Received: 2 September 2023 | Accepted: 13 September 2023

DOI: 10.1002/lrh2.10396

GUEST EDITOR COMMENTARY

Learning Health Systems

Transforming health and well-being through publishing computable biomedical knowledge (CBK)

Güneş Koru[1,2]

[1]Department of Health Policy and Management, University of Arkansas for Medical Sciences (UAMS), Fayetteville, Arkansas, USA

[2]Department of Biomedical Informatics, University of Arkansas for Medical Sciences (UAMS), Fayetteville, Arkansas, USA

Correspondence
Güneş Koru, University of Arkansas for Medical Sciences (UAMS), Northwest Regional Campus, Fayetteville, AR, USA.
Email: gkoru@uams.edu

Abstract

Computable biomedical knowledge artifacts (CBKs) are software programs that transform input data into practical output. CBKs are expected to play a critical role in the future of learning health systems. While there has been rapid growth in the development of CBKs, broad adoption is hampered by limited verification, documentation, and dissemination channels. To address these issues, the Learning Health Systems journal created a track dedicated to publishing CBKs through a peer-review process. Peer review of CBKs should improve reproducibility, reuse, trust, and recognition in biomedical fields, contributing to learning health systems. This special issue introduces the CBK track with four manuscripts reporting a functioning CBK, and another four manuscripts tackling methodological, policy, deployment, and platform issues related to fostering a healthy ecosystem for CBKs. It is our hope that the potential of CBKs exemplified and highlighted by these quality publications will encourage scientists within learning health systems and related biomedical fields to engage with this new form of scientific discourse.

1 | INTRODUCTION

Computable knowledge refers to the knowledge codified and presented via software programs.[1] *Computable biomedical knowledge artifacts* (CBKs, henceforth) are software artifacts containing machine-interpretable or executable instructions that transform input data into practical outputs. CBKs are expected to play a critical role in the future of learning health systems, enabling mass action in moving from knowledge to performance, truly transforming health and well-being in the digital age. Among other purposes, CBKs can be used to:

- Supplement the traditional human-readable knowledge presented in articles and books, supporting health education and research.
- Perform diagnostic and prognostic predictions by processing and analyzing health data to support clinical decisions (see Box 1 for examples).
- Perform simulations for comparing the future impacts of certain choices, instrumental in dealing with public health crises and policymaking.

In recognition of CBKs' essential roles, the Learning Health Systems (LHS) Journal has encouraged the submission of peer-reviewed CBKs as archival scholarly contributions. The first CBK publication in the LHS Journal presented an open-source immunization calculation engine by Arzt et al.[2] This CBK uses a set of immunization rules and patient data to evaluate and return the validity of each immunization in a patient's history. After evaluating a patient's immunization history, the core of this CBK generates the appropriate immunization recommendations.

The growing interest in CBKs led the LHS Editorial Board to organize a dedicated track[3] in the journal, inaugurated by this Special Issue with eight articles. Four manuscripts report a functioning software

BOX 1 Example CBKs for clinical decision support.

Traditionally, the knowledge presented in articles and books has been static and non-computable. For example, a health professional consulting with printed or printable clinical practice guidelines would leverage static knowledge. The same guidelines could be programmed into software (eg, as a mobile app) constituting a CBK that accepts symptoms and other data as input and provides recommendations as output. This CBK could dynamically update its algorithms and clinical guidance based on the most recent evidence. A predictive model, another CBK with one or more software components, could further support the clinician by leveraging big data to provide personalized recommendations or risk scores for rehospitalization based on patient history, social determinants of health, environment, and other data. Consequently, these two CBKs would constitute "computable knowledge" available to clinicians to supplement the knowledge disseminated through traditional static formats.

artifact—a CBK—developed to support clinical decisions. Four additional manuscripts tackle methodological, policy, deployment, and platform issues related to fostering a healthy ecosystem for CBKs. Before introducing these articles, let us start with the reasons for establishing a dedicated track for publishing peer-reviewed CBKs.

2 | WHY SHOULD CBKS GO THROUGH A PEER-REVIEWED PUBLICATION PROCESS?

Following the Health Information Technology for Economic and Clinical Health (HITECH) Act,[4] software's role in the US healthcare ecosystem has been rapidly transformed from simply record-keeping to harnessing the power of data and information through data analytics and artificial intelligence. In this flourishing technology landscape, a large group of individuals—from professional software engineers to researchers to individual entrepreneurs—have developed CBKs for various healthcare applications. Some CBKs have been described in scientific journals, some were shared broadly as open-source software, and others were made available as proprietary software to their users. Despite the proliferation and availability of CBKs, several questions arise whenever a CBK is considered for adoption in clinical or policy settings:

2.1 | Does it exist?

