

How artificial intelligence and machine learning can help healthcare systems respond to COVID-19

Lead Author: Prof. Mihaela van der Schaar, John Humphrey Plummer Professor of Machine Learning and AI in Medicine, Cambridge University; Director, *Cambridge Centre for AI in Medicine*; Turing Faculty Fellow, The Alan Turing Institute; Chancellor’s Professor, UCLA

Contributing Authors: Dr. Ahmed Alaa, Prof. Andres Floto, Dr. Alexander Gimson, Prof. Stefan Scholtes, Dr. Angela Wood, Dr. Eoin McKinney, Dan Jarrett, Prof. Pietro Lio, Dr. Ari Ercole; *Cambridge Centre for AI in Medicine*

Abstract

The COVID-19 global pandemic is a threat not only to the health of millions of individuals, but also to the stability of infrastructure and economies around the world. The disease will inevitably place an overwhelming burden on healthcare systems that cannot be effectively dealt with by existing facilities or responses based on conventional approaches.

We believe that a rigorous clinical and societal response can only be mounted by using intelligence derived from a variety of data sources to better utilize scarce healthcare resources, provide personalized patient management plans, inform policy, and expedite clinical trials.

In this paper, we introduce five of the most important challenges in responding to COVID-19 and show how each of them can be addressed by recent developments in machine learning (ML) and artificial intelligence (AI). We argue that the integration of these techniques into local, national, and international healthcare systems will save lives, and propose specific methods by which implementation can happen swiftly and efficiently.

We offer to extend these resources and knowledge to assist policymakers seeking to implement these techniques.

Challenge 1: Managing limited healthcare resources

Healthcare systems in the UK and elsewhere in the world will soon face a severe scarcity of resources—in particular testing kits, hospital beds, ICU beds, ventilators and personnel. The UK, for example, currently has capacity for only a limited amount of tests to be conducted per day [1] and many of the intensive care unit (ICU) beds are already occupied [2]. This scarcity of resources is aggravated by the fact that individuals with COVID-19 are known to experience widely different patterns of disease progression and outcomes: while some patients are asymptomatic, others manifest flu-like symptoms of varying severity, and some experience complications such as pneumonia and fatal multi-organ failure [3]. Since resources are limited and risk and disease progression so heterogeneous, it is crucial to identify—as early as possible—*which individuals* are most likely to have been infected by the virus, *which infected individuals* may experience adverse events, *which types of medical resources* those individuals will require, and *when* these resources will be required.

The scarcity of healthcare resources will be exacerbated by the need to employ those resources to deal with both COVID-19 cases and with other patients who require—or will require—medical care. There is thus a clear and urgent need to deploy systems that can provide early warnings for personalized risk and disease progression of individuals, and that can inform medical personnel and healthcare systems about which patients would benefit from what resources and when [5].

How AI can help

Mature AI-based support systems for a number of chronic diseases already exist, but a pandemic such as COVID-19 presents a problem of a different nature.

While chronic diseases jeopardize *individual health*, pandemics jeopardize *public health* because they present both *individual* and *systemic* risks. Pandemics, by nature, undermine the underlying social structures that connect individuals together.

Dealing with these systemic risks requires merging clinical data with a variety of diverse social data. Electronic health records (EHR) hold data that can be used to pinpoint individuals' clinical risk factors [4], and which can be linked to the multitude of "big data" pertaining to human-to-human interactions (e.g. data from cellular operators, traffic, airlines, and social media). ML and AI are especially suited for merging these various sources of data to issue accurate predictions of risk and help uncover the social structures through which systemic risks manifest and spread.

One of the strengths of ML is its ability to "learn" how an individual's features (risk factors), including clinical and social information, can be mapped into *personalized predictions of risk* [6]. While standard epidemiological approaches—such as the Cox proportional hazards model—are unable to effectively combine data from diverse data sources and modalities (e.g. demographic, social, longitudinal, imaging, multi-omics), modern techniques for representation learning based on neural networks can readily and effectively learn from such diverse data to issue personalized predictions of risk as well as to update these predictions as features evolve over time. For instance, learning can be informed by dynamic models of social interaction [7][42], and hence can predict the likelihood that an individual has been, or will be, in contact with a coronavirus carrier. This makes it possible to better allocate testing kits such that individuals most likely to have been exposed to the virus can be identified and tested most promptly. Such risk prediction mechanisms can be used for disease prevention, monitoring and detection.

