

NIGMS Sandbox: A Learning Platform toward Democratizing Cloud Computing for Biomedical Research

Ming Lei^{1*}, Lakshmi K Matukumalli¹, Krishan Arora¹, Nick Weber², Rachel Malashock², Fenglou Mao³, Susan Gregurick³, and Jon Lorsch¹

¹National Institute of General Medical Sciences; ²Center for Information Technology; ³Office of Data Science Strategy, Office of Director; National Institutes of Health. * Current Address: West Virginia University Health Science Center, G102, 108 Biomedical Road, P O Box 9104 Morgantown, WV 26506-9104

Corresponding Author: Ming Lei; (304) 293-7206; ming.lei@hsc.wvu.edu

ABSTRACT

Biomedical data is growing exponentially in both volume and levels of complexity, thanks to the rapid advancement of technologies and research methodologies. Analyzing these large datasets, referred to collectively as “big data”, has become an integral component of research that guides experimentation-driven discovery and a new engine of discovery itself as it uncovers previously unknown connections through mining of existing data. To fully realize the potential of big data, biomedical researchers need access to high-performance-computing (HPC) resources. However, supporting on-premises infrastructure that keeps up with the consistently expanding research needs presents persistent financial and staffing challenges, even for the most resource-rich institutions. For other institutions, including primarily undergraduate institutions (PUIs) and minority serving institutions (MSIs), that educate a large portion of the future workforce in the United States, this challenge presents an insurmountable barrier. Therefore, new approaches need to be explored to provide broad and equitable access to HPC resources to researchers and students who are the next generation of biomedical researchers.

INTRODUCTION

Cloud computing, i.e., *on-demand and highly scalable computing and storage* on clustered high-power servers managed by service providers such as Google Cloud, Amazon Web Services, and Microsoft Azure, holds the promise of releasing institutions from the burden of owning and continuously upgrading information technology infrastructure to provide HPC capacity for research. Cloud service providers (CSPs), with resources as well as incentives to invest in updating and innovation, offer the public their services in a pay-as-you-go manner. As consumers, academic and/or research institutions pay for needed resources and computing time based on actual usage. Paying on a per-use-basis for needed cloud capacity and computing time, instead of paying the premium of standing up, maintaining, updating, and operating data centers, is more manageable financially for many institutions. In addition to cost savings, cloud computing has the potential to increase the research productivity of research-intensive institutions. For example, the time needed to complete a project on servers of an on-premises data center can be dramatically reduced by the ‘near-infinite’ parallelization available on commercial cloud infrastructure. Furthermore, user fees incurred by sponsored research activities or training activities are also payable by most research or training grants, respectively,

enabling institutions to support research and/or training activities without carrying the financial burden for unnecessary infrastructure or capacity. To strengthen public cloud's service capacity for biomedical research and encourage biomedical researchers to explore use of the cloud, the National Institutes of Health (NIH) hosts more than 250 petabytes of biomedical research data and secured group discounts with major cloud service providers [1–3].

Despite the vast potential cloud computing holds in providing broad access to big data and scalable data analytic capacity, most biomedical researchers have not been able to take advantage of this technology for their research. Equally if not more concerning is that biomedical students, including graduate students interested in independent research careers, are lacking opportunities or support to learn skills that enable them to use cloud computing. To better understand the biomedical research community's needs and look for ways to overcome the cloud accessibility barriers, NIH's National Institute of General Medical Sciences (NIGMS) issued [a Request for Information \(RFI\)](#) and organized [a follow-up Workshop](#) for institutional leaders, faculty researchers, postdoctoral fellows, and students of institutions either supported by the NIGMS Institutional Development Award (IDeA) program or that have a historical mission in serving minority students. Through the RFI responses and the workshop, we learned that faculty investigators primarily relied on their institutional servers to access and analyze data, and that, in general, they had not been trained to navigate the cloud. Consequently, biomedical students lacked exposure and mentoring in learning cloud computing skills for research. The consensus among the responses to the RFI and the workshop participants was that new approaches are needed to help democratize cloud computing for biomedical research.

DEVELOPMENT OF THE NIGMS SANDBOX

We hypothesized that a cloud-based learning platform may provide a gateway for biomedical investigators and students to start their journey to the cloud. The platform would include a collection of technologies and/or methodologies, each covered in a free-standing learning module, that is commonly used in biomedical research but not offered through curricula of all institutions. This platform would provide access to many users, enabling them to choose the modules that meet their needs and execute them on-demand. The modules would deliver the learning materials in an interactive format and use appropriate cloud resources for data access and analyses so that users would learn the materials at their own pace while gaining experience and skills in navigating the cloud. The platform would also pilot a cloud account management mechanism to assess costs and explore affordable means to sustain the democratization of cloud computing for biomedical research. The platform is named "NIGMS Sandbox for Cloud-based Learning" <https://github.com/NIGMS/NIGMS-Sandbox> to capture the essence of a learning environment that encourages users to learn through "playing".

The Sandbox was built through a collaboration among NIH scientists and technical staff (from NIGMS, the Office of Data Science Strategy (ODSS), and the Center for Information Technology (CIT)), NIGMS-funded investigators, cloud engineers, bioinformaticians, and project managers from Google Cloud and Deloitte Consulting. To map out a path for the large collaborative team to achieve the aims of this program, two investigator groups, each supported by a supplemental award to an IDeA Network of Biomedical Research Excellence (INBRE) grant,

teamed up with cloud engineers and bioinformaticians from Google Cloud and Deloitte to convert an RNA-Seq and a proteomic analysis curriculum, respectively, into cloud-based workflows. They then developed [instructional videos](#) to introduce the scientific concepts of these technologies or methods and created interactive demos to walk through the workflow step-by-step and to explain the objectives and cloud tools used for each step. To help users solidify their learning, the team also designed practicums that allow users to practice and solve problems. These efforts yielded two functional modules (4,5) through which users can learn these technologies and analyze their own research data (Figure 1). Importantly, the NIGMS-funded biomedical investigators gained software development and cloud engineering skills through working with the software and cloud engineers, which should enable them to develop more cloud-based tools in the future as well as train others to do so (Figure 1).

Following a similar pathway, 10 more investigator teams were funded to work with cloud engineers and bioinformaticians from Google Cloud and Deloitte to develop 10 additional learning modules covering a wide range of topics relevant to biomedical research, ranging from fundamentals of bioinformatics to biomarker discovery. The 12 modules in total (Table 1) collectively represent the first version of the NIGMS Sandbox, which can be accessed publicly on GitHub <https://github.com/NIGMS/NIGMS-Sandbox>. The Sandbox is an open platform that, with continued improvements and researcher contributions, can support many more learning modules to serve the broader biomedical research community. This overarching GitHub repository includes an introduction of the overall design of the Sandbox, documents that support the use of all modules, such as those on how to find a Google Cloud Project and how to spin up a Vertex AI notebook, and links for twelve GitHub repositories (one for each of the modules).

It is important to note that the computational requirements of each module differ. For example, some modules use the base Google Cloud Platform (GCP) machine image with CPU-based machines with Python kernels, while others require R kernels; some modules launch virtual machines from custom Docker images, while others require GPU-based machines. As such, users are encouraged to review and evaluate the required compute environment for each module before using a module. Furthermore, when switching from one module to another, it is recommended to review differences in configuration in order to provision the appropriate environment. The development of each module and corresponding instructions for using it are described in each one of the 12 manuscripts included in this special edition issue of *Briefings in Bioinformatics* (Refs. 4-15).