This might sound like a simple question with an answer taken for granted. Still, external verification mechanisms are lacking to ensure a CBK exists, even when presented in scientific manuscripts. This lack of verification is partly due to the supportive role CBKs have played. Traditionally, CBKs have served as "a means to an end," for example, to prove a scientific point, to support healthcare delivery, or for business purposes. Even when a CBK itself is the main subject, say in a manuscript or presentation, the authors or presenters might only present evidence about its properties, such as usefulness, predictive power, or run-time performance. Without obtaining the CBK and confirming that it runs, the reviewers and audience can only *assume* its existence. Situating the CBK itself at the center of a peer-reviewed publication process confirms its existence at one point in time because the editors and reviewers can receive and execute it. Given growing concerns about data falsification, plagiarism, and predatory publishing, verifying the existence of CBKs is an essential step toward improving the quality of scientific discourse.

2.2 | Did it work as described?

Developing a CBK involves writing software programs, which, to this date, remains a complex and human-intensive task prone to errors.[5] While software quality assurance and control techniques such as inspections and testing are helpful, the evidence shows that these controls are not consistently applied in biomedical software projects for various reasons.[6] Furthermore, even when quality assurance measures are used, controlling software quality becomes increasingly challenging as program size and complexity increase.[7,8] The long-standing evidence and wisdom in software engineering suggest that most CBKs may unavoidably include multiple bugs, including "showstoppers," that cause system crashes or freezes. Given these realities, we should not assume that any CBK will work as described, and hands-on testing and validation of CBKS is warranted. While it may not be feasible to execute (run) all possible execution paths or input combinations for a given CBK, it is possible to validate the critical functionality by checking whether a CBK can execute to provide the expected outputs for a selected set of inputs determined by the authors. In addition to the functional verification, some essential non-functional properties, such as performance (eg, execution time), can be observed and commented on during peer review.

2.3 | Is it documented?

Another unwarranted assumption could be that published CBKs can easily be used by interested parties with reasonable effort. However, using a CBK requires learning about it, which can be challenging without adequate documentation. The development of metadata for digital artifacts has not been a routine practice in biomedical fields. The Standards Workgroup under the Mobilizing CBK (MCBK) initiative recently identified 13 metadata categories to describe CBKs.[9] Some categories, such as type, biomedical domain, purpose, persistent identifiers, and location, are critical for successful adoption. Requiring minimal essential metadata for the peer review process will make them more Findable, Accessible, Interoperable, Reusable, and Trustable (FAIR+T).[10]

3 | WHAT ARE THE BENEFITS OF PEER-REVIEWED PUBLICATION OF CBKS?

Putting CBKs through a peer-reviewed publication process can answer the above questions and contribute to advancing science by addressing at least four important challenges: reproducibility, reuse, trust, and recognition:

3.1 | Reproducibility

Unfortunately, reproducibility remains a challenge in many scientific disciplines.[11,12] A peer-review process, which subjects CBKs to observation of their functional and other properties and constructive criticism by other researchers, is expected to improve CBKs and their documentation, facilitating their execution by third parties. Consequently, refinement and improvement of CBKs through successive studies become easier, advancing science.

3.2 | Reuse

Over time, specific CBKs may prove to be useful to a group of scientists or professionals. The peer-review process prepares CBKs for reuse, and publication makes them widely available. As a result, scientists do not have to reinvent the wheel by testing and documenting again, saving significant resources that would otherwise have to be spent on further development, testing, and documentation. Often, artifacts built from verified components via reuse can achieve higher quality, such as better reliability, performance, and security.

3.3 | Trust

Peer review confirms that CBKs existed, worked as described, and were sufficiently documented. Establishing this basic degree of trust is crucial for relying on CBKs for research or other purposes. Furthermore, trusted CBKs can serve as reliable building blocks in composing complex CBKs.

3.4 | Recognition

Numerous CBKs developed through scientific investigations have often not received proper preparation for sharing beyond the initial research group. One key factor is the lack of support or acknowledgment, apart from the useful commercialization support provided by the technology development offices, for the work needed to ready a CBK for wider utilization and reuse. By establishing an academic platform for CBKs, we provide fresh motivation to allocate resources for CBK development and documentation.

4 | CAN WE RELY ON OPENNESS INSTEAD OF SYSTEMATIC PEER REVIEW?