With regard to individuals who have already been infected with SARS-CoV-2, it is important to be able to predict, on a personalized basis, the risk of experiencing adverse events (including mortality) as well as the specific healthcare resources a patient will/might need, and the timing of these needs. Key factors determining risk include patient age and/or the existence of comorbidities, such as hypertension, cystic fibrosis, or immunosuppression associated with transplants and chemotherapy for cancer [8]. However, the current risk-scoring methods for infected patients employ only a few factors such as Clinical Frailty Score, the number of comorbidities, or use of medication or administrative data from previous hospital admissions, to predict which patients are at highest risk. These relatively crude approaches (such as using the Charlson score [9]) may be appropriate in the frenetic acute stages of this epidemic, but they fail to capture and account for the subtle interactions between age, specific comorbidities (including their duration), organ failure scores and the impact of the wider determinants of health. Moreover, reliance on such measures may also result in failure to discover new risk factors—which may be especially important for a disease about which so little is known. Methods based on automated ML [6][43][44] can address both of these failings, offering increasingly accurate subtypings of the COVID-19 disease as more data is gathered, while simultaneously uncovering new risk factors as well as interactions among new and established risk factors.

Moreover, patients with severe comorbidities may experience diverse adverse events, especially because their usual treatments may be suspended in order to avoid compromising their immune response. In these cases, ML methods for assessing competing risks at the individual level—estimating the probabilities of a patient experiencing different adverse events and how these probabilities are changing over time—have shown improved prediction accuracy compared to conventional epidemiological methods. These methods can play an essential role in guiding decisions for individual patients.

As noted, patients diagnosed with COVID-19 experience a wide variety of disease trajectories and outcomes, such as heterogeneous incubation periods and variable rates of in-hospital pulmonary decline [3]. As a consequence, risk assessment must involve more than just a single time point prediction of final outcomes, and must also take account of the evolution of disease severity. This provides not only a way to predict a patient's disease trajectory but—by including follow-up data—it enables understanding of how specific events (e.g. symptoms worsening) may move patients from one trajectory or cluster/temporal phenotype to another. This is fundamental to the delivery of systems that allow the *dynamic* management of limited healthcare resources.

There are two routes by which resource allocation should be an active choice guided by AI-based risk assessment:

(a) *Patient-level* resource allocation concerning decisions made on an individual basis. This includes such decisions as who is most likely to benefit from early admission to an ICU, which patients can be safely stepped down from intensive support, or which patients can be safely discharged from the hospital and when.

Such questions of timely admission and discharge are often answered on the basis of simple rules of thumb (e.g. examining the most recent measurements for a select few variables, and making decisions based on threshold values). A more tailored approach to triaging should involve incorporating information from the individual patient's history using all available observations at any point in time. Such triaging systems should be informed by state-of-the-art AI-based early warning systems trained on temporal (longitudinal) data. These AI-enabled systems have been developed and tested, and have shown their efficacy in practice [10, 11]. They can anticipate a patient's future healthcare needs, and can also handle irregularly collected data and impute missing features [12, 13], as well as inferring important predictors of patient outcomes (e.g. specific comorbidities that are more likely to worsen the prognosis of a COVID-19 patient) [14, 15]. Importantly, ML and AI approaches are well-suited for handling the irregularly-sampled temporal data common in this setting, employing mixtures of Gaussian processes [45], recurrent neural networks [18], state-space models [19], neural network-based point processes [20], Bayesian non-parametric approaches [46].

(b) *Measurement-level* resource allocation (*value of information*). Making observations and collecting healthcare data is costly in terms of time, materials, and personnel – and not all data is equally valuable. We need to actively decide what we should measure, when it should be measured, and on which patients; making observations on each patient should be an active and informed process, so that the information gained can be most valuable. Developments in the field of active sensing offer tools for optimally trading off the value of information against the cost of acquisition [16, 17].

