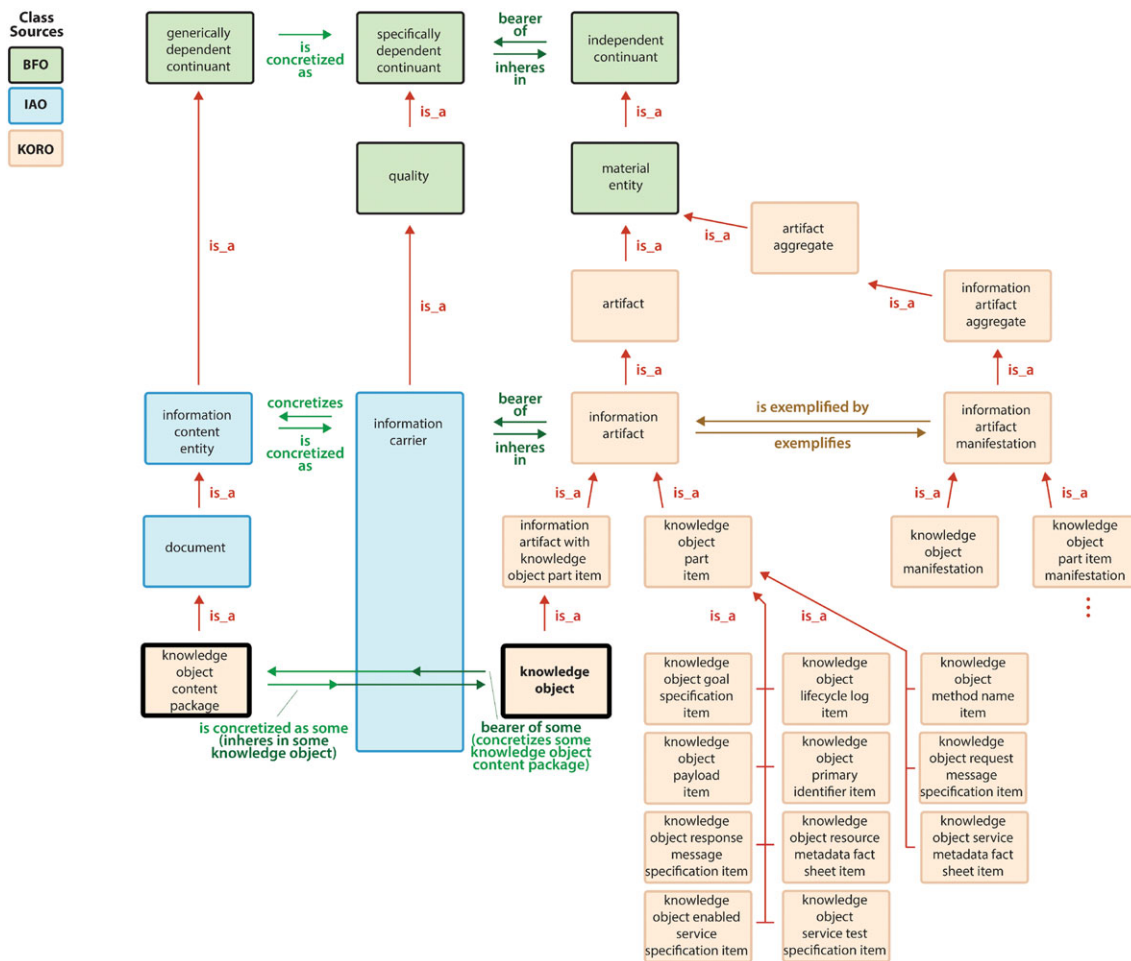


FIGURE 4 Portion of Knowledge Object Reference Ontology (KORO) indicating the whole and parts of a knowledge object with key relationships. BFO, Basic Formal Ontology; IAO, Information Artifact Ontology