The journal will ask the authors of CBKs to share their source code and executables to the greatest extent possible, contributing to openness. However, the lessons learned from the open-source movement remind us that openness per se might not be sufficient to ensure the quality of software artifacts. The Linus Law of the 1990s—"*Given enough eyeballs, all bugs are shallow*"[13]—implied that bugs in publicly available code would be more likely to be detected, leading to higher software quality. To this date, however, no evidence has been found to support the Linus Law.[14] Most useful software artifacts are sizeable and complex, and understanding, executing, and reviewing them takes time and effort. For this reason, all open-source software projects that took quality seriously found that they had to adopt systematic review and testing practices to ensure and control quality.[15]

Fortunately, a positive observation from the field of software engineering is that a small number of reviewers, typically between 2 and 4—leads to efficient improvements. The marginal gains increasingly decrease with further, extensive testing or review.[16] Based on these lessons, a peer-review mechanism for CBKs should contribute to their quality and likelihood of adoption. In fact, the quality gains from peer review can increase the benefits of openness by providing more trustable and documented CBKs to open science communities.

5 | WHO WILL BENEFIT FROM PEER-REVIEW AND PUBLICATION OF CBKS?

5.1 | Scientists

Reliability and accuracy are essential in scientific projects. In digital systems, simple bugs can create disproportionately serious failures: Systems can freeze, crash, miscalculate, or distort data; arguments in the wrong order can result in skewed predictions. These examples can be multiplied. The ability to access peer-reviewed and broadly published CBKs, reproduce their results, and reuse them greatly benefits scientists by presenting new research directions and opportunities while boosting their productivity. As discussed above, publishing widely referenced and used CBKs can positively contribute to scientists' careers by bringing academic credit and recognition.

5.2 | Research-funding agencies

Recently, the National Institutes of Health (NIH) and the National Science Foundation have revised their policies to require data-sharing plans that maximize sharing. In the future, extending such policies to cover CBKs can further benefit the scientific community. Peer-reviewed CBKs would significantly contribute to achieving a higher return on investment from research dollars by making CBKs more trustable and shareable.

5.3 | Governments

Biomedical research has implications for local, state, and federal governments interested in addressing health problems while controlling costs. Governments may be interested in adopting or developing many different types of CBKs that support their missions, such as simulation and predictive models. Furthermore, governments might require developing, deploying, and broadly sharing peer-reviewed CBKs as a condition in future contracts and awards.

5.4 | Industry

Industry stakeholders can benefit from increased scientific productivity by developing and sharing CBKs. For example, a company interested in commercializing existing predictive models can achieve greater success if they leverage peer-reviewed and verified CBKs. In addition, companies can also contribute to the literature by creating new CBKs, submitting them for peer review, and publishing them with appropriate intellectual property rights. This broad CBK publishing mechanism can benefit both business and scientific purposes.

5.5 | Patients and their families

Finally, a broader development and adoption of CBKs can benefit patients and their families through improvements in healthcare quality and advances in health research. Some CBKs, such as personal health records and medication adherence apps, may also benefit patients directly; however, there needs to be a stricter check on addressing patient safety and ethical issues for those CBKs.[17]

6 | IN THIS SPECIAL ISSUE…

6.1 | CBKs for decision support

Under the "Sync for Genes" program supported by the Office of the National Coordinator, Dolin et al. standardized the sharing of dynamically annotated gene variants (also called mutations). In their article, "*Sync for Genes Phase 5*: *A standards-based approach to sharing dynamically annotated genomic variants*, "Dolin et al. discuss how their CBKs promote the development of clinical decision support (CDS) applications in surfacing up-to-date annotations to clinicians. The authors published the CBKs' source code in a GitHub repository, and created a backend database with synthetic data. The authors also provided a Swagger API interface that allows interaction with the code and data through any web browser. Finally, they present two proof-of-concept applications demonstrating the use of their CBKs. With this work, Dolin et al. provide examples of how CBKs facilitate moving from knowledge to practice in the LHS cycles.

A common CDS problem is matching patient characteristics with the criteria that apply to the decision support system or approach. To address this problem, Alper et al. developed a CBK called Strike-a-Match Function and described it in their article "Striking a Match between FHIR-based patient data and FHIR-based eligibility criteria." Written in Javascript, the CBK uses JSON input and provides JSON outputs based on HL7 FHIR. The CBK is made available by the authors on GitHub. In addition, the authors offer an Eligibility Criteria Matching software library on fevir.net to share functions, rate them, and collaborate with others.