Taken together, techniques in patient- and measurement-level prediction and active sensing can integrate information across time in order to effectively allocate resources among patients while prioritizing measurements and treatments, thereby enabling institutions to dynamically allocate hospital resources depending on their patients' needs. This, in turn, allows physicians to make informed decisions to postpone treatment for conditions that are not immediately life-threatening, so that resources can be used for critically-ill COVID-19 patients.

Challenge 2: Developing personalized patient management and treatment plans

Vaccines and therapeutics for COVID-19 will take some time to be developed and introduced [28]. Although a number of antiviral medications are being trialed, none are currently known to be effective. Due to the heterogeneous nature of patient disease trajectories, one-size-fits-all treatment plans (which are developed at the population level, rather than individualized patient level) are very likely to be ineffective in terms of both patient outcomes and use of limited resources. To save lives and effectively use scarce resources, it is therefore essential to develop management and treatment plans that are personalized for each patient, depending on their unique characteristics and needs.

How AI can help

AI-based models can efficiently use available observational data to learn the effects of existing treatments on individuals, given their specific features (characteristics) including existing comorbidities. This can help identify the best management approach for each patient based on their features, including the best policies for employing mechanical ventilation and/or experimental treatments.

Learning personalized treatment *effects* can be carried out using recent advances in machine learning models tailored for learning in the observational data setting. Models based on Gaussian processes [27], generative adversarial networks [30], and deep neural networks [31] can learn individual-level effects of experimental treatments on the basis of observational data (without resorting to randomized trials), and can potentially be integrated with omics data. Using such methods, we can identify patient subgroups for whom treatments such as Remdesivir, Lopinavir/ritonavir and Chloroquine may be effective [32].

Learning personalized *management and treatment plans* requires modeling the evolution of a patient's health state as they respond to different sequences of treatments and supportive care over time during hospitalization. AI-

based models [27][30][31] can answer what-if questions such as: “what would happen to a given patient if mechanical ventilation were delayed for a few hours?” or “what would a given patient’s outcomes be if we switched from supportive care to a specific experimental treatment a week earlier?”

Because each patient is different, answering *counterfactual* questions such as these requires efficient use of regularly-updated observational data from patients who have experienced different management plans since the beginning of the viral outbreak; this can help us learn what the best plan is for each individual given the available healthcare resources. The counterfactual questions under consideration here pertain not only to *which* intervention should be applied for each patient, but also *when* the patient needs this intervention. This is especially relevant given the limited resources available.

Leveraging longitudinal patient observational data, the counterfactual recurrent network in [33] learns to predict counterfactual health trajectories under different management plans for each individual. Such predictions can be made in a dynamic fashion by modeling and updating the patient’s state over time to account for information in all past medical events. Doing this can pinpoint a timeframe during which intervention would be most effective, as well as choosing the specific sequence of interventions that would result in the best outcome for each patient. In addition, this method has the potential to show when scarce interventions—such as the mechanical ventilator—are not critically needed by a patient, and for how long their use can be postponed without the patient suffering significant deterioration.

Challenge 3: Informing policies and enabling effective collaboration

Governmental and local responses to the outbreak have been remarkably varied [38]. At the national, regional, institutional and even individual practitioner level, very different decisions are being made in terms of prevention, diagnosis, treatment and disease management. In uncertain settings—like the ones presented by COVID-19—policymakers and clinical practitioners make different observations, take different measurements, and act differently on the basis of evidence. For example, differences in triaging protocols across institutions and practitioners could mean that two patients with similar profiles end up receiving different courses of treatment simply because of where they happen to be seen.

There is little consensus as to what policies and decisions are most effective; as a result, there is substantial variation in efficiency of resource usage, quality of care, and patient outcomes. It is imperative that we employ systematic, data-driven methods in learning from the policies already implemented in the wild. Coordinating experiences across different actors allows us to more systematically quantify policies and their efficacies, thereby informing the design of improvements.

