Data Collection, Management, Analysis and Decision Support During COVID-19: A Retrospective from The Ohio State University

Namrata Banerji
banerji.8@osu.edu
The Ohio State University
Columbus, Ohio, USA

Steve Chang
schang@osc.edu
Ohio Supercomputer Center
Columbus, Ohio, USA

Andrew Perrault
perrault.17@osu.edu
The Ohio State University
Columbus, Ohio, USA

Tanya Y. Berger-Wolf
berger-wolf.1@osu.edu
The Ohio State University
Columbus, Ohio, USA

Mikkel Quam
quam.7@osu.edu
The Ohio State University
Columbus, Ohio, USA

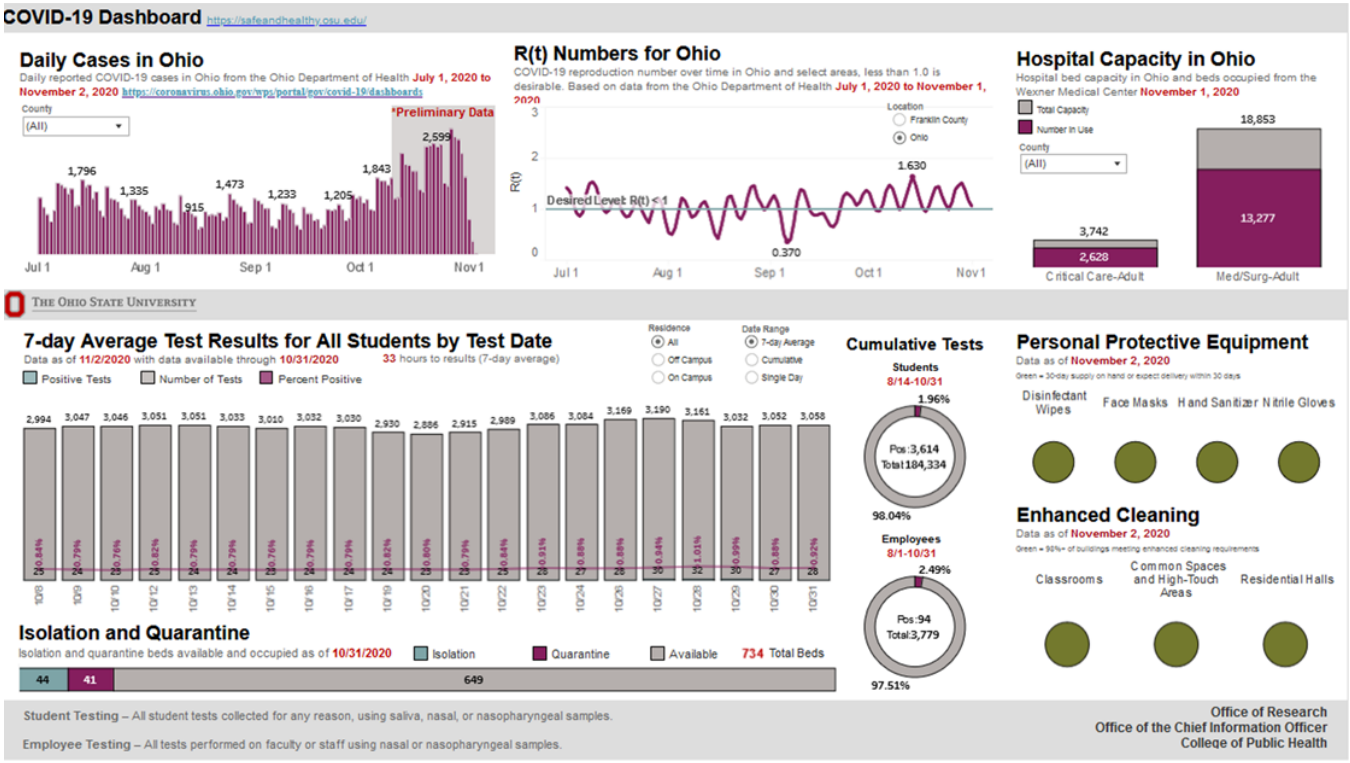

Figure 1. Archived OSU Safe & Healthy COVID-19 Dashboard for November 2, 2020

Abstract

The COVID-19 pandemic has disrupted societies worldwide, challenging governments, healthcare systems, and individuals to adapt rapidly to new circumstances. During this crisis, data has emerged as a critical resource for understanding, monitoring, and mitigating the impact of the disease. This

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

epiDAMIK 6.0: The 6th International workshop on Epidemiology meets Data Mining and Knowledge discovery August 07, 2023, Long Beach, CA,
,

© 2023 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

paper explores the multifaceted role of data usage in the context of COVID-19, focusing on its collection, management, analysis, and application in decision-making.

We present The Ohio State University's data-driven framework for comprehensive monitoring of the COVID-19 pandemic, including the data sources, methods and technologies to find cases, collect primary case data and conduct contact tracing, as well as provide subsequent visualizations and analyses for a large public university system based in an urban center. We discuss the challenges associated with data collection, including privacy concerns, data quality, and the need for harmonization across different sources. We also investigate the roles and limitations of data analysis in supporting intervention choice and implementation strategies amid the complexities of the pandemic as it unfolded. There are also ethical considerations to be weighed in surrounding data usage during the pandemic. Balancing privacy, consent, and transparency is crucial in maintaining public trust and ensuring the responsible handling of sensitive information. We examine privacy-preserving techniques, ethical frameworks, and legal regulations aimed at safeguarding individuals' rights while harnessing the power of data.

In our experience, conscientious data architecture provided a foundation for meaningful ethical applications of data products, which not only to helped mitigate the current crisis, but also can provide valuable insights for better addressing future public health emergencies.

CCS Concepts: • **Information systems** → **Database administration**; • **Applied computing** → *Health care information systems.*

Keywords: datasets, public health, data management, ethics

ACM Reference Format:
Namrata Banerji, Steve Chang, Andrew Perrault, Tanya Y. Berger-Wolf, and Mikkel Quam. 2023. Data Collection, Management, Analysis and Decision Support During COVID-19: A Retrospective from The Ohio State University. In *Proceedings of epiDAMIK 6.0: The 6th International workshop on Epidemiology meets Data Mining and Knowledge discovery August 07, 2023, Long Beach, CA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

The onset of the COVID-19 pandemic in early 2020 was one of the most significant and life changing events for everyone on the planet, from small businesses to entire countries. In case of educational institutions, the indefinite suspension of classes, upending of every traditional aspect of academic and student life, and the transition to virtual education was stressful for students, staff, and faculty alike. The Ohio State University (OSU), a large urban educational institution, undertook a massive policy response to support the continuing function of the university by monitoring and managing the dynamics of the pandemic on and around its campuses. Putting together a workforce of epidemiologists, data scientists, public health policy makers was only the first step of what shaped up to be at least a three year marathon. Data was at the center of the whole process, both as the decision enabler and as the product of many of the contributing efforts. To make data actionable required the work of many teams and several iterations of cleaning, analysis and inference, and visualization. In this paper, we present the overall data-focused aspects of the process, highlighting the achievements and the hindrances, as well as the major takeaways, so that we are better prepared for future public health emergencies or other large scale collective responses. This manuscript, besides serving as a piece of institutional memory, communicates in detail the various obstacles encountered in the handling of the mammoth data for the data science community to be aware of. Among the main takeaways we consider the effectiveness of the data driven approaches for managing the pandemic response, the need for an institutional data infrastructure, and the importance of a well organized team of experts and professionals working together towards a well-defined goal.

We organize this paper into Challenges, Implementation, and Takeaways from our experience with the COVID-19 Comprehensive Monitoring Team at The Ohio State University.

2 Overview

The Ohio State University stood up the Comprehensive Monitoring Team (CMT) to include a framework of support for data driven decisions for pandemic management, including robust case finding (via serial mass administration of individual PCR tests with rapid in-house processing), locally administered isolation of cases, contact tracing and quarantine of close contacts, as well as data integration, analysis, modelling, risk evaluation, policy recommendations, and intervention implementation based upon knowledge derived from individual case management, subsequent viral (genomic) sequencing, large scale syndromic surveillance and evidence of environmental (wastewater and dust) shedding. Here we present the core of the data component of this system that integrated data from various testing centers, conducted daily analyses, and represented data in formats usable by the leadership to support both individual level contact tracing and the university's policy response to the public health emergency. In the coming sections, we discuss the goal of setting up such a system, the implementation pipeline, data sources and some of the challenges and takeaways.

3 Goals

Building and maintaining such a huge framework and employing a whole workforce including faculty, students, healthcare workers consumes university resources at a large scale.

The goals were the result of several rapid iterations of convergent conversations between the university administration and members of the CMT, as well as the consultations with external experts. The specific aims of the data components of the framework were as follows:

- **Tracking the positivity rate.** Positivity rate emerged early in the pandemic as the agreed upon indicator of the state of the population and the basis for comparing different populations.
- **Contact tracing.** Removal of positive and potentially positive cases from the population is key for suppressing the spread of the virus. It was necessary to provide contact information for people who tested positive and to identify and contact their close contacts in order to isolate and quarantine them, respectively.
- **Understanding the micro trends and risks based on events.** To understand the dynamics, the risks, and the implications of the pandemic for various subpopulations it was necessary to provide the ability to zoom in on specific time intervals and subgroups in the data. Examples of the questions asked include: How does fall break or Halloween behaviour change impact infection rates? Is there an increased risk of students in a 4-person suite over a 2-person dorm room? How do the risks associated with in-person classes compare with hybrid or remote classes?
- **Supporting daily policy decisions of a large urban university.** Daily decisions supported by data included the choice of a testing strategy and protocol, transition to hybrid vs online only classes, occupancy in classrooms, vaccination and masking requirements, etc. Having access to the right data was essential.
- **Reducing burden in the wider population.** OSU Columbus campus is a large urban campus with highly permeable boundary in the center of a city. In order to contain the pandemic, the infection rates needed to be controlled both on and around campus. Moreover, the university sought to mitigate the export of infections to communities beyond beyond their campuses. College students mix with the city population and visit their family over breaks, potentially increasing the risk of transmission to vulnerable community members. Recommending and at times requiring testing before breaks was one such measure taken to reduce the burden on vulnerable immuno-compromised population outside the university.

4 Implementation

OSU has 68,000 students, 12,000 of which reside in residence halls in a regular year. During the pandemic, about 8,000 students were in residence halls and were required to test weekly. Additional students, faculty, and staff were testing voluntarily. At its peak, more than 30,000 tests per week were processed.

Multiple teams across Information Technology support, Student Life, Translational Data Analytics Institute, Infectious Disease Institute, University Medical Centers, College of Public Health, and many more were responsible for standing up a system that would be in place for at least the next 3 years. The data environment was a secure and flexible environment that allowed for dynamic data definition and integration of data from at least 56 sources when it was introduced. (The number of data sources grew to over 100 by the end of 2022.) Initial data sources included testing data, student information, residence and permanent addresses, demographics, class registration, residence layout, class and college affiliations, WiFi access point information, and much more. The pipeline is illustrated in Figure 2 and is described very briefly below.

- Primary test data was transmitted into the internal secure data environment via electronic file transfer multiple times a day.
- Additional attributions from other internal OSU systems (IDM, SIS, Student Life, etc.) were preloaded and updated according to the system's change protocol (mostly on a semester basis).
- Test results and internal data were combined into a cohesive reusable dataset (AKA the "gold table").
- Analysts and dashboard builders utilized a common source for all reports and visualizations.
- Data was also sent to Helpspot/Salesforce to support case investigation and contact tracing efforts.

5 Challenges

The novelty, scale, and duration of the recent and ongoing pandemic were major challenges. Data collection, management, and analysis pipelines at this scale had no modern precedent and had to be designed as they were beginning to be used. Moreover, the timelines were drastically compressed and the requirements initially were changing frequently. In addition, some areas lacked data collection and some critical data streams were initially completely absent. Further, we lacked the full understanding of how to translate the questions into data and how to prioritize the variables and the analysis for decision support, particularly in the context of human behavior. Below are some of the issues that posed significant challenges to the team.

5.1 Data cleaning

The data was collected from numerous sources, some of which were manual entries and consequently had unavoidable human error. For example, a table of people in the database had the OSU unique identification (name.#) as the primary key, and the table of test results was supposed to have the same as foreign key. Typographical errors or null values

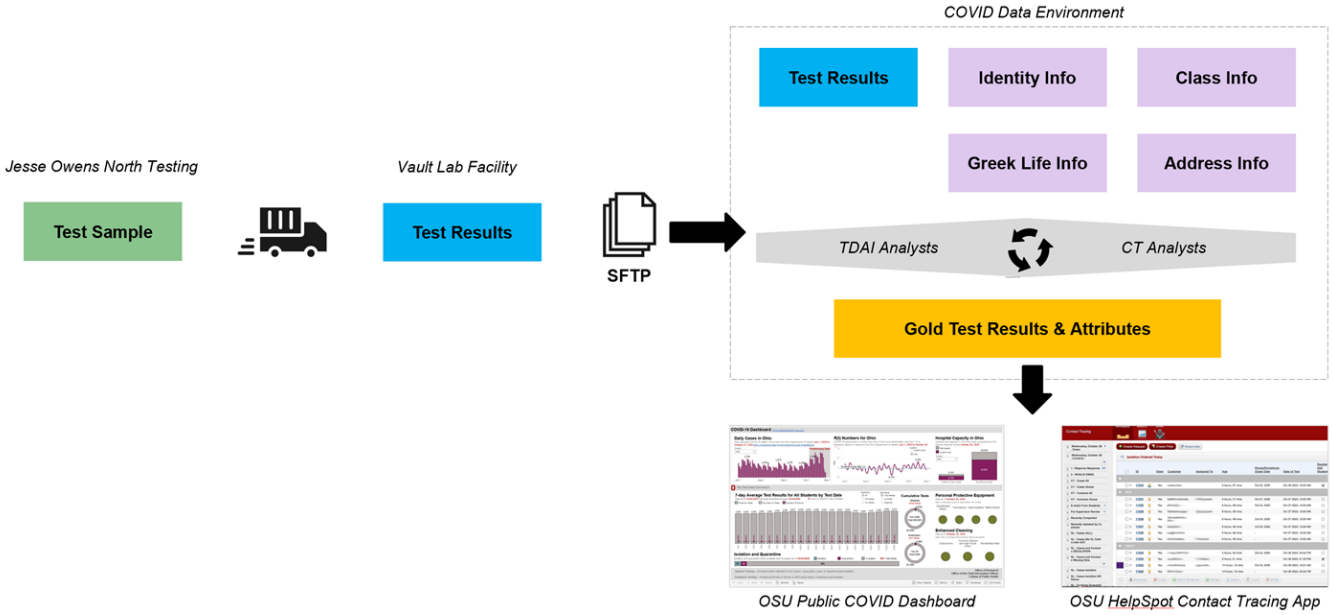

Figure 2. Data flow in the OSU COVID-19 monitoring pipeline.

in this identifier column resulted in a non negligible shift in the summary statistics, given the enormous number of tests conducted. Once the problem had been identified, there was joint effort to clean it up, combining more than four data streams and reducing the number of unidentified tests to a number that would not change the inference. Yet, there were still a few individually unidentifiable entries in the datasets, albeit not enough a number to raise a concern. Minimizing manual entry to data sources can reduce such issues by a considerable amount.

A similar problem was found in the table for employee vaccination records, with clearly wrong dates of doses. While most were due to errors, in some cases, employees were actually part of vaccination trials and had received a dose before any vaccination received emergency use authorization on approval for distribution to the general public. These cases were indistinguishable from the erroneous cases without careful manual investigation and knowledge of the regulatory frameworks and timing of numerous vaccine candidates from all over the world.

These examples illustrate the need for expert data curation, close scrutiny of outputs of any script that consumed these data sources, efforts to minimize manual data entry where possible, and for close collaboration with domain experts at every step.

5.2 Data storage, backup, documentation, and recovery

The volume of data generated by testing mandates as well as voluntary testing required careful consideration of large, yet quickly accessible and continuously backed up data storage. The ability to look up prior data was critical to understanding trends and the dynamics of trends, as well as comparing the outcomes of various past decisions. For continuously changing data, such as the daily updated test data, it is needed to maintain regular snapshots, checkpoints, and versions. This aspect was not fully appreciated initially and required significant efforts to redesign data architecture. The 'gold' data of people, consisting of faculty, staff, students, and everyone else affiliated in some way with the university, updates significantly every semester overwriting previous data in the database (S3 environment). We would save a snapshot of the data every semester, but unfortunately initially the snapshots were taken towards the end of the semesters when students had already started leaving the campus. As a result of this, recently when we wanted to get a time series of positivity rates in residence halls, it was significantly different from the original since we do not have the correct denominator. Recovering this information is possible, but requires integration of other data sources, demanding significant investment of resources, effort, and time. Majority of the people who were part of the university supporting the CMT and were responsible for setting up the system are no longer working at OSU. Moreover, early in the reopening of the university, the primary focus was on managing the pandemic and bringing down the positivity rate, and detailed documentation was not prioritized.

Mid semester migration from one homegrown case data management solution to an outside vendor was a major issue that required major investment and retraining and are

continuing to deal with today from a data and analysis perspective. Roughly from August 2020 to November 2020, we had our positive test (case) data ingested and case investigation/contact tracing notes stored in a secured instance of a HelpSpot database integrating in some instances with REDCap surveys and pushing out to several communication platforms, but later we shifted to a Salesforce Health Cloud build, which assisted with future testing data variations, vaccine information, as well as some automatic reminder communications. The data had been migrated from the old table to the new one in theory, but in part user generated heterogeneity, as well as version control issues in the HelpSpot source data meant there continued to be gaps in the data ingested by Health Cloud (Salesforce) which do not have simple workarounds for analysis of all variables. We maintain several tables for the test information storage, but there are inconsistencies across those tables. More than one tables exist mainly because we derived simpler versions of tables with many columns that are not relevant for day-to-day analysis. One of the (intermediate) mother tables recently had one of its very important columns (the test specimen collection time/date column) dropped from an integration during an update, and it should have been okay to just look it up in a derived or other related testing table had there not been major differences in the number of entries in the others.

5.3 Data governance and ethical considerations

The university has a complex set of data governance regulations as do individuals' private health information, whether used in the healthcare or public health applications. While special authorization was granted to use some of the data in the pandemic emergency, security and privacy remained strict requirements. Each team member had training in handling secure and private data.

In addition to the standard data governance issues, dealing with the high resolution personal data has its own set of ethical issues. Ultimately, the main question was: what is the benefit of using a particular data source or performing a particular analysis and would it change the decisions or the pandemic dynamics? If so, was it necessary to use individual and identifiable data for decision making or could aggregate or coded information have similar utility? For example, while it is within the rights of the university to use the WiFi access point information to "follow" an individual or to understand who is within the same room, such information has a high 'icky factor' and should be used sparingly. Moreover, while initially it seemed that WiFi data would provide a good proxy for contact tracing, it turned out that the resolution of the data did not correspond well to the physical definitions of a contact. Ultimately, it was decided to use WiFi data in aggregate to assess population movements rather than individuals' proximity to other individuals. For example, WiFi data was used to estimate the number of students leaving campus over the weekend or the number of students present in an "in person" classroom. Moreover, the aggregate trends proved to be much more robust than the individual-based analysis and were significantly less time consuming. Additionally, adherence to the current applicable statutory guidelines for case investigation, subsequent case management, and/or contact tracing may require some variation depending upon individuals' occupation, travel history, personal risk factors, immunocompetence, vaccination status, which could include certain specific preexisting conditions, medications, clinical care received, viral (variant/sub-variant) lineage, and/or disease severity. However, specific individuals' health information related to their experience with COVID-19 would largely not meaningfully determine macro-level prevention policy or interventions in the university context independently from aggregate trends and information in the wider public health policy guidance, which are separately informed by individuals' public health, laboratory testing and clinical health records. Therefore, particularly those sensitive individual level data, especially health data were collected and subsequently shared only to the extent they would have 'meaningful use,' within the data user groups' spheres of control, stated goals, and purview.

6 Takeaways

7 Behavior over analytics

The main takeaway of our data-supported pandemic monitoring framework is the same as the main takeaway for dealing with the COVID-19 pandemic world-wide: ultimately, the main determinant of the success of the system hinges on modifiable human behavior, rather than the sophistication of the analysis. No improvement in the accuracy of the analysis of the effect of masking in a given setting (i.e. library, classroom, laboratory, or healthcare setting) is meaningful if people would not (continue to) comply with an indoor mask mandate. Similar limitations became apparent with both pharmaceutical and non-pharmaceutical interventions, even as evidence increasingly substantiated benefits and new sub-variants emerged, populations' apparent risk tolerance grew and spread.

7.1 Communication is key

Working with a team this large, with people from vastly diverse backgrounds, communication between the teams becomes an essential component. A major part of the analysis was being carried out by graduate student employees, who were sometimes not aware of things like floor structure in dorms, testing protocols, vaccination mandates, etc., which were important analysis components. Similarly, the modelling team was involved in building risk models, models for testing strategy development, etc. that relied on domain knowledge outside of mathematics or computer science. Clearly, experts in every relevant domain (epidemiology,

public health, student residence life, university logistics and operations, etc.) need to be constant partners in the analysis.

7.2 Equity Considerations and singling out demographic groups

When patterns appear to be emerging within specific groups or sub-demographic, there may be an equity oriented opportunity for targeting or strengthening an intervention but there may also be a bias in the observed signal. One group may in fact be more often in situations involving exposure to infectious persons, or engaged in more risky behavior than others, as we occasionally discovered from data analysis. However, available policy level changes may not have been feasible solutions and were not always ultimately enacted. What we started to see in the data raised questions on ethics and trustworthiness of data enabled interventions, without context or corroboration. Some solutions aimed to address one groups perceived or real deficiency in access to resources or excessive exposure could foster stigma, or loss of other resources in unanticipated ways. After careful consideration, it was agreed that singling out a group was often not enough of a value addition or could do more harm than good. In some cases, trends observed initially in one population or group were indicative of larger trends that could be addressed by policy shifts relevant to the whole community, which would address both the observed inequity and mitigate for known unintended consequences.

7.3 Micropatterns significant, but not usable in hindsight

The reflections on the decisions made over the course of three years showed that the micropatterns and the microtrends observed in the data had little to no effect on those decisions. Observations that a certain subgroup engaged in activities that increased the risk of the spread of the infection did not prompt the authorities to take measures to shut down those activities in many cases because it was either not cost effective or unethical to do so. These data nuances did provide information but it was not actionable. In retrospect, however, the information's main utility was in the fact that no single critical subgroup was the key to the solution. The scale of the phenomena did not lend itself to a single pathway of solution or a single target group. Patterns that we learned in settings like an early long term care facility were also observed later in dorms, sorority and fraternity houses and athletics teams and they led to better population level responses. A good example would be the limitations of certain kinds of tests for transmission suppression. The Big10 testing program involved daily testing of athletes during their competition season, given that team members were often unable to mask and physically distance in some sports. Unfortunately, when transmission started to increase rapidly in late autumn 2020 as sports teams re-started their compressed seasons, even daily testing with rapid results was insufficient to suppress transmission, largely because the particular test used did not detect all infectious individuals immediately. By the time one tests positive on an antigen test, like those in use at that time, a person may have already been infected and infectious for a few days, meaning potentially exposing others and continuing transmission chains. Antigen tests are useful for rapid diagnosis particularly when symptomatic but are not always ideally suited for early enough detection to reduce spread in a serial testing model. OSU opted for developing and deploying routine swift, minimally invasive testing, eventually with its own PCR testing lab and although they were not as fast as antigen tests, the average turnaround time was less than 24 hours. This was a scenario where tracking a micropattern gave us really good information of what and how we should be testing.

7.4 Data infrastructure

The overall data infrastructure consists of cyberinfrastructure (compute, storage, networking, cloud and web services), information infrastructure (data and metadata management, search, archiving, cataloging, and digital services), and analytics infrastructure (data integration, harmonization, and analysis). The large volume of data collected, collection rate, distributed team setting, potential errors, inconsistencies, and variations in reporting standards, and changing objectives all strained and challenged existing data infrastructure at OSU and necessitated expansion of that infrastructure. Moreover, COVID-19 management provided a great case-study and emphasis on the fact that data infrastructure integrates cyber-, information, and data services infrastructures through **human infrastructure**. Building the human infrastructure is both the most critical aspect and the hardest to implement of any data infrastructure.

7.5 Scale

The sheer volume of COVID-19 data generated from testing and vaccination overwhelmed existing data management systems of the university as well as the state. Scaling up data infrastructure and analytical capabilities to handle large-scale data collection and analysis proved to be a significant challenge, but one that can definitely be overcome.

8 Conclusion

We hope that the COVID-19 pandemic was a one-off unique event, never to be repeated. Yet, we should be prepared to respond to a similar event by learning from our experience. We hope that the OSU CMT work presented here can serve not only as a blueprint, but as a good foundation for aspects that should be considered in using or modifying the blueprint and pitfalls to avoid, should the response at this scale be ever needed.

Acknowledgments

We would like to acknowledge the work of many people who have contributed to the effort of enabling the data driven approach to monitoring and managing the COVID-19 pandemic at the Ohio State University: the entire Comprehensive Monitoring Team, contact tracing Team, CMT student analysts, dashboard developers, and the OTDI team, including D& A data engineers, data governance team, network administrators, and enterprise security.