Scheduling, a traditional problem in outpatient settings, is addressed by Azad et al. in their paper titled, "Application of Computable Biomedical Knowledge to Transform Patient Centered Scheduling." To tackle the scheduling difficulties posed by appointment variability and uncertainty, the authors developed an algorithm for selecting the optimal group of appointments. They coded a linear integer program in Python and posted a Jupyter Notebook on a GitHub repository. Their notebook includes the source code with the sample inputs and outputs. In addition to this reproducible CBK, the authors discuss the technical infrastructure needed to deploy this CBK with a mainstream Electronic Health Records (EHR) system.

Obtaining clinical quality indicators based on real-world data is essential for a Learning Health System to move from performance to knowledge. In "A Computable Biomedical Knowledge Object for Calculating In-hospital Mortality for Patients Admitted with Acute Myocardial Infarction (AMI)," Sadsad et al. report a CBK that operates on data consistent with the OMOP (Observational Medical Outcomes Partnership) common data model. OMOP compatibility increases the reusability of this CBK. The authors implemented this CBK as workflow in the workbench of Piano, a system for data automation to generate insights from data. Through a demo account provided by the authors, the workflow is executable, and the results are reproducible.

6.2 | The CBK ecosystem: Methods, platforms, and policy

Ebben et al. present a novel structured approach to evaluate clinical practice guidelines using real-world data. In "A novel method for continuous measurements of clinical practice guideline adherence," the authors discuss how this method was applied to endometrial cancer patients in the Netherlands.

Wyatt et al. report the results from a workshop tackling a significant problem, the regulation of CBKs. In "Principles guiding the regulation of computable biomedical knowledge libraries and software in the UK," the authors report the participants' view that software executables for medical purposes should be in the scope of regulation. On the other hand, knowledge objects that cannot be directly executed, for example, software-neutral objects written as algorithms using pseudocode, would not be in the scope. Their work is an excellent example of bringing together a group of highly relevant stakeholders to engage in productive discussions, leading to organized views by a CBK community.

In another example of bringing CBK communities together, Scott et al. report results from two "collaborathons" that explored ways of representing the clinical guidelines in a computable format. In their article titled "Modelling clinical narrative as computable knowledge: the NICE computable implementation guidance project," the authors report that modifying the Digital Adaptation Kit developed by the World Health Organization (WHO) is a feasible approach in technology-neutral logical specifications of the recommendations made by the National Institute for Health and Care Excellence (NICE).

In "EvidenceHub: A Place to exchange medical knowledge and form communities," Hong et al. present a novel online platform for sharing various CBKs. The platform advances medical knowledge by facilitating dissemination and discussions around CBKs. In EvidenceHub, the CBKs are published under GPL 2.0 and maintained by volunteers. Hong et al.'s approach to publishing CBKs uniquely supports knowledge growth and gradual refinement. Beyond serving only as a repository, EvidenceHub provides sophisticated tools for the peer review and quality assurance of CBKs. Investigating how to motivate volunteers to contribute to peer review and QA is a fruitful research direction for future studies.

7 | CONCLUSION

Traditional peer review processes for scientific work have served us well over many centuries. However, the problems with reproducibility, reuse, trust, and recognition have steadily increased as we started developing CBKs of growing size and complexity. Peer-reviewed CBK allows us to verify CBKs in the biomedical domain and augments the traditional scientific discourse by facilitating reproducibility, reuse, and increasing trust.

To support this new scientific discourse, we should reconsider the traditional measures of academic productivity, which can exclude or undervalue the development of quality software artifacts. Despite their convenience, we must ask ourselves whether these measures serve the scientific enterprise well to lead us toward the best overall results benefiting our health systems and patients. Over the following decades, do we want to see a continuation of papers with unverified CBKs, further contributing to the so-called "reproducibility crisis"? In a recent Nature survey of 1500 scientists,[11] 90% agreed that reproducibility was a problem, and more than half reported significant reproducibility problems in their fields.

While the characterization of reproducibility problems as crises may be debated, there is a general agreement on the advantages presented by improving reproducibility.[18] Through the new track, the LHS journal and its contributors are taking one important step towards that goal. The CBK track will build a library of trusted CBKs supported by a solid peer-review and verification process. The steps taken by the journal to publish peer-reviewed CBKs are aligned with movements in scientific communities outside of healthcare. For example, the Association of Computing Machinery (ACM) recently encouraged publishing and peer-reviewing computing artifacts and badging them based on their availability, evaluation, reproducibility, and replicability status.[19] If the scientific community benefits from this approach and continues to play a leadership role, other stakeholders, such as government agencies, may also follow our lead with support for this and similar initiatives.