IS EXEMPLIFIED BY def. a particular *information artifact manifestation* is exemplified by a particular *information artifact* when the particular *information artifact* is one of potentially many material bearers of the same *information content entity*, and the particular *information artifact manifestation* is the collection of all material bearers of that particular *information content entity*

EXEMPLIFIES def. a particular *information artifact* exemplifies a particular *information artifact manifestation* when the particular *information artifact* is one of potentially many material bearers of the same *information content entity*, and the particular *information artifact manifestation* is the collection of all material bearers of that particular *information content entity*

TEXTUAL ENTITY def. a part of a *manifestation* (FRBR sense) that is a *generally dependent continuant* whose concretizations are patterns of glyphs intended to be interpreted as words, formulas, etc.

However, in contrast to the definition of “textual entity” above from IAO, because in FRBR manifestations and items embody “material entities,” we interpret a manifestation otherwise to be an “independent continuant” and not a “generally dependent continuant.” Our interpretation is reinforced by the FRBR-aligned Bibliographic Ontology, wherein the relation “has manifestation” has as a sub-property “has embodiment” to indicate that a manifestation embodies (ie, makes tangible) some expression.⁴⁰

The answer to RQ2 can now be summarized. By definition, every KO is a “material entity” that bears a concretization of an instance of “knowledge object content package,” which in turn is an instance of “document.” Thus, the content of every KO is constrained by the definition of “knowledge object content package” above, which stipulates that every instance of “knowledge object content package” has 5 required parts. When the defined parts of a “knowledge object content package” are concretized as instances of “information carrier,” then they must inhere in a material instance of “knowledge object part item.” As shown in Figure 4, the class “knowledge object part item”

These inverse “is exemplified by” and “exemplified” relations are directly inspired by the Functional Requirements for Bibliographic Records (FRBR) conceptual model of the International Federation of Library Associations.³⁹ We find that the creators of IAO have also been influenced by FRBR.³¹ The IAO definition of a “textual entity” indicates this:

has 11 subclasses to account for this, one for each type of “knowledge object part item” in KORO.

Finally, as Figure 4 also indicates, because instances of “knowledge object” and “knowledge object part item” are also instances of “information artifact,” they exemplify some “information artifact manifestation.” Hence, there are subclasses of “information artifact manifestation” corresponding to the class “knowledge object,” and also to the class “knowledge object part item” and its subclasses (not shown).

RQ3. What questions specifically about Knowledge Object parthood can KORO answer competently to demonstrate its logic and potential utility?

When given an instance of an arbitrary entity, described in the form of OWL axioms, KORO supports automated reasoning to infer things about the makeup and parts of the arbitrary entity. By so doing, it enables effective answering of 3 competency questions CQ1, CQ2, and CQ3. The results of KORO's performance in this regard are organized by competency question and given in the following text:

CQ1. According to its parts, is a given instance of an arbitrary entity a Knowledge Object? (Answer: Yes or No)

To answer CQ1, KORO provides sufficient logic to the Pellet reasoner to infer from axioms specifying the parts that an arbitrary entity has, whether or not the arbitrary entity is a member of the class “knowledge object.” KORO actually enables a machine to answer this question in 2 different ways. In 1 way, KORO supports reasoning over axioms that indicate an entity has parts that are instances of “knowledge object part item” to determine whether or not those instances include all 5 required parts comprising a “knowledge object.” In another way, KORO supports reasoning over axioms that indicate an entity has parts that are instances of a “knowledge object content package” to determine whether or not those instances include all 5 required parts comprising a “knowledge object content package.” In the latter case, KORO supports further reasoning to correctly infer that a “knowledge object” entity exists when it is asserted that an arbitrary entity is a bearer of a concretization of an instance of “knowledge object content package.”

CQ2. If an arbitrary entity IS an instance of Knowledge Object, what Knowledge Object part items does it have?

To answer CQ2, KORO provides sufficient logic to infer that an arbitrary entity has parts that are instances of “knowledge object part item” and to identify precisely what particular types of “knowledge object part item” the arbitrary entity has. To do this, KORO takes advantage of the logic of the inverse relations “has part” and “is part of.”