How AI can help

Evidence-based decision-making entails collecting costly observations about the effects of a treatment or resource allocation policy, and subsequently committing to an informed decision on the basis of accumulated observations. AI methods tackle the following important and inter-related questions, and offer a common language through which healthcare practitioners can exchange information and improve policies over time:

What are the different diagnostic and treatment policies?

First, we can quantify the *decision rules* that underlie observed behavior (e.g. the guidelines and protocols often implicitly encoded in different institutions). For a concise, standardized *description* of behavior, what is learned here is an accessible mapping from information (e.g. patient characteristics and measurements) to actions (i.e. diagnostic tests and treatments)—a mapping that is immediately indicative of which individual actions are more or less likely to be taken in any given scenario. For instance, precisely how badly does a patient have to be doing before a ventilator is employed? How does this compare across hospitals? Modeling policies transparently enables identification and comparison of actions by different decision-makers.

What are the priorities demonstrated by different policies?

Second, we can quantify the *preferences* and potential biases implicit in observed behavior (e.g. what treatments are preferred *ceteris paribus*, and how time-sensitivity and monetary costs weigh on decisions). For a concise, standardized *interpretation* of different policies, what is obtained is an accessible mapping from behavior to preferences—a mapping that is immediately indicative of what aspects of the decision-making process (e.g. expense, accuracy, time pressure, and patient populations) appear more or less important for a given decision-maker. For instance, we may expect that an institution diagnoses and treats older patients more quickly than younger ones. But do they actually? And if so, by how much?

What are the best improvements to existing policies?

Third, we can quantify the *efficacy* of both new and existing policies, so as to identify the best policy to implement under different priorities and in different circumstances. While answering the two directly preceding questions can facilitate transparent comparison and understanding of variations in clinical practice, here the culminating objective is in the *prescription* of better diagnosis and treatment policies. This can be accomplished by integrating the results of answering both of the preceding questions with novel techniques in reinforcement learning to inform and instruct policy evaluation and policy improvement. For instance, knowing that we desire to optimize for a certain agreed-upon set of priorities, precisely what is the best decision rule to implement going forward?

What is the level of confidence in learned policies?

Finally, we can quantify the *uncertainty* involved in these tasks of identifying, understanding, and improving decisions and policies. This is critical in new environments where existing predictive expertise is lacking—such as in responding to COVID-19. For instance, novel diagnostic tests will certainly have errors, and the overall effectiveness of different treatments is not known *a priori*. Quantifying uncertainty permits a clear understanding of which components of our models we can have more confidence in, and which components we cannot. For instance, when can we assert that a given action is a no-brainer, and when does the choice require more nuanced deliberation?

In the presence of constraints on medical personnel, the potential for decision support through policy evaluation, improvement, and well-calibrated uncertainty estimates enables routine decisions to be taken in a more timely manner—thus freeing up valuable medical attention to the cases that warrant real-time expertise. Recent work in tackling some of these challenges involves quantifying and explaining behavior in the context of early diagnosis [34], modeling and understanding sequential decisions under uncertain deadline pressure [35], and learning interpretable parameterizations of behavior via counterfactual reasoning to highlight cost-benefit tradeoffs associated with different actions over time [36]. There is, therefore, enormous opportunity for AI to facilitate healthcare professionals in making better decisions by learning from the “experimentation” currently taking place in different settings for different patients.

Challenge 4: Understanding and accounting for uncertainty

Because it is novel, the COVID-19 pandemic is being managed with very limited prior experience and limited data. Our knowledge of its behavior is very uncertain [21]. Moreover, since the coronavirus may undergo further natural selection and evolve geospatially as the pandemic progresses, understanding this evolution and the uncertainties that accompany it will be key to mitigation.

To work effectively within this environment, we need to quantify the *degree of uncertainty* in our knowledge about the disease, and then apply this to understand the risks and benefits of the clinical and social policies employed to mitigate it. This is especially important because unproven hypotheses about COVID-19 have already circulated rapidly online and through various media outlets, and are likely to impact individual behavior and hence affect systemic risks.