Let us finish this commentary by explaining what the interested authors should submit their CBKS to the LHS journal. The authors must submit two components:

1. *A functioning CBK*. The authors will be expected to demonstrate that their CBK executes and produces identical outputs given identical sample inputs like functions with contained and controlled side effects.[20] Both low- and high-level programs can constitute a CBK. Low-level programs might be provided as assembly language or machine code. High-level programs might be written in system programming languages such as C, C++, or Java, and translated to executable machine code for execution. CBKs might also be written in scripting languages such as Python, Tcl, or R, to be executed by an interpreter. Some CBKs can include websites with a graphical user interface, mobile applications, web services, or Application Programming Interfaces (APIs). Some CBKs can compose other relevant CBKs and add functionality.
2. *A manuscript accompanying the CBK is also required for each submission*. The written manuscript will provide the background and documentation to explain the CBK within 2000 to 4000 words.

Further information about the CBK and manuscript submission requirements can be found at the links below:

- Computable Knowledge Papers Track[21]
- Computable Knowledge Publication Author Instructions[3]

ORCID

Güneş Koru https://orcid.org/0000-0003-1693-4986

REFERENCES

1. Friedman CP, Flynn AJ. Computable knowledge: an imperative for learning health systems. *Learn Health Syst*. 2019;3:e10203.
2. Arzt NH, Chertcoff D, Nicolary S, Suralik M, Berry M. Immunization calculation engine: an open source immunization evaluation and forecasting system. *Learn Health Syst*. 2022;6:e10285.
3. Friedman C, Koru G. Learning Health Systems: Computable Knowledge Publications - Instructions for Authors. 2021.
4. HITECH: Health Information Technology for Economical and Clinical Health Act.
5. Brooks FP. No silver bullet, essence and accidents of software engineering. *IEEE Comput*. 1987;20:10-19.
6. Koru G, El Emam K, Neisa A, Umarji M. A survey of quality assurance practices in biomedical open source software projects. *J Med Internet Res*. 2007;9:e8.
7. Kitchenham B, Pfleeger SL. Software quality: the elusive target. *IEEE Softw*. 1996;13:12-21.

8. Tian J. *Software Quality Engineering: Testing, Quality Assurance, and Quantifiable Improvement*. Hoboken, NJ: John Wiley & Sons, Inc.; 2005.
9. Alper BS, Flynn A, Bray BE, et al. Categorizing metadata to help mobilize computable biomedical knowledge. *Learn Health Syst*. 2022;6:e10271.
10. Wilkinson MD, Dumontier M, Aalbersberg IJJ, et al. The FAIR guiding principles for scientific data management and stewardship. *Sci Data*. 2016;3:160018.
11. Baker M. 1,500 scientists lift the lid on reproducibility. *Nature*. 2016;533:452-454.
12. ATCC. Six factors affecting reproducibility in life science research and how to handle them. 2017 https://www.nature.com/articles/d42473-019-00004-y
13. Raymond ES. The cathedral and the bazaar. *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*. Sebastopol, CA: O'Reilly and Associates; 1999.
14. Favato D, Ishitani D, Oliveira J, Figueiredo E. Linus's law: more eyes fewer flaws in open source projects. Proceedings of the XVIII Brazilian Symposium on Software Quality. Association for Computing Machinery, New York, NY, USA; 2019:69-78. doi:10.1145/3364641.3364650
15. Koru G, Tian J. Defect handling in medium and large open source projects. *IEEE Softw*. 2004;21:54-61.
16. Glass R. *Facts and Fallacies of Software Engineering. in 174–177*. Boston, MA: Addison Wesley Professional; 2003.
17. Platt J, Spector-Bagdady K, Platt T, et al. Ethical, legal, and social implications of learning health systems. *Learn Health Syst*. 2018;2:e10051.
18. Fanelli D. Is science really facing a reproducibility crisis, and do we need it to? *PNAS*. 2018;115:2628-2631. doi:10.1073/pnas.1708272114
19. Association for Computing Machinery. Artifact Review and Badging - Current. https://www.acm.org/publications/policies/artifact-review-and-badging-current
20. *Pure Function*. San Francisco, CA: Wikimedia Foundation; 2022. https://en.wikipedia.org/wiki/Pure_function
21. *Computable Knowledge Papers*. Hoboken, NJ: Wiley Online Library; 2021. https://onlinelibrary.wiley.com/page/journal/23796146/homepage/ckp. doi:10.1002/(ISSN)2379-6146

How to cite this article: Koru G. Transforming health and well-being through publishing computable biomedical knowledge (CBK). *Learn Health Sys*. 2023;7(4):e10396. doi:10.1002/lrh2.10396