CQ3. If an arbitrary object instance IS NOT an instance of Knowledge Object, does it partially fulfill the requirements of Knowledge Object, and, if so, what specific Knowledge Object Part Items does it have?

Building on the logic used to answer CQ1 and CQ2, KORO includes a class “information artifact with knowledge object part item” to enable reasoning that results in answers to CQ3. Using this class, KORO provides sufficient logic to infer from axioms that assert an arbitrary entity has parts that are instances of “knowledge object part item” whether or not that entity is an instance of “knowledge object.” When, an entity having instances of “knowledge object part item” is determined not to be an instance of “knowledge object,” then KORO enables inferences about the part items it does have. This reasoning supports automated determination of the parts of a KO the entity is missing by comparison with a predefined list of all required “knowledge object part items.”

RQ4. What are necessary implementation decisions that have to be made to successfully implement a Knowledge Object that bears a concretization of computable knowledge within its required instance of a ‘Knowledge Object Payload Item’ part?

To answer RQ4, we have built a simple working example of an instance of “knowledge object.” It is referred to from hereon as KO¹. KO¹ has, as one of its necessary parts, an instance of “knowledge object payload part item.” It is referred to as payload-of-KO¹. Payload-of-KO¹ bears a concretization of a very simple mathematical model, which is an example based on a previously published lung cancer risk predictive model.¹³

For this simple example, required KO parts for KO¹ are listed in Table 2, column 1. Corresponding examples of the information content constituting the parts of KO¹ are listed in Table 2, column 2.

Creators of KOs have innumerable options for formatting the information content concretized and born by the material parts of their KOs. For each of the 5 first-degree required “information content entity” parts making up a KO (and for all other parts besides), 3 common implementation questions pertain:

1. What will be the information model or ontology used to represent the content contained in each KO part?
2. How will the information in each KO part be serialized so that it can be managed as a resource in an instance of the Knowledge Grid's digital Library component⁵?
3. How will the information in each KO part be serialized for use as a service-enabler by the Knowledge Grid's Activator component, which allows its users to quickly stand up KO-based web services?

Since 3 common questions pertain to each KO part, to create the 5 first-degree required KO parts for any given KO involves answering 15 different implementation questions. For our simplified example of KO¹, answers to these 15 implementation questions are provided in the third, fourth, and fifth columns of Table 2.

We are currently testing the performance of the actual lung cancer risk prediction model packaged as a resource and made available to support a web service by using a KO. Using an instance of the KGrid's activator, we load our KO into the activator and then the activator uses it to engender a web service. The activator exposes the computable payload in a way that allows external systems to call on it. Eventually, we plan to deploy KOs to provide knowledge as a service. We hope

TABLE 2 Implementation decisions made to create an instance of “knowledge object” called KO¹

Required Part	Contents	Information Model	Library Serialization for Use as a Resource	Activator Serialization to Enable a Service
Knowledge object primary identifier item	Ark://99999/fk4jh3tk9s	Archival resource key ⁴¹	As RDF triple	As JSON
Knowledge object payload item	Risk score (%) = $V1^{0.4} + V2^{1.4} - V3^2$	Python 2.7	As binary file	As JSON
Knowledge object resource metadata fact sheet	Title: Risk model Description: Predictive model Creator: A. Person	Dublin core metadata set ⁴²	As RDF triples	As JSON
Knowledge object enabled service specification	Accepts: V1:integer; V2:integer; V3:integer Returns Y:float	Custom	As RDF/XML file	As key-value pairs in JSON
Knowledge object lifecycle log	Created on: 1-1-2017	PROV-O Ontology ⁴³	As RDF triples	~ N/A ~

Abbreviations: JSON, Javascript Object Notation; RDF, Resource Description Framework; XML, eXtensible Markup Language.

this method will provide a highly scalable way to generate consistent, actionable health information for decision makers.