There are two particular sources of uncertainty encountered when dealing with the current COVID-19 pandemic. First, models trained using data from the population tested for SARS-CoV-2 may be biased, since patients who have been tested are not necessarily representative of the overall UK population. Relying on data from other countries (e.g. China, Italy, Iran or South Korea) where the outbreak took off earlier may lead to biased models that do not generalize well to the UK population. Second, there is simply a limited amount of data available about

COVID-19 patients. In light of the enormous heterogeneity of population characteristics and outcomes, blanket hypothesis tests about the disease may appear erroneously inconclusive.

How AI can help

ML techniques provide a wide variety of solutions for modeling uncertainty in these two settings:

The problem of biased models can be addressed via *transfer learning* methods, which implicitly cope with uncertainty in predictions made for individuals from populations other than those used for training. International collaboration could produce a large dataset comprising data from multiple infected populations, and a single two-level risk assessment model can be trained on such a dataset and deployed worldwide. However, since the characteristics of the UK population differ from other populations (e.g. the median age in the UK is 40.5 years, whereas for Iran it is 31.3 years), the model needs to “learn” its uncertainty in predictions made for specific individuals in the UK. Transfer learning models [23, 24] accomplish this task effectively, and have been previously shown to work well for prognostication of disease across different countries or healthcare systems.

Transfer learning is relevant not only for transferring models *across populations*, but also for updating ML models for a single population as the diagnosed cohort evolves *over time*. Specifically, as diseases and populations evolve over time, observed distributions of patients—and their characteristics and outcomes—may change. Using successive batches of data as they arrive, scalable Bayesian optimization techniques have been developed to update existing models with new information for continued performance, while efficiently leveraging past optimizations [25].

Importantly, we need to know what we *do* know and what we *do not* know when providing decision-makers with actionable information based on AI-systems. For this, new technologies have been developed to systematically quantify uncertainty in the predictions made by machine learning models [26]. We need to mandate that all predictions generated by AI-enabled systems for decision support also come with confidence estimates. Decision-makers will then know when and for which patients they can trust the predictions issued by the AI system, and when they cannot have the same degree of certainty.

Challenge 5: Expediting clinical trials

In order to compare the effectiveness of a new treatment to an existing one (or to a placebo), randomized clinical trials (RCTs) are considered the gold standard. However, it is a well-known fact that RCTs can be slow and costly, and may fail to uncover specific sub-populations for which a treatment would be most effective. Moreover, conclusions drawn from RCTs are typically valid only for the types of patients recruited for that RCT. This is an especially significant problem in the case of COVID-19, as elderly patients and patients with comorbidities, who are known to be at higher risk, are typically excluded from RCTs.

How AI can help

Recent work on improving the design of adaptive clinical trials has demonstrated that both the efficiency and effectiveness of RCTs can be significantly enhanced through the application of ML.

While most RCTs simply allocate patients to treatment and control groups through uniform randomization, this procedure can be highly sub-optimal in terms of learning. ML-enabled adaptive clinical trials [37][47] can recruit patients in cohorts (rather than all at once), the effects on each cohort can be observed before recruiting the next cohort, and may be heterogeneous across identifiable subgroups of patients (as would be expected for COVID-19 treatments). Rather than recruiting and assigning subjects at random, ML methods can recruit subjects from identifiable subgroups, and assign subjects to treatment or control groups in a way that speeds up learning. These methods have been shown to significantly reduce error and achieve a prescribed level of confidence in findings, while requiring many fewer patients.

Another important application of ML in the domain of RCTs is the post-hoc analysis of clinical trials, and, in particular, the identification of sub-populations that are most similar in characteristics (array of features) and treatment response. The identification of such sub-populations is informative in itself because it permits inference about the extent of heterogeneity and response across the entire population. Moreover, quantifying the uncertainty associated with the treatment of specific sub-populations makes it easier to target particular treatments to those sub-populations and understand what treatments are suitable for the population as a whole.

Recommended means of implementing the techniques

For the variety of AI-based models discussed, deployment can take advantage of existing data infrastructure. Information about patients from EHR records, can be combined—through collaboration between public and private sectors—with data from cellular operators, airlines, traffic applications, and social media. Such collections of static and temporal data streams can be integrated by AI models to enable more personalized predictions of risk and treatment effects, patient management policies, and better trial design.

The UK is in an excellent position to capitalize on the available data and employ the aforementioned AI technologies to improve outcomes. Recent emergency privacy legislation has allowed the collection, linkage and analysis of record-level, patient-identifiable data by the national health and social care data aggregator, NHS Digital. In addition, all COVID-19 testing data is collated nationally by Public Health England, while patients suspected to have been infected are also triaged through a national health triage service, with the data from calls being made available in real time. Daily service-level reporting of resource utilization across the acute sector is being expanded, while similar capacity monitoring is being introduced across social care. Taken together, this nationally-collated and linked dataset provides the ideal environment to implement the techniques outlined in this paper.

Implementation of proposed AI-based methods can be accomplished through three levels of human-machine interface:

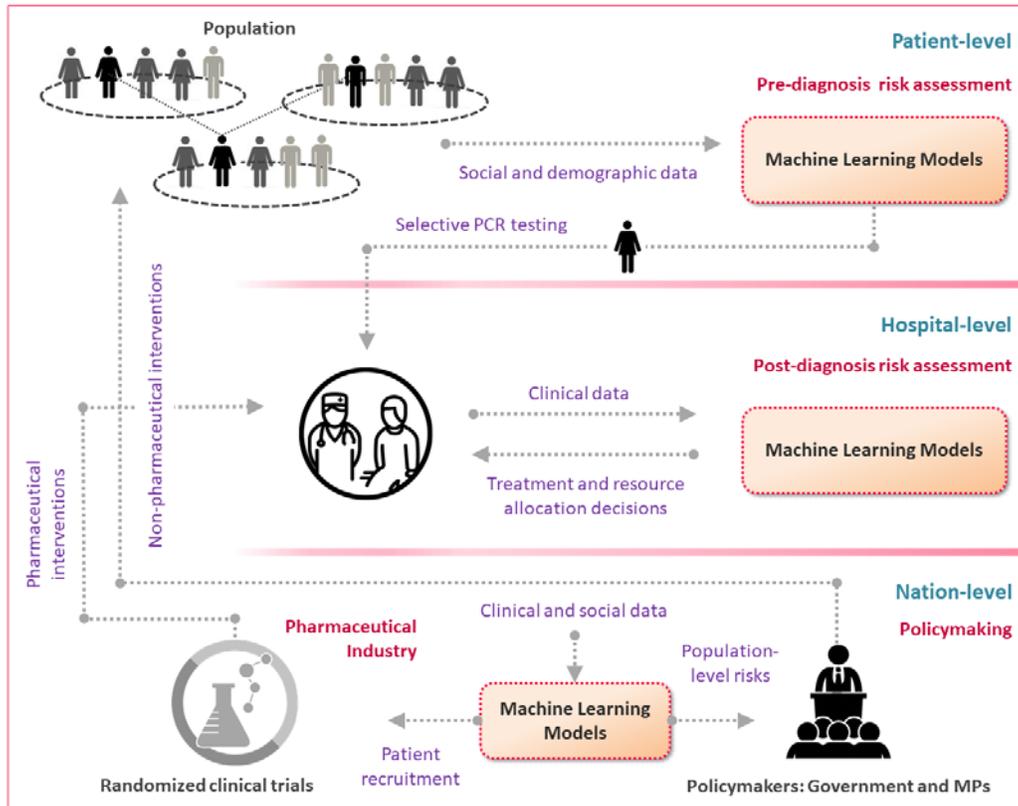
(a) Individual-level interface. Based on individual health risk assessments issued by AI models, patients can be alerted to undergo PCR testing via email, mobile phone calls, and text messages. GPS-powered mobile applications installed by individuals on their cell phones can be used to guide social distancing by notifying individuals in high-risk geographical locations where a significant number of COVID-19 cases have been diagnosed, or where a significant number of high-risk potential carriers of the virus are currently present. Implementation of this system can be primarily based on software applications that can be promptly executed and linked to NHS-held GP patient registries.

(b) Hospital-level interface. For public health risk assessments, existing in-hospital EHR infrastructure can be complemented with risk-modeling applications that display