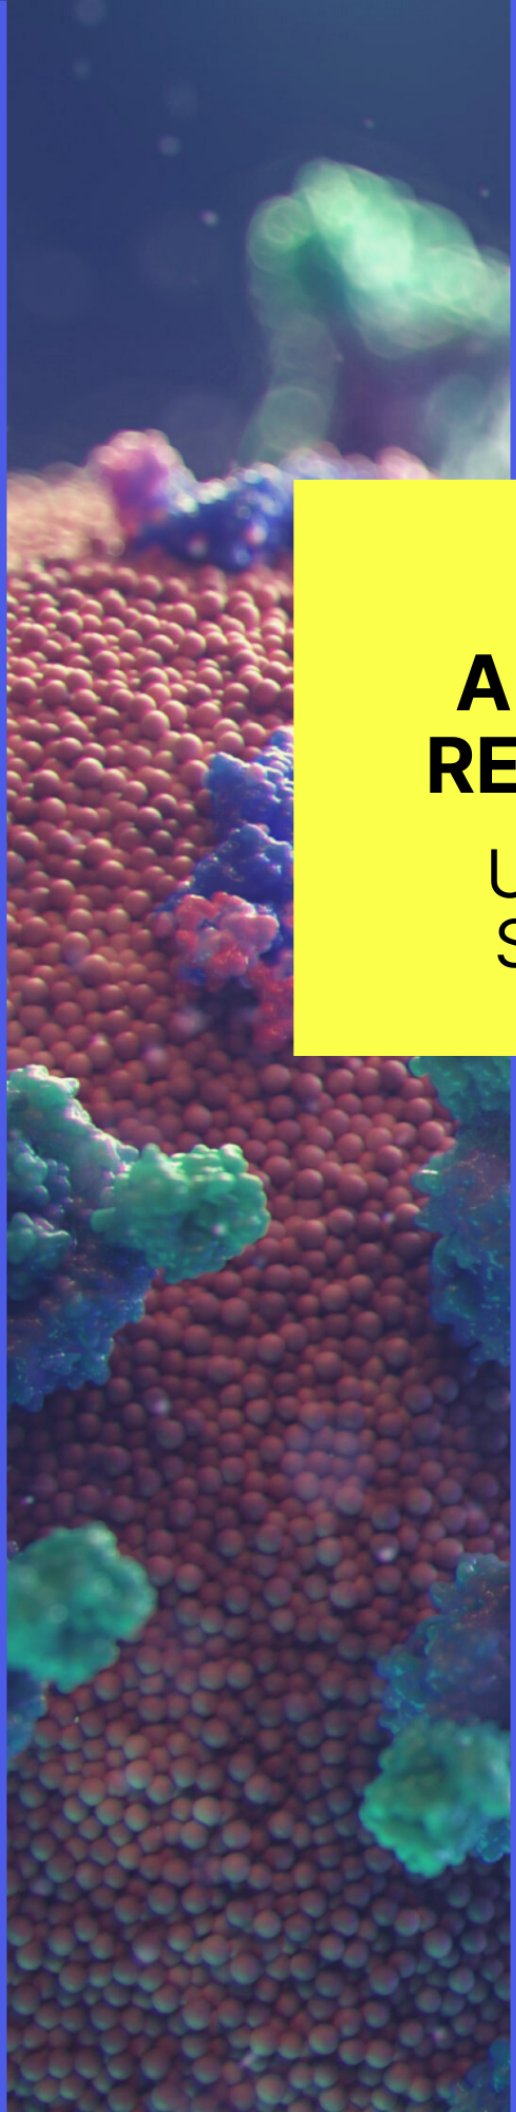

OCTOBER 2021

THE FUTURE SOCIETY

AI-POWERED IMMEDIATE RESPONSE TO PANDEMICS

Update on Living Repository: Summaries of Top Initiatives

PRODUCED BY:

The Future Society (TFS)

IN COLLABORATION WITH:

Global Partnership on AI (GPAI) AI & Pandemic Response Subgroup

Contributors

The Future Society
Samuel Curtis, Project Manager
Adriana Bora, Analyst
Rui Wang, Senior Technical Expert
Nicolas Miailhe, Research Director
Sacha Alanoca, Research Advisor

Acknowledgments

GPAI AI and Pandemic Response Subgroup Co-Chairs
Alice Oh
Paul Suetens

GPAI AI-Powered Immediate Response to Pandemics Project Co-Leaders
Michael O'Sullivan
Paul Suetens

GPAI AI-Powered Immediate Response to Pandemics Project Steering Group
Cyrus Hodes
Michael O'Sullivan
Alan Paic
Margarita Sordo
Paul Suetens

CEIMIA
Jacques Rajotte
Edward Teather

OECD AI Policy Observatory
Karine Perset
Francesca Sheeka

Table of Contents

Introduction

For the 2021 Summit, the Immediate Response Committee of the AI & Pandemic Response Subgroup commissioned The Future Society (TFS) to develop an updated and upgraded catalogue of AI initiatives with potential to combat COVID-19 and other future pandemics, and transform it into a Living Repository. The research builds upon last year's report, *Responsible AI in Pandemic Response*.

An impact assessment has been conducted yielding a subset of initiatives that show promise in terms of their potential to scale, for the purpose of identifying those that could benefit most from partnership to deliver on their promise.

The research began in collaboration between The Future Society, the OECD, and GPAI by applying the OECD Framework for the Classification of AI Systems to classify AI initiatives based on their technical characteristics. Those associated with the development of AI systems created or repurposed to aid in COVID-19 responses were invited to complete a survey shared in a public announcement. A total of 66 initiatives were identified via TFS desktop research and survey responses.

The Immediate Response Committee then worked with TFS to build upon the OECD Framework to develop a more impact-focused set of criteria, including:

- Background of the initiative (name, sources, objective/purpose)
- Origin (including organization(s), locality)
- Categorisation (type of approach / AI method)
- Scope (domain, target users/operators and beneficiaries, geographic coverage)
- Data (description of the dataset in use including demographics, target population, size, collection timeframe, any public access links)

The 66 AI systems have been classified using this framework to create the Living Repository. This is being shared in an open 'work-in-progress' format in reflection of the immediate needs of the pandemic for those that may find it useful as a resource.

Using these classifications, the Immediate Response Committee has conducted assessments of initiatives' intrinsic scalability and their potential to mitigate the current and future pandemics, to narrow the 66 identified initiatives into a shortlist of 26, of which 11 have initially been selected as candidates for potential partnerships with the AI & Pandemic Response Subgroup and GPAI more widely.

In this document, we are pleased to share summaries of the initial selected 11 initiatives. They include AI systems that have been trained to:

- predict the distances and angles between pairs of proteins' amino acid residues;
- determine the effectiveness of non-pharmaceutical interventions (NPIs) on COVID-19

- identify individuals who are at the greatest risk of heightened vulnerability to COVID-19, based on individuals' pre-existing medical conditions;
- provide users with personalized daily COVID-19 "risk scores" associated with regular activities;
- organize both structured and unstructured COVID-19 data into a knowledge graph that can be navigated and queried to retrieve information;
- provide a country-level risk modeling framework intended to assist the government and individuals in making informed decisions;
- quickly and accurately detect the presence of COVID-19 in thoracic CT scans;
- model the spread of COVID-19 based on the prevalence of mask-wearing in a population;
- identify, track, and analyze events associated with COVID-19 via mentions on online news articles and social media posts;
- aggregate and clean various sources of US pandemic-related raw data to produce COVID-19 "indicators" for "nowcasting" (situational awareness) and short-term forecasting;
- allow users to view current occupancy rates of hospitals across the US and recommendations for intra-state patient transfers based on current occupancy rates.