6 | DISCUSSION

We developed KORO to overcome issues that arose during early development of the KGrid infrastructure platform. Before KORO, we were unable to consistently answer simple questions about KOs, such as “What are the mandatory and optional parts of a Knowledge Object?” Now, with KORO, these and other questions about KOs and their parts can be answered in a consistent way.

In addition, KOs, as we conceive of them, are capable of bearing any form of concretized knowledge content entity. There is no restriction as to the computer language used to encode the information in a KO payload item. The KGrid infrastructure platform is being built in ways that make it computer-language agnostic where KO payload items are concerned. As much as possible, KGrid activator technology is being made extensible so that web services arise from computable biomedical knowledge encoded in a wide array of languages including Python, R, Java, and Javascript.

We envision future work that will leverage KORO to help us build more capable KGrid components. Soon, we plan to further upgrade the KGrid's digital library in accordance with KORO. Among other things, this digital library upgrade will focus on better support for KO versioning and archiving. We also look forward to further investigating how the W3C Open Annotation Model may inform our work to add and manage resource and service-related metadata to KOs.

Because the scope of the present paper is limited to defining and describing KOs and their parts, here, we have not covered the definitions of the dual roles KOs play as resource and service enablers. We have also ignored for now any defined processes that involve KOs, especially instances of the class “service interaction” in KORO. We look forward to further work on KORO that will focus more attention on formalizing roles and processes involving KOs.

Finally, we believe that the development of KORO is one of many necessary steps towards computable knowledge interoperability. By formally defining what KOs are in an ontology, future knowledge

management systems can gain the capability to determine automatically whether or not arbitrary entities are KOs and to delineate whether the parts those entities have are formally defined KO part items.

7 | CONCLUSION

KORO is a realist ontology built using the BFO and the IAO. It supports automated reasoning to determine whether or not an entity is a KO. A KO is a material information artifact that bears a concretized instance of “knowledge object content package.” Knowledge objects have a number of required parts, including an instance of “knowledge object payload item.” In turn, an instance of “knowledge object payload item” is a material information artifact that bears a concretized instance of “knowledge content entity.” Thus, KORO specifies a workable package for managing and sharing any instance of computable knowledge, including instances of computable biomedical knowledge that are needed to sustain and support LHSs.

ACKNOWLEDGMENTS

The following people at the University of Michigan helped with the development of the Knowledge Object Reference Ontology: Rocky Fischer, Frank Manion, and Wei Shi, from Cancer Center Informatics; Andrew Cullen, Lisa Ferguson, Oliver Gadabu, Kristen McGarry, George Meng, and Maria G. Trinidad from the Department of Learning Health Sciences; and Jeff Cowall, Bob Riddle, and John Walsh of Health Information Technology Services.

ORCID

Allen J. Flynn  <http://orcid.org/0000-0002-3471-3063>

REFERENCES

- Olsen L, Aisner D, McGinnis JM. Institute of Medicine (US). Roundtable on evidence-based medicine. The learning healthcare system: workshop summary. 2007; doi:<https://doi.org/10.17226/11903>.
- Friedman CP, Allee NJ, Delaney BC, et al. The science of learning health systems: foundations for a new journal. *Learning Health Syst.* 2017;1:e10020. <https://doi.org/10.1002/lrh2.10020>