DEPLOYMENT OF THE NIGMS SANDBOX

A learning platform for biomedical research like the NIGMS Sandbox, while publicly available, is unlikely to be an effective gateway to cloud computing on its own, because the intended users typically do not have an institutionally sponsored cloud account to access the platform and pay for the cloud computing time required for learning. For the platform to reach students and investigators broadly, we need to bring the platform to them and develop account management methods to keep track of usage and pay user fees to cloud computing service providers. Taking advantage of the ready infrastructure supported by the NIGMS INBRE grants,

we utilized a new offering called [NIH Cloud Lab](#) to provide students and investigators supported by these grants access to the Sandbox. NIH Cloud Lab is a program developed by NIH's CIT, and sponsored by ODSS as part of the NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) Initiative. The STRIDES initiative provides NIH and NIH-funded users secure access to cloud services, on-demand technical assistance, open-source GitHub documentation and tutorials, and tracking of the cloud service usage by users. Through NIH Cloud Lab, users have full access to the cloud console, including all offered services by Google Cloud, Amazon Web Services, and Microsoft Azure, with service usage and incurred costs tracked. (Note: NIGMS Sandbox currently only uses Google Cloud.)

How users access the Sandbox through NIH Cloud Lab is illustrated in Figure 2. In this experiment, the cost incurred by each user is paid by the NIH as an investment in NIH's strategic goal of democratizing cloud computing for biomedical research. By October 3, 2023, more than 250 students, staff, and faculty, across three cohorts from NIGMS-funded programs, have used over 800 times the NIGMS Sandbox modules through the NIH Cloud Lab platform. Figure 3 shows the costs for cloud computing time incurred by each user from July 11 to October 3, 2023. The outcome of the experiment, to be concluded in 2024, will inform future improvements of the Sandbox, including refinement and expansion of its' contents, more effective deployment to all students and investigators in higher education and research institutions, and sustainable models including possibly institution-NIH partnerships to pay the cloud computing costs.

CONCLUSION

The ability to access and analyze big data is becoming essential for biomedical research. Cloud computing through commercial cloud services may provide this capacity broadly with manageable costs and significant economies of scale, thus enabling more students and investigators to participate in biomedical research and analyze big data. We hope our effort in developing a cloud-based learning platform will help catalyze the democratization of cloud computing for biomedical research.

ACKNOWLEDGEMENTS

We thank Dave Belardo, Anne Billak, Marcia Price, and Sam Russ of Google; Antej Nuhanovic, Thad Carlson, Madeline Cowan, Kyle O'Connell, and Ross Campbell of Deloitte; Joel Mills of Covalent Solutions; Dana Gaffney of NIH/CIT STRIDES; and Alison Lin and Raphael Isokpehi of NIH/ODSS for supporting the development and deployment of the NIGMS Sandbox, as well as Kyle O'Connell for providing input on the manuscript.

AUTHOR CONTRIBUTIONS

All authors were involved in the conception and execution of the NIGMS Sandbox project. ML, LM, and KA wrote the manuscript, and all authors approved of the final version of the manuscript.

REFERENCES

1. Dahlquist JM, Nelson SC, Fullerton SM. Cloud-based biomedical data storage and analysis for genomic research: Landscape analysis of data governance in emerging NIH-supported platforms. *HGG Adv* 2023; 4:100196.
2. Holko M, Weber N, Lunt C, Brenner SE. Biomedical research in the Cloud: considerations for researchers and organizations moving to (or adding) cloud computing resources. *Pac Symposium Biocomputing* 2023.
3. Navale V, Kaeppler DV, McAuliffe M. An overview of biomedical platforms for managing research data. *J Data Inform Manag* 2021; 3:21–27.
4. Allers S, O’Connell KA, Carlson T, Belardo D, King BL. Reusable Cloud-Based Computing Environments for the Analysis of Bacterial Gene Expression Data Using Bulk RNA Sequencing. In Review.
5. O’Connell KA, Kopchick B, Carlson T, Belardo D, Byrum SD. Understanding Proteome Quantification in an Interactive Learning Module on Google Cloud Platform. In Review.
6. Wilkins OM, Campbell R, Yosufzai Z, Doe V, Soucy SM. Cloud-Based Introduction to BASH Coding for Biologists. In Review.
7. Qin Y, Maggio A, Hawkins D, Beaudry L, Kim A, Pan D, Deng Y. Whole Genome Bisulfite Sequencing Data Analysis Learning Module on Google Cloud Platform. In Review
8. Seaman RP, Campbell R, Doe V, Yosufzai Z, Graber JH. A Cloud-based Training Module for Efficient de novo Transcriptome Assembly Using Nextflow and Google Cloud. In Review
9. Veerappa AM, Rowley MJ, Maggio A, Beaudry L, Hawkins D, Kim A, Sethi S, Sorgen PL, Guda C. Cloud ATAC: A Cloud-based Interactive Pipeline for the Analysis of Pooled-cell and Single-cell ATAC-seq Data. In Review.
10. Nguyen H, Tran B, Nguyen H, Petereit J, Nguyen T. CCPA: Cloud-based, self-learning modules for Consensus Pathway Analysis using GO and KEGG. In Review.
11. Transcriptomics and Epigenetic Data Integration Learning Module on Google Cloud. In Review.
12. Gnimpieba E, Lushbough C, Do T, Ramana V, Rajesh S, Zylla J, Aryal S, Haas S, Doe V, Yosufzai Z, Pan D, Campbell R, O’Connell KA. Biofilm Marker Discovery with Cloud-based Dockerized Metagenomics Analysis of Microbial Communities. In Review.
13. Moctezuma L, Rivera LB, Nouhuijs FV, Orcales F, Kim A, Campbell R, Fuse M, Pennings PS. Using a decision tree to predict COVID case numbers: a tutorial for beginners. In Review.

14. Hemme CL, Beaudry L, Yosufzai Z, Kim A, Pan D, Campbell R, Price M, Cho BP. A Cloud-Based Learning Module for Biomarker Discovery. In Review.
15. Woessner AE, Anjum U, Salman H, Lear J, Turner JT, Campbell R, Beaudry L, Zhan J, Cornett LE, Gauch S, Quinn KP. Identifying and Training Deep Learning Neural Networks on Biomedical-related Datasets. In Review.




Figure 1. Building Cloud-based learning modules

Access & Use the NIGMS Sandbox via NIH Cloud Lab




Figure 2. Access and use of NIGMS Sandbox via the Cloud Lab

First, institutions supported by NIGMS-funded programs identify users among their students, faculty, or staff; Second, the institutions forward requests to the NIH Cloud Lab team; Third, the NIH Cloud Lab team grants each authenticated user an NIH Cloud Lab account, which enables users to access modules in the Sandbox, identify module(s) of interest, and clone them into the Cloud Lab environment for learning.




Figure 3. The daily billing graph provided by the Google Cloud Console serves as a proxy for account usage. Each colored band represents the spending of a unique user for that particular day. Engagement per user has increased through time.

Table 1. List of the 12 NIGMS Sandbox Learning Modules

- 1) RNA-seq Differential Expression Analysis (Ref 4)
- 2) Proteome Quantification (Ref 5)
- 3) Fundamentals of Bioinformatics (Ref 6)
- 4) DNA Methylation Sequencing Analysis with WGBS (Ref 7)
- 5) Transcriptome Assembly Refinement and Applications (Ref 8)
- 6) ATAC-Seq and Single Cell ATAC-Seq Analysis (Ref 9)
- 7) Consensus Pathway Analysis in the Cloud (Ref 10)
- 8) Integrating Multi-Omics Datasets (Ref 11)
- 9) Metagenomics Analysis of Biofilm-Microbiome (Ref 12)
- 10) Introduction to Data Science for Biology (Ref 13)
- 11) Analysis of Biomedical Data for Biomarker Discovery (Ref 14)
- 12) Biomedical Imaging Analysis using AI/ML approaches (Ref 15)