We now look forward to the progress being discussed at Summit 2021, and will publish an update including a revised Living Repository and an additional 15 descriptive summaries in early 2022.

Our intention is that the analysis will then be used to help inform the Immediate Response Committee's partnerships approach in 2022, but should also provide a useful tool and model for the critical evaluation of AI initiatives within the ongoing and in future pandemics.

Researchers tried to make these summaries comprehensive and accurate with information that was publicly available, but we acknowledge that they may contain errors or details that are outdated. If you are a developer of one of these initiatives and would like to correct or add information, please contact the International Centre of Expertise in Montreal on Artificial Intelligence ("Centre d'expertise international de Montréal en intelligence artificielle"; CEIMIA) at info@ceimia.org.

Summaries

AlphaFold

AlphaFold is a deep learning model trained to make predictions of the distances and angles between pairs of proteins' amino acid residues, which are then used to construct accurate predictions of the shapes of proteins [1]. AlphaFold is a product of DeepMind, an artificial intelligence research laboratory based in the United Kingdom. In August 2020, AlphaFold's developers shared the predicted structures of six under-studied proteins associated with SARS-CoV-2—created with AlphaFold—with the scientific community.

In their blog post announcing the release of COVID-19-related protein structure predictions, DeepMind researchers shared their desire to contribute to the scientific community's interrogation of the functions of viruses and for their model to serve as a hypothesis generation platform for future experimental work in developing therapeutics [2].

Since their first announcement in 2018, the AlphaFold model has gone through numerous stages of development. When announcing the release of the six SARS-CoV-2-associated protein structure predictions, DeepMind referred to their model entered in the 13th biennial Critical Assessment of Protein Structure Prediction (CASP13), dubbed "AlphaFold v1.0" [2]. This version of AlphaFold was trained on publicly-available data consisting of approximately 170,000 protein structures from the professionally-curated Research Collaboratory for Structural Bioinformatics Protein Data Bank (RCSB PDB) [3] and large, open-access databases of protein sequences derived from genome sequencing projects, such as UniProt [4]. The sequences of proteins associated with SARS-CoV-2 were also obtained from UniProt.

AlphaFold v1.0's model consists of two stages: (1) a two-dimensional dilated convolutional residual network that takes an amino acid sequence and, using training data, outputs the prediction of distance and torsion between amino acid residues; and (2) a differentiable model that performs gradient descent using the output of the first stage to optimize the 3-dimensional shape of a protein towards its lowest energy potential (in other words, closest to equilibrium) [1].

Following the release of the structure predictions associated with SARS-CoV-2 in August 2020, in July 2021, DeepMind published a validated, redesigned version of AlphaFold's earlier model ("AlphaFold v2.0"). This iteration replaced the convolutional neural network in the prior model with a transformer-based architecture—the "Evoformer"—which treats the prediction of protein structures as a graph inference problem in 3D space, processing inputs through repeated layers of a neural network block to produce an array that represents the inputs in a lower dimension. This new model had the best performance by a significant margin at CASP14 held in 2020 [5].

Upon the publication of the paper describing AlphaFold v2.0's architecture, Deepmind also openly released its source code, trained weights, and an inference script to the research community [6]. It also partnered with European Molecular Biology Laboratory's European

Bioinformatics Institute to release the AlphaFold Protein Structure Database, which includes structural predictions of all of the 20,000 proteins in the human proteome, as well as those from other biologically significant organisms, such as E. coli, yeast, drosophila, and mice.

The developers identified limitations in predicting parts of the human proteome, such as proteins that are unable to be accurately modeled with single-chain structure prediction, and must be modeled in complex or in cellular milieu. They also note a bias towards the human proteome for health and medicinal research, while other biologically, medically, or economically important organisms are underrepresented.

The open source nature of this initiative's source code and the model's noteworthy accuracy in predicting protein structures suggests a high potential for impact in pandemic response, as protein structure prediction is critical for understanding viral biology and pharmaceutical design.

More information can be found in the Living Repository.

Bayesian hierarchical semimechanistic model

This initiative developed a Bayesian hierarchical semimechanistic model to determine the effectiveness of non-pharmaceutical interventions (NPIs) on COVID-19 transmission [7]. The initiative was developed by researchers at the University of Oxford, Australian National University, the Quantified Uncertainty Research Institute, Harvard University, the University of Bristol, the University of Manchester, the London School of Hygiene and Tropical Medicine, the London School of Economics and Political Science, the University of Cambridge, Tufts University, and Imperial College London. The initiative is presented in an academic paper published in *Science* in February 2021 [7].

The rationale behind this initiative was to provide an alternative to simulation studies, which tend to make strong assumptions that are relatively difficult to validate, by developing a data-driven, cross-country model that compares national interventions to the subsequent numbers of cases or deaths within those respective regions.

NPI data were collected across 41 countries from January 22nd to May 30th, 2020. To mitigate errors, all NPI data were entered independently by two of the authors, using primary sources, and then manually compared with two public datasets: the Epidemic Forecasting Global NPI [8] and the Oxford COVID-19 Government Response Tracker [9]. Data on confirmed COVID-19 cases and deaths were obtained from the COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University [10]. To prevent bias, data were pre-processed by neglecting COVID-19 case numbers before a country had reached 100 cases, and fatality numbers before 10 deaths [7].

The model presented in this paper was built upon the semimechanistic Bayesian hierarchical model developed by Flaxman *et al* [11], which estimated the effects of NPIs on COVID-19 transmission in Europe. Similar to Flaxman's approach, this model used COVID-19 case and death data to make a 'backward' inference of the number of new cases for each country, which was then used to infer daily reproduction numbers [7]. The reproduction number and the occurrence of NPIs were then used to estimate NPI effects. To account for cross-country variations in effectiveness, reporting, and fatality rates, as well as uncertainty in the generation interval and delay distributions, researchers utilized a Markov chain Monte Carlo (MCMC) sampling algorithm [12] to infer posterior distributions of each NPI's effectiveness.

The researchers found that NPIs demonstrated highly consistent trends across countries [7]. For instance, closing both schools and universities was consistently highly effective at reducing COVID-19 transmission, as was banning gatherings of 10 people or fewer, whereas targeted closures of face-to-face businesses with a high risk of infection, such as restaurants, bars, and nightclubs, had a small-to-moderate effect. They furthermore found that when most NPIs were already in place, stay-at-home orders had only a small additional effect; thus, by using effective interventions, some countries could effectively control COVID-19 spread while avoiding stay-at-home orders.

The researchers point out numerous limitations of their approach, such as an inability to factor in country demographics, regional differences in interpretations or implementations of NPIs, and a lack of data on some NPIs not captured in this study, which may restrict the feasibility of scaling up the tool [7]. Some of these limitations may resolve with time, however, if NPIs were to become more standardized across larger geographies, and as more COVID-19 case and death data becomes available before and after the implementation of various NPIs.

More information can be found in the Living Repository.

C-19 Index

The C-19 Index is an open source, AI-based predictive model designed to identify individuals who are at the greatest risk of heightened vulnerability to COVID-19, based on individuals' pre-existing medical conditions [13]. The C-19 Index was developed by researchers at ClosedLoop.ai, a private healthcare software company, and one researcher affiliated at Healthfirst, a New York-based health insurance company. A research article describing their models was uploaded to medRxiv in March 2020 [14] and was published in the Journal of Medical Artificial Intelligence in December of 2020 [13].

In their research article, the researchers note that identifying who is most vulnerable to COVID-19 complications or death is not straightforward; however, patterns that were emerging in data from Wuhan and the US (in early 2020) suggested that the risk of death increased with age, for those who have diabetes, heart disease, blood clotting problems, or have shown signs of sepsis. Researchers believed that building predictive models based on these known risks could be useful for outreach campaigns targeted to those most at-risk of severe COVID-19 complications [13].

Researchers used data from two different datasets to train their models: the Center for Medicare & Medicaid Services Limited Data Set for 2015 and 2016 [15], and a medical claims dataset containing 2.5 million Healthfirst insurance beneficiaries. Each dataset represented different US demographics: the former contained data for those over the age of 65 or disabled who receive Medicare, while the latter contained data from overall healthier adults enrolled in Medicaid. Cohorts were created from each data set, and then the resulting cohorts were combined, such that the combined cohort had an age profile consistent with the overall US population [13].

Three different models, which output a person's "C-19 Index" score—the percentile risk of near-term severe complications from an upper respiratory infection, were then trained on the combined cohort's data: (1) a "survey risk factors" logistic regression model that outputs a person's percentile risk score based on responses to a web-based survey; (2) a "diagnosis history model," which train gradient-boosted trees in a time-delayed fashion, allowing the model to use current claims data by simulating the 3-month delay in claims processing that usually occurs in practical settings; and (3) an "expanded feature model," a model built within ClosedLoop—a software system for creating machine learning models—that uses additional engineered features from peer-reviewed studies (not disclosed in their publication). The key differences between each model is the number of features each employs, and thus, their ease of implementation.

The validation dataset contained 14,000 COVID-19 cases in New York City from February 2020 until mid-May 2020. The logistic regression used the fewest features and delivered the lowest performance, with an AUROC (Area Under the Receiver Operating Characteristics—a measurement of a model's ability to distinguish between classes) of .731. In comparison, both the diagnosis history model and the expanded features models obtained AUROCs of .810.

The C-19 Index has already been utilized by at least two healthcare organizations: Medical Home Network, an Illinois-based accountable care organization [16], and Healthfirst, the aforementioned New York-based health insurance company . However, authors note several limitations of their study which impact the feasibility of using this approach on a larger scale: no real COVID-19 cases were used in the model's training, the approach relied on claims data instead of clinical data, and data excluded those under 18 years of age. Therefore, moving forward, possible technical enhancements could be to validate the proxy outcome and determine their validity based on COVID-19 data, to build models on COVID-19 vulnerability on COVID-19 data (without having to use other upper respiratory diseases as proxies), and to test on data of those under 18 years of age.

More information can be found in the Living Repository.

COVI

COVI is a research project that culminated in the development of an AI-enabled contact-tracing mobile application. This application aimed to provide users with personalized daily COVID-19 “risk scores” associated with regular activities (such as taking public transportation and socializing with friends). These scores were based on users’ demographic and health profiles [17]. The project commenced in May 2020, led by researchers at Mila (a research institute specializing in artificial intelligence), Canada, with affiliations including the University of Ottawa, Universite de Montreal, The Alan Turing Institute, University of Oxford, University of Pennsylvania, McGill University, Borden Ladner Gervais LLP, The Decision Lab, HEC Montreal, Max Planck Institute, Libeo, and the University of Toronto.

The primary objective of the mobile application was to help members of the general public make informed, risk-reducing decisions in a manner that preserved individual privacy [18]. COVI extends beyond a contact-tracing mobile app by combining contact-tracing information with other user data (e.g., user demographics, health information, symptoms) to predict daily personal risk factors for each user. In addition, COVI translates these personal risk scores into recommendations based on public health guidelines. Finally, collected data is used to define epidemiological models and intervention simulations, which could then be shared with public health officials to help them preempt resurgence of the virus and inform reopening strategies [18].

COVI was developed to utilize a variety of data from users, all of which would be obtained by consent. Upon opening the app for the first time, users would be provided with an overview of how the app works and the privacy implications of sharing data with COVI. It then would ask for consent for the collection, use, and disclosure of IP-based geolocation history, random “contact” IDs (generated when a phone is within 2 meters of another phone with COVI installed), and users’ current risk levels — all necessary for the app to function properly [18]. If a user were to start presenting symptoms or be diagnosed for COVID-19, they could report accordingly. Then, contacts made with that user within the past 14 days would be notified, and the symptoms/diagnosis would be factored into the computation of the contacts’ risk scores [18].

COVI also asks for consent for collection and use of data pertaining to a user’s age, sex, health conditions, active symptoms, ongoing relevant behavior, coarse geographical location, and app analytics information; all of these data (except analytics information) would be fed into the application’s risk assessment function, which would compute locally on the user’s device [18]. Data remains on a user’s device unless the user opted to allow COVI Canada to receive encrypted, pseudonymized data packets and heat-map information (in aggregated form), which would be used by COVI’s underlying ML model and assist in epidemiological research by government or other third parties [18].

Once collected, data can be used to train deep learning ML models to predict contagiousness risks, and to fit an epidemiological model [18]. COVI deploys architectural scaffolding for deep learning around a Transformer architecture, which draws upon information pertaining to

demographics, behaviors, health conditions, symptoms, and contact with other users [19] to dynamically refine the ML model. In addition, the data shared by the app users also enables AgentSim, an agent-based simulator that offers flexibility in designing contact-tracing and epidemiological simulations, to identify new patterns and specific parameters (such as distance, sex, and age) to model how the virus spreads [20].

Design and development of COVI had been finalized, and Mila was aiming for it to be endorsed for use by the Canadian government, in early June of 2020. However, in spite of putting forward a demonstrated effort towards building a “privacy-conscious app” [17] and fostering public trust by explaining the rationale behind their app design decisions in their white paper, the Canadian government decided to endorse a different application which collected less personal data, citing privacy concerns by provincial and territorial leaders [21].

COVI’s code and documentation remain accessible on Mila’s website (hosted on GitHub [19] and arXiv [18]) as open-source, with a non-exclusive and royalty-free license, “should [others] wish to deploy an AI-enabled health app inspired by our approach” [17]. The open-access nature, human-centric privacy protocols, and consensual use of encrypted, pseudonymized user data suggest that this tool has a high potential to scale to other geographies.

More information can be found in the Living Repository.

COVID-19 Hospital Capacity Management

COVID-19 Hospital Capacity Management is an publicly-accessible online dashboard that allows users to view current occupancy rates of hospitals across the US and recommendations for intra-state patient transfers based on current occupancy rates. It also provides an interactive tool to modify model parameters (eg. patient type, transfer budget, transfer distance threshold) to obtain more customized recommendations within any state or hospital system within the US [22]. The tool was developed by a team from Johns Hopkins Center for Systems Science and Engineering and Malone Center for Engineering in Healthcare, and is affiliated with the Center for Data Science in Emergency Medicine and the Department of Civil and Systems Engineering at Johns Hopkins University [23]. The initiative was first publicly announced on October 27, 2020 [24] and the team's first related academic preprint was uploaded to arXiv on November 6, 2020 [25].

In their preprint, the developers state that the motivation behind their effort was to minimize resource shortages, which would, in turn, improve the overall quality of patient care and prevent early discharges and cancellations of elective surgeries [25]. They note a few instances of patient transfers occurring ad-hoc in the COVID-19 pandemic, but remark that treating this issue at a more protracted, system level—across hospitals, counties, and states—will spur more efficient resource use. They also recognize the alternative approach of hospitals individually and reactively responding by creating surge capacity, but point to studies suggesting that such an approach can lead to a reduced quality of care compared to hospitals working in coordination to make use of existing resources.

For data on past hospital occupancy and COVID-19 hospitalizations, the dashboard relies on statistics provided by the US Department of Health and Human Services [26]. To make future projections, the team uses the US Center for Disease Control's county-level forecasts of COVID-19 cases [27]—an ensemble of models from many forecasting teams, which the researchers behind COVID-19 Hospital Capacity Management then disaggregate to the hospital level.

To make recommendations for patient redistribution, the researchers constructed a series of linear optimization (linear program and mixed-integer linear program) models to solve a multi-period demand problem: "given a set of nodes and time periods, with nominal demand (ie. COVID-19 patients) at each node during each period and fixed capacity at each node, determine the optimal quantity of demand to transfer between each pair of nodes during each time period" [25]. To better reflect constraints in reality, they extended the model by adding a number of parameters to each node, including the type of patient (ICU vs acute care), the per-transfer hospital budget, the total transfer budget, percentage of capacity reserved for COVID-19, transfer distance limits, and lengths of stay, among others. On the COVID-19 Hospital Capacity Management site, the output of the models—recommended intra-state patient transfers—are presented on dynamic graphics for each US state. In their paper, the researchers also presented an analogous multi-period method to model critical redistribution (rather than patient transfers), however, this was not presented on their dashboard.

The online dashboard remains accessible and up-to-date with data updated on a weekly basis, and the source code publicly accessible on their GitHub repository [28]. Limitations to scaling up this approach include the accessibility and quality of data pertaining to present hospital capacity, the aforementioned related parameters, and the accuracy of forecasts of COVID-19 hospitalizations within a given geographical region.

More information can be found in the Living Repository.

COVIDcast

COVIDcast is a site that aggregates and cleans various sources of US pandemic-related raw data to produce COVID-19 “indicators,” which are intended to inform decision making by a broad range of users—public health authorities, the healthcare industry, the public and private sectors, epidemiologists, data journalists, and the general public [29]. COVIDcast was created by the Delphi Research Group at Carnegie Mellon University—one of the US Centers for Disease Control and Prevention’s (CDC) Influenza Forecasting Center of Excellence [30]—with support from Amazon, the US CDC, Change Healthcare, the US Defense Threat Reduction Agency, Facebook, Uptake, Optum, and Google.org [29]. In addition to data from the aforementioned sources, the group also uses their own indicators to create forecasting models at the US state and county level. COVIDcast was launched in May of 2020 and its first academic paper was published on June 25th, 2021 [31].

The Delphi Research Group’s motivation is to develop the theory and practice of epidemic tracking and forecasting. In doing so, they procure data streams that reflect epidemic (or pandemic) activity, define relevant indicators, and make them available for public consumption. They then use these indicators for “nowcasting” (situational awareness) and short-term forecasting [29].

COVIDcast collects and shares a wide range of COVID-related data, categorized into “public behavior,” “early indicators,” and “late indicators.” Public behavior includes: the frequency of bar and restaurant visits, obtained via SafeGraph; the degree to which people wear masks and are willing to get vaccinated, obtained via a survey administered through Facebook; and Google searches pertaining to COVID-19 symptoms, obtained from Google. Early indicators include: COVID-19-related doctor visits provided by partnering health system organizations, such as Change Healthcare; and COVID-19 symptoms present in individuals or communities, obtained via Facebook surveys. Late indicators include: COVID-19 antigen test positivity rates, provided by Quidel; hospital admissions data, provided by partnering health system organizations; and COVID-19 cases and deaths, provided by Johns Hopkins University and USAFacts. Even though COVIDcast did not launch until May of 2020, they were able to collect data retrospectively beginning in February of that year, and they continue to collect and update data on a nearly daily basis.

The Delphi Research Group aggregates, cleans, and then displays these data on the COVIDcast site throughan interactive dashboard. On the dashboard, users may browse indicators’ daily trends, or explore correlations between indicators, at the US state or county level. Using their own indicators, the Delphi Research Group implemented time series models at the state and county levels. At the state level, a basic autoregressive time series model was developed using only two features: COVID case and death counts. At the county level, a quantile regression model was developed using four features: case counts, self-reported symptom rates, doctor’s visits rates, and population [32]. These models are included in an ensemble model developed in collaboration with the Reich Lab at the University of Massachusetts, and accessible at the COVID-19 Forecast Hub [33].

The COVIDcast dashboard is updated with new data nearly daily, and indicators are publicly accessible via an API [34]. The Delphi Research Group also tracks and reports (via their API) revisions made to datasets [35]. As of July 2021, the API was reported to have been accessed by "thousands of users every day, requesting hundreds of thousands of pieces of information" [36]. COVIDcast's indicators have reportedly been used by numerous organizations responding to the pandemic, including COVID Act Now, COVID Exit Strategy, DeepCOVID, and the Institute for Health Metrics and Evaluation (IHME). Furthermore, the COVID-19 Forecast Hub's ensemble model, which integrates Delphi Research Group's forecast models, serves as the basis of the US CDC's COVID-19 forecasting communications [29].

In terms of technical scalability, some COVIDcast indicators may be difficult to repurpose for geographies outside of the US. For instance, some indicators that rely on proxy measurements for behaviors in the US—mobility data collected via smartphone activity, survey data obtained via surveys on social media sites, and health care data obtained from health system organizations—may be unsuitable in areas where smartphones are not as widely used, social media use or literacy are less pervasive, and health care infrastructure is relatively weak. COVIDcast's forecasting models, however, do not rely on all of their indicators; the state-level autoregressive time series model, for instance, uses only COVID-19 case counts and death counts. Though COVIDcast relies on Johns Hopkins University and USAFacts for these data within the US, the WHO is a relatively accurate and reliable source of data for such information from other countries (as this reporting is mandated for the 194 WHO member states by the International Health Regulations of 2005) [37]. Thus, an autoregressive time series model of this type may be more easily repurposed for other geographies.

More information can be found in the Living Repository.

GeoSpark Analytics Hyperion COVID-19 Live Dashboard

Disclaimer: Details pertaining to technical specifications of this tool were not publicly available; the information below was obtained from blog posts on the GeoSpark Analytics site. Furthermore, whereas the tool was accessible to the general public in November 2020, at some point before October 2021, access was restricted to those with an ArcGIS account associated with the Geospark Analytics team.

The Hyperion COVID-19 Live Dashboard is a dashboard that uses machine learning to identify, track, and analyze events associated with COVID-19 via mentions on online news articles and social media posts [38]. This tool was developed by GeoSpark Analytics, a private computer software company, in partnership with Esri, a private geographic information systems software supplier (most famous for their product, ArcGIS). The dashboard was publicly announced on a blog post in April 2020, but was built upon the GeoSpark Analytics Hyperion platform, which was developed prior to the COVID-19 pandemic [39].

Specific information pertaining to the data or models used by the Hyperion COVID-19 Live Dashboard are not publicly available. A blog post describing the Hyperion platform, upon which the dashboard was built, describes three functionalities: (1) categorizing disparate forms of information into classes of activities using a machine learning model that learns patterns in unstructured data to automatically recognize and categorize data from social media, news media and other sources into themes such as social unrest, conflict and terrorism; (2) modeling patterns of human activity by evaluating news, social media, and other information in the location of the anomaly; and (3) continuously assessing levels of "stability," by comparing current activity against long-term trends (which are used to define "normalcy") within geographic regions [40]. In a separate blog post, they illustrate how their dashboard has been built with these functionalities [39]. However, they do not describe exactly what data are utilized to perform these tasks; details pertaining to the machine learning models employed are also not available.

Geospark Analytics claims that their technology detected anomalous activity levels in Wuhan, China and categorized them as a disease outbreak on December 31st 2019, eight days before the WHO announced concern over the pneumonia outbreak [40]. In April 2020, they claimed that their dashboard had been viewed more than 15,000 times in the previous month, and that it had been integrated into other applications [39], but did not describe in detail what these applications were.

As recently as November 2020, the dashboard was accessible to the general public on their website. As of October 2021, however, it appears that users need to have an ArcGIS account associated with the GeoSpark Analytics organization in order to access the dashboard.

More information can be found in the Living Repository.

IBM COVID-19 Deep Search

The IBM COVID-19 Deep Search platform organizes both structured and unstructured COVID-19 data into a knowledge graph that can be navigated and queried to retrieve information. The Deep Search platform was developed by IBM Research Europe, based in Zürich, Switzerland, in the first half of 2020, by members of the Scalable Knowledge Ingestion group [41].

IBM's Deep Search access page states that the purpose of the platform is to allow scientists and academics to "unlock the knowledge" of published unstructured and structured data pertaining to COVID-19 [42]. Users can do so by either navigating the knowledge graph manually or building query workflows to extract specific answers from the data.

Deep Search incorporates data from various unstructured and structured data sources. Unstructured data is obtained via the COVID-19 Open Research Dataset (CORD-19) [43], a large resource of scientific papers on COVID-19 and related historical coronavirus research sourced from PubMed Central (PMC) [44], the World Health Organization (WHO) COVID-19 Database [45], as well preprints from bioRxiv, medRxiv, and arXiv [46, p. 19]. Structured data included pharmaceutical and genetic databases from DrugBank [47] and Genbank [48], as well as clinical trials from Clinicaltrials.gov [49] and the World Health Organization International Clinical Trials Registry Platform [50]. In total, the platform is claimed to have ingested 158,524 COVID-19-related papers from the aforementioned sources (as of October 4th, 2021) [51], and the resulting knowledge graph contains approximately 4 million nodes and 50 million edges [42].

Deep Search is an integration of two IBM technologies: Corpus Conversion Service (CCS) and Corpus Processing Service (CPS). The development of both tools preceded the COVID-19 pandemic (IBM notes their "extensive use" in the materials science, automotive and energy industries) but were combined and made accessible to support pandemic response. CCS is a cloud-based platform that allows users to convert PDFs or bitmap documents into a structured representation of the original data. CCS parses documents (using optical character recognition to parse images), applies ML models on parsed documents to assign semantic labels to content, and reassembles documents into a machine-readable data format, such as JSON [52]. CPS then integrates this data into a knowledge graph, allowing users to navigate structured data manually on the knowledge graph interface, execute queries for specific information, and delve more deeply into specific topics by accessing source documents [42].

Access to Deep Search is granted to scientists and academics. Those interested in using the tool may apply for access on the IBM site [53], and according to IBM, there are 647 registered users, as of October 18, 2021.

According to IBM, CCS is capable of ingesting 100,000 PDF pages per day on a single server with an accuracy above 97%. This capability to structure, parse, and navigate large amounts of scientific data suggests a high potential for Deep Search to scale into areas of research with large amounts of published (or preprinted) research and accessible data.

More information can be found in the Living Repository.

Johns Hopkins US Risk Model

The Johns Hopkins US Risk Model is a county-level COVID-19 risk modeling framework intended to assist the US government and individuals in making informed decisions. The project was announced in September 2020 by researchers at Johns Hopkins Center for Systems Science and Engineering, with funding from the US National Science Foundation, National Institute of Allergy and Infectious Diseases, and NASA [54].

In a blog post announcing this initiative, researchers shared that the goal of their modeling is to identify at-risk populations and to learn the locations and attributes of those that are most exposed to risk of infection and death from COVID-19 [54]. To this end, researchers claim to have constructed their risk-modeling framework using a “flexible approach” that would allow them to model different risk indicators for different use cases [54]; however, a more in-depth explanation of how this was achieved is not provided.

In building the risk modeling framework, the initiative relied on US epidemiological, mobility, and demographic data from a number of sources [54]. Epidemiological data are drawn from the Johns Hopkins COVID-19 Data Repository, which aggregates authoritative, publicly-available COVID-19 case, death, and recovery rates from across the globe at various levels of granularity—from country-wide to city-wide, depending on availability of data [10]. Mobility data were sourced from mobile phone usage data and provided by SafeGraph. It appears, however, that while SafeGraph provided social-distancing metrics for free at the peak of the pandemic, such data have now been wrapped into their Weekly Patterns product, for purchase [55]. Both population and health indicators were gathered from the US census (population totals, demographic percentages, and age breakdowns), County Health Rankings (smoking percentages, poverty, and chronic disease), and the Definitive Healthcare Dataset published by ESRI (Statistics on hospital beds and availability) [54].

Models developed for forecasting COVID-19 risks at local, state, and national level use different statistical methodologies, such as multiple linear regression, logistic regression, random forest regression/classification, and curve fitting [54]. Researchers explored techniques that could further improve predictive capabilities, such as ensemble approaches, input clustering, and deep learning [54]. They claimed to have modeled several different aspects of the outbreak, including cases and deaths over different time horizons, case and death curves’ deviations from current trends, case and death rates per person, risk categories based on time-dependent rates of change, and categorical epidemiological classifications [54]. Details describing which sets of data were used for any particular model were not disclosed.

This initiative’s site displays a map that visually compares projected quantiles of new cases in each county during the first two weeks of August 2020 (output from the model) to observed cases reported, with striking similarities [54]. However, it is unclear whether this initiative is still under development, whether it is or was used and by whom, and what the process is for obtaining access to the model or its predictions. It is difficult to assess the technical scalability of this tool for numerous reasons, including the ambiguity with respect to the licensing of this tool,

the minimum amount of data required for any particular model, the accessibility of its data (namely, mobility data which is no longer available for free), and the quality of the required datasets at different geographic scales from across the globe.

More information can be found in the Living Repository.

RADLogics Deep Learning CT Image Analysis

RADLogics Deep Learning CT Image Analysis is an AI-assisted tool designed to quickly and accurately detect the presence of COVID-19 in thoracic CT scans [56]. The tool was developed by RADLogics, a healthcare software company based in New York, USA and Tel Aviv, Israel [57] with support from Tel-Aviv University, Affiliated Taizhou Hospital of Wenzhou Medical University, Mount Sinai Hospital and The University of Maryland School of Medicine [56]. The first academic article associated with this tool was uploaded to arXiv on March 10, 2020 [56].

RADLogics Deep Learning CT Image Analysis was developed early in the COVID-19 pandemic to respond to the growing need to quickly evaluate large numbers of thoracic CT scans for COVID-19 detection, measurements, and the tracking of disease progression.

The model was trained on 50 thoracic CT scans of patients in China, collected between January and February 2020, which were diagnosed by a radiologist as suspicious for COVID-19 [56]. The cases were extracted by querying a cloud picture archiving and communication (PACS) system for cases that were referred for laboratory testing. Each 2D slice was annotated as normal (n=1036) versus abnormal (n=829).

The Deep Learning CT Image Analysis tool consists of two subsystems that analyze thoracic images at a 3- and 2-dimensional level [56]. Subsystem A is a 3D analyzer for nodules and focal opacities, implemented with off-the-shelf software. Subsystem B detects coronavirus abnormalities using a 2D deep-learning model built on a deep convolutional neural network architecture with ResNet-50 (pre-trained using the ImageNet dataset [58]). Each subsystem makes predictions independently and the overall classification (the “corona score”) is computed based on the ratio of slices determined to be COVID-19-positive out of the total slices of lung images from the outputs from each subsystem.

The Deep Learning CT Image Analysis tool’s classification accuracy was tested on 107 thoracic CT scans—56 COVID-19-positive patients confirmed by RT-PCR, and 51 patients without any abnormal findings in a radiologist’s report—and achieved an AUC of 0.996 (95%CI: 0.989-1.00) [56]. Since this article was uploaded to arXiv in March of 2020, it is not entirely clear whether or how often the tool has been updated, or to how many hospitals the tool has been deployed.

More training and validation would seem beneficial for assessing the transferability of this tool. The training dataset could use a wider variety of clinical data, a larger-scale validation could be conducted and peer reviewed, and the tool’s capability to distinguish between COVID-19 pneumonia and non-COVID-19 pneumonia could be evaluated (as it is not covered by the available article). Furthermore, instructions for accessing the Deep Learning CT Image Analysis tool are not public-facing; it appears as though interested users need to contact RADLogics directly for access. Access to other RADLogics medical imaging tools require purchasing credentials to install or access (via the cloud) RADLogics’s patented workflow software; we assume this to be the case for this tool as well.

More information can be found in the Living Repository.

Universal Masking is Urgent in the COVID-19 Pandemic: SEIR and Agent Based Models, Empirical Validation, Policy Recommendations

Universal Masking is Urgent in the COVID-19 Pandemic: SEIR and Agent Based Models, Empirical Validation, Policy Recommendations is a preprinted research article [59] and an associated online, interactive, agent-based simulation [60] that models the spread of COVID-19 based on the prevalence of mask-wearing in a population. The researchers involved are affiliated with the Hong Kong University of Science and Technology, the International Computer Science Institute, Ecole de Guerre Economique, the University of Cambridge, Manifold research, University College Longon, ELU AI Ltd, the Royal Free Hospital, London and the Population Research Institute at The Family Federation of Finland. Both the research article and online masking simulator became accessible in April 2020.

The objective of this study was to evaluate the effectiveness of mask-wearing in preventing the spread of COVID-19 with new theoretical models and empirical data-analysis techniques. Researchers aimed to build a base of evidence to support urgent implementation of universal masking in regions that had not yet adopted it as policy or as a broad cultural norm.

The research article presents two models for predicting the impact of universal face mask wearing upon the spread of the SARS-CoV-2 virus during the pandemic: a stochastic, dynamic, network-based, compartmental, susceptible-exposed-infectious-recovered (SEIR) approach; and an individual agent-based modeling (ABM) Monte Carlo simulation [59]. For the former approach, researchers used a SEIR model implemented on a stochastic dynamic network, rather than a deterministic SEIR model, as it more closely represented interactions between individuals in a large population. Parameters were tuned to model different degrees of social distancing, lockdown stringency and mask wearing, and the empirical characteristics of COVID-19 spread as documented in the “SEIRS+” COVID-19 notebooks [61]. For the latter approach, researchers created a square wraparound two-dimensional environment, within which a population of individuals could exist in one of four SEIR states. The wraparound feature allowed the environment to represent an arbitrarily large space, giving more accurate dynamics without boundary effects from small spaces. Parameters were tuned to best approximate known COVID-19 dynamics, and the impact of masking was modeled by allowing for variation in mask wearing and *mask characteristics,* with mask transmission rate (T) and mask absorption rate (A) denoting the proportion of viruses that are stopped by masks during exhaling (transmission) versus inhaling (absorption), respectively [59].

The SEIR and ABM predictive models demonstrated that: (1) near-elimination of COVID-19 transmission when least 80% of a population is wearing masks, versus minimal effect on transmission when only 50% or less of the population is wearing masks, and (2) a significant impact when universal masking is adopted early (by day 50 of a regional outbreak), versus minimal impact when universal masking is adopted late (after day 50) [59].

To validate their models, the researchers compared their results with what little (at the time) historical macro-scale empirical data were available. They collected a data set describing the "degree of success" in managing COVID-19 by countries or regions and by the prevalence or enforcement of universal masking. The dataset contained the *number of detected COVID-19 cases from Jan 23 to April 10, 2020* and the *characteristics of universal masking culture and/or universal masking mandates or government recommendations* within 38 countries/provinces in Asia, Europe, and North America with similarly high levels of economic development. This empirical data validated the predictive models' findings for the need for universal *and* early masking.

Since this research involved predictive, simulated models, it would be relatively easy to reproduce: models could be tuned to more accurately simulate COVID-19 spread, considering the far greater amount of empirical data on COVID-19 variant characteristics and on geographical masking culture/mandates/recommendations than when the study was initially conducted in April 2020. There also exists a much greater amount of COVID-19 transmission data against which these models may be validated. It should be noted, however, the degree of uncertainty with regards to the influence that a larger base of evidence would have in changing norms or policies concerning mask wearing.

More information can be found in the Living Repository.

References

[1] A. W. Senior *et al.*, “Improved protein structure prediction using potentials from deep learning,” *Nature*, vol. 577, no. 7792, Art. no. 7792, Jan. 2020, doi: 10.1038/s41586-019-1923-7.

[2] “Computational predictions of protein structures associated with COVID-19,” *Deepmind*. https://deepmind.com/research/open-source/computational-predictions-of-protein-structures-associated-with-COVID-19.

[3] R. P. D. Bank, “RCSB PDB: Homepage.” https://www.rcsb.org/.

[4] “AlphaFold: a solution to a 50-year-old grand challenge in biology,” *Deepmind*. https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology.

[5] J. Jumper *et al.*, “Highly accurate protein structure prediction with AlphaFold,” *Nature*, vol. 596, no. 7873, pp. 583–589, Aug. 2021, doi: 10.1038/s41586-021-03819-2.

[6] “Highly accurate protein structure prediction with AlphaFold | Nature.” https://www.nature.com/articles/s41586-021-03819-2.

[7] J. M. Brauner *et al.*, “Inferring the effectiveness of government interventions against COVID-19,” *Science*, vol. 371, no. 6531, Feb. 2021, doi: 10.1126/science.abd9338.

[8] “Datasets.” http://epidemicforecasting.org/datasets.

[9] “COVID-19 Government Response Tracker.” https://www.bsg.ox.ac.uk/research/research-projects/covid-19-government-response-tracker.

[10] CSSEGISandData, *COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University*. 2021. https://github.com/CSSEGISandData/COVID-19

[11] S. Flaxman *et al.*, “Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe,” *Nature*, vol. 584, no. 7820, pp. 257–261, Aug. 2020, doi: 10.1038/s41586-020-2405-7.

[12] M. D. Hoffman and A. Gelman, “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo,” *JMLR*, vol. 15, p. 31, Apr. 2014.

[13] D. DeCaprio *et al.*, “Building a COVID-19 vulnerability index,” *Journal of Medical Artificial Intelligence*, vol. 3, no. 0, Art. no. 0, Dec. 2020, doi: 10.21037/jmai-20-47.

[14] D. DeCaprio, J. Gartner, C. J. McCall, T. Burgess, S. Kothari, and S. Sayed, “Building a COVID-19 Vulnerability Index,” p. 9, Mar. 2020, doi: 10.1101/2020.03.16.20036723.

[15] “Limited Data Set (LDS) Files | CMS.” https://www.cms.gov/Research-Statistics-Data-and-Systems/Files-for-Order/LimitedDataSets.

[16] “Illinois ACO deploys AI to identify high risk COVID-19 patients.” https://www.beckershospitalreview.com/artificial-intelligence/illinois-aco-deploys-ai-to-identify-high-risk-covid-19-patients.html.

[17] B. Prud’homme *et al.*, “COVI,” *Mila*. https://mila.quebec/en/project/covi/.

[18] H. Alsdurf *et al.*, “COVI White Paper,” *arXiv:2005.08502 [cs]*, Jul. 2020, http://arxiv.org/abs/2005.08502

[19] “COVI-ML.” Mila. https://github.com/mila-iqia/COVI-ML

[20] P. Gupta *et al.*, “COVI-AgentSim: an Agent-based Model for Evaluating Methods of Digital Contact Tracing,” *arXiv:2010.16004 [cs]*, Oct. 2020, http://arxiv.org/abs/2010.16004

[21] “Federal government rules out adoption of Mila Institute’s COVID-19 contact-tracing app,” *The Logic*, Jun. 04, 2020. https://thelogic.co/news/federal-government-rules-out-adoption-of-mila-institutes-covid-19-c

ontact-tracing-app/.
[22] "COVID-19 Patient Redistribution." https://covid-hospital-operations.com/patients-interactive.
[23] "COVID-19 Patient Redistribution." https://covid-hospital-operations.com/about#aboutus.
[24] C. G. / P. Oct 27 and 2020, "New models can help hospitals stay ahead of COVID-19 surges," *The Hub*, Oct. 27, 2020. https://hub.jhu.edu/2020/10/27/models-help-hospital-systems-allocate-resources-during-covid-19-surge/.
[25] F. Parker, H. Sawczuk, F. Ganjkhanloo, F. Ahmadi, and K. Ghobadi, "Optimal Resource and Demand Redistribution for Healthcare Systems Under Stress from COVID-19," *arXiv:2011.03528 [cs, math]*, Nov. 2020, http://arxiv.org/abs/2011.03528
[26] "COVID-19 Reported Patient Impact and Hospital Capacity by Facility | HealthData.gov." https://healthdata.gov/Hospital/COVID-19-Reported-Patient-Impact-and-Hospital-Capa/anag-cw7u.
[27] CDC, "Coronavirus Disease 2019 (COVID-19)," *Centers for Disease Control and Prevention*, Feb. 11, 2020. https://www.cdc.gov/coronavirus/2019-ncov/science/forecasting/forecasts-cases.html.
[28] F. Parker, *COVID-19 Patient Redistribution Site*. 2021. https://github.com/flixpar/covid-patient-redistribution-site
[29] D. Group, "Delphi Research Group." https://delphi.cmu.edu/about/.
[30] D. Group, "Delphi National Center of Excellence for Epidemic Forecasting." https://delphi.cmu.edu/about/center-of-excellence/.
[31] D. McDonald *et al.*, *Can Auxiliary Indicators Improve COVID-19 Forecasting and Hotspot Prediction?* 2021. doi: 10.1101/2021.06.22.21259346.
[32] "Model: CMU-TimeSeries - Zoltar." https://zoltardata.com/model/307.
[33] "Home - COVID 19 forecast hub." https://covid19forecasthub.org/.
[34] "COVIDcast Epidata API," *Delphi Epidata API*. https://cmu-delphi.github.io/delphi-epidata/api/covidcast.html.
[35] D. Group, "Accessing Open COVID-19 Data via the COVIDcast Epidata API," Oct. 07, 2020. https://delphi.cmu.edu/blog/2020/10/07/accessing-open-covid-19-data-via-the-covidcast-epidata-api/.
[36] "COVIDcast API: An open repository of real-time COVID-19 indicators," *News-Medical.net*, Jul. 21, 2021. https://www.news-medical.net/news/20210721/COVIDcast-API-An-open-repository-of-real-time-COVID-19-indicators.aspx.
[37] "International Health Regulations (2005)." https://www.who.int/publications-detail-redirect/9789241580410.
[38] G. Analytics, "Coronavirus Dashboard." https://knowledge.geospark.io/coronavirus-dashboard.
[39] A. Brownfield, "Not Just Another Coronavirus Map," *Medium*, Apr. 03, 2020. https://blog.geospark.io/not-just-another-coronavirus-map-c0a62e9bd2d8.
[40] M. Franklin, "From Localized Incident Reporting to Global Monitoring — How Hyperion has Processed and Analyzed…," *Medium*, Jan. 31, 2020. https://blog.geospark.io/from-localized-incident-reporting-to-global-monitoring-how-hyperion-has-processed-and-analyzed-840d0e48fdb2.
[41] "IBM's New AI Tool Parses A Tidal Wave of Coronavirus Research," *IEEE Spectrum*, Jul. 01, 2020. https://spectrum.ieee.org/ibms-new-ai-tool-parses-a-tidal-wave-of-coronavirus-research.
[42] "Use deep search to explore the COVID-19 corpus," *IBM Research*, Jun. 05, 2018. https://research.ibm.com/covid19/deep-search/.
[43] "CORD-19 | Semantic Scholar." https://www.semanticscholar.org/cord19.

[44] “Home - PMC - NCBI.” https://www.ncbi.nlm.nih.gov/pmc/.
[45] “Global research on coronavirus disease (COVID-19).” https://www.who.int/emergencies/diseases/novel-coronavirus-2019/global-research-on-novel-coronavirus-2019-ncov.
[46] L. L. Wang *et al.*, “CORD-19: The COVID-19 Open Research Dataset,” presented at the ACL-NLP-COVID19 2020, Online, Jul. 2020. https://aclanthology.org/2020.nlpcovid19-acl.1
[47] D. S. Wishart *et al.*, “DrugBank 5.0: a major update to the DrugBank database for 2018,” *Nucleic Acids Research*, vol. 46, no. D1, pp. D1074–D1082, Jan. 2018, doi: 10.1093/nar/gkx1037.
[48] “GenBank Overview.” https://www.ncbi.nlm.nih.gov/genbank/.
[49] “Home - ClinicalTrials.gov.” https://clinicaltrials.gov/.
[50] “International Clinical Trials Registry Platform (ICTRP).” https://www.who.int/clinical-trials-registry-platform.
[51] “IBM COVID-19 Knowledge Graph.” https://ds-covid19.res.ibm.com/datasources.
[52] P. W. J. Staar, M. Dolfi, C. Auer, and C. Bekas, “Corpus Conversion Service: A Machine Learning Platform to Ingest Documents at Scale,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, London United Kingdom, Jul. 2018, pp. 774–782. doi: 10.1145/3219819.3219834.
[53] “Deep Search - COVID-19- IBM Research.” https://research.ibm.com/interactive/covid19/deep-search/.
[54] M. Marshall and L. Gardner, “Predicting COVID-19 Risk in the U.S.,” Sep. 03, 2020. https://systems.jhu.edu/research/public-health/predicting-covid-19-risk/.
[55] “Social Distancing Metrics | SafeGraph Docs,” *SafeGraph*. https://docs.safegraph.com/docs/social-distancing-metrics.
[56] O. Gozes *et al.*, “Rapid AI Development Cycle for the Coronavirus (COVID-19) Pandemic: Initial Results for Automated Detection & Patient Monitoring using Deep Learning CT Image Analysis,” *arXiv:2003.05037 [cs, eess]*, Mar. 2020, http://arxiv.org/abs/2003.05037
[57] “About Us,” *RADLogics*. https://www.radlogics.com/about-us/.
[58] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2009, pp. 248–255. doi: 10.1109/CVPR.2009.5206848.
[59] D. Kai, G.-P. Goldstein, A. Morgunov, V. Nangalia, and A. Rotkirch, “Universal Masking is Urgent in the COVID-19 Pandemic: SEIR and Agent Based Models, Empirical Validation, Policy Recommendations,” *arXiv:2004.13553 [physics, q-bio]*, Apr. 2020, http://arxiv.org/abs/2004.13553
[60] “Agent-Based Model for COVID-19 Masking.” http://dek.ai/masksim/.
[61] R. S. McGee, *SEIRS+ Model Framework*. 2021. https://github.com/ryansmcgee/seirsplus

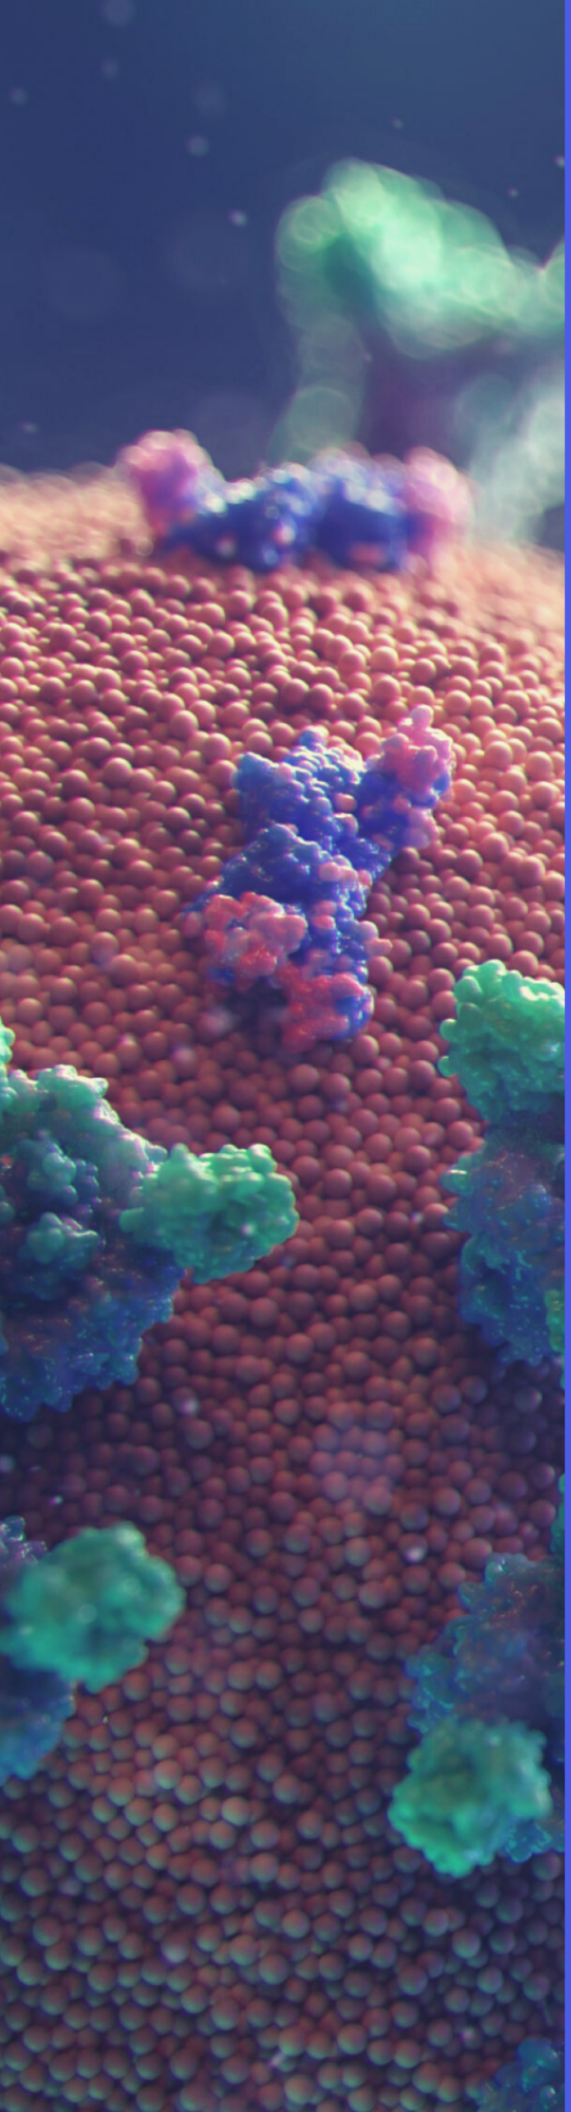

THE FUTURE SOCIETY

The Future Society is a nonprofit think-and-do tank with the mission to advance the responsible adoption of Artificial Intelligence and other emerging technologies for the benefit of humanity.

The Future Society
867 Boylston Street, 5th Floor, Boston MA 02116, United States
www.thefuturesociety.org