3. Friedman CP, Rubin J, Brown J, et al. Toward a science of learning systems: a research agenda for the high-functioning learning health system. *J Am Med Inform Assoc.* 2014. <https://doi.org/10.1136/amiajnl2014-002977>
4. Friedman CP, Rubin JC, Sullivan KJ. Toward an information infrastructure for global health improvement. *IMIA Yearbook.* 2017. <https://www.thieme-connect.com/products/ejournals/pdf/10.1055/s-0037-1606475.pdf>
5. Flynn AJ, Bahulekar N, Boisvert P, et al. Architecture and initial development of a digital library platform for computable knowledge objects for health. In *Informatics for Health: Connected Citizen-Led Wellness and Population Health* (Vol. 235, pp. 496-500). (Studies in Health Technology and Informatics; Vol. 235). IOS Press. <https://doi.org/10.3233/978-1-61499-753-5-496>
6. Lagoze C, Payette S, Shin E, Wilper C. Fedora: an architecture for complex objects and their relationships. *Int J on Digit Libr.* 2006 Apr 1;6(2):124-138.
7. Flynn AJ, Shi W, Fischer R, Friedman CP. Digital knowledge objects and digital knowledge object clusters: unit holdings in a learning health system knowledge repository. In *system sciences (HICSS), 2016 49th Hawaii international conference on 2016 Jan 5* (pp. 3308-3317). IEEE
8. Ruttenberg A, Courtot M, Peters B, Rees J, Malone J. The information artifact ontology. 2008. Accessed 5 Jan 2017 at: <http://www.obofoundry.org/ontology/iao.html>
9. Ashburner M, Ball CA, Blake JA, et al. Gene ontology: tool for the unification of biology. *Nat Genet.* 2000 May;25(1):25-29.
10. Menzel C. Reference ontologies-application ontologies: either/or or both/and?. In *KI Workshop on Reference Ontologies and Application Ontologies.* 2003 Sep 16;16.
11. Van de Sompel H, Treloar A. A perspective on archiving the scholarly web. *IPRES 2014 proceedings 2014*:194.
12. Wilkinson MD, Dumontier M, Aalbersberg IJ, et al. The FAIR guiding principles for scientific data management and stewardship. *Sci Data.* 2016;3:160018.
13. Tammemägi MC, Katki HA, Hocking WG, et al. Selection criteria for lung-cancer screening. *N Engl J Med.* 2013 Feb 21;368(8):728-736.
14. Carstensen PH, Schmidt K. Computer supported cooperative work: new challenges to systems design. In: Itoh K, ed. *Handbook of Human Factors/Ergonomics.* Tokyo: Asakura Publishing; 2003:619-636.
15. Ackerman MS, Dachtera J, Pipek V, Wulf V. Sharing knowledge and expertise: the CSCW view of knowledge management. *Comput Support Coop Work.* 2013 Aug 1;22(4-6):531-573.
16. Faden RR, Kass NE, Goodman SN, Pronovost P, Tunis S, Beauchamp TL. An ethics framework for a learning health care system: a departure from traditional research ethics and clinical ethics. *Hastings Cent Rep.* 2013 Jan;43(s1):1.
17. Ackerman MS, McDonald DW. Answer Garden 2: merging organizational memory with collaborative help. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work 1996 Nov 16.* ACM.
18. Dourish P, Edwards WK, LaMarca A, et al. Extending document management systems with user-specific active properties. *ACM Trans Inf Syst.* 2000 Apr 1;18(2):140-170.
19. Laqua S, Sasse MA, Greenspan S, Gates C. Do you know dis?: a user study of a knowledge discovery tool for organizations. In *proceedings of the SIGCHI conference on human factors in computing systems 2011 may 7* (pp. 2887-2896). ACM
20. Polsani PR. Use and abuse of reusable learning objects. *J Digit Inform.* 2006 Feb 27;3(4).
21. Brügger N. Website history and the website as an object of study. *New Media Soc.* 2009 Feb;11(1-2):115-132.
22. Sinclair J, Joy M, Yau JY, Hagan S. A practice-oriented review of learning objects. *IEEE Trans Learn Technol.* 2013 Apr;6(2):177-192.
23. Poltrack J, Hruska N, Johnson A, Haag J. The next generation of SCORM: innovation for the global force. In *The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC) 2012* (Vol. 2012, No. 1). Orlando: National Training System Association
24. Bechhofer S, De Roure D, Gamble M, Goble C, Buchan I. Research objects: towards exchange and reuse of digital knowledge, 2010.
25. Smith B, Ashburner M, Rosse C, et al. The OBO foundry: coordinated evolution of ontologies to support biomedical data integration. *Nat Biotechnol.* 2007 Nov 1;25(11):1251-1255.
26. Soiland-Reyes, S., Bechhofer S., Eds., *Research object ontology, Release 2016-01-28* Available at: <http://wf4ever.github.io/ro/2016-01-28/> Accessed on: December 10, 2017.
27. Malone J, Brown A, Lister AL, et al. The Software Ontology (SWO): a resource for reproducibility in biomedical data analysis, curation and digital preservation. *J Biomed Semant.* 2014 Jun 2;5(1):25.
28. Smith B, Kumar A, Bittner T. Basic formal ontology (BFO) for bioinformatics. *J Inform Syst.* 2005;1-6.
29. Arp R, Smith B, Spear AD. *Building ontologies with basic formal ontology.* MIT Press; 2015 Jul 31.
30. Smith B, Ceusters W. Aboutness: towards foundations for the information artifact ontology. In *International Conference on Biomedical Ontology, Lisbon, Portugal 2015 Jul* (pp. 47-51).
31. Information Artifact Ontology. Available at: <https://raw.githubusercontent.com/information-artifact-ontology/IAO/master/releases/2017-03-24/iao-merged.owl> Accessed on: December 10, 2017.
32. Courtot M, Gibson F, Lister AL, et al. MIREOT: the minimum information to reference an external ontology term. *Appl Ontology.* 2011 Jan 1;6(1):23-33.
33. Ren Y, Parvizi A, Mellish C, Pan JZ, Van Deemter K, Stevens R. Towards competency question-driven ontology authoring. In *European Semantic Web Conference 2014 may 25* (pp. 752-767). Springer
34. Gennari JH, Musen MA, Fergerson RW, et al. The evolution of Protégé: an environment for knowledge-based systems development. *Int J Hum Comput.* 2003 Jan 31;58(1):89-123.
35. Sirin E, Parsia B, Grau BC, Kalyanpur A, Katz Y. Pellet: a practical OWL-DL reasoner. *Web Semantics: Sci, Services and Agents on the World Wide Web.* 2007 Jun 30;5(2):51-53.
36. Floridi L. Semantic information and the correctness theory of truth. *Erkenntnis.* 2011 Mar 1;74(2):147-175.
37. Floridi L. Semantic information and the network theory of account. *Synthese.* 2012 Feb 1;184(3):431-454.
38. Weinberger D. *Too Big to Know: Rethinking Knowledge Now That the Facts Aren't the Facts, Experts Are Everywhere, and the Smartest Person in the Room Is the Room.* Basic Books; 2011.
39. Tillett B. Functional requirements of bibliographic records (FRBR): final report. *Int Federation of Library Assoc Instit (IFLA).* 1998.
40. Shotton D, Peroni S. FaBiO, the FRBR-aligned bibliographic ontology. Available at: <http://www.sparontologies.net/ontologies/fabio/source.html> Accessed on: December 12, 2017.
41. Kunze J. Towards electronic persistence using ARK identifiers. In *Proceedings of the 3rd ECDL Workshop on Web Archives 2003 Aug 21.*
42. Weibel S. The Dublin core: a simple content description model for electronic resources. *Bulletin of the American Society for Information Science and Technology.* 1997 Oct 1;24(1):9-11.
43. Missier P, Belhajjame K, Cheney J. The W3C PROV family of specifications for modelling provenance metadata. In *Proceedings of the 16th International Conference on Extending Database Technology 2013 Mar 18.* ACM.

How to cite this article: Flynn AJ, Friedman CP, Boisvert P, Landis-Lewis Z, Lagoze C. The Knowledge Object Reference Ontology (KORO): A formalism to support management and sharing of computable biomedical knowledge for learning health systems. *Learn Health Sys.* 2018;2:e10054. <https://doi.org/10.1002/lrh2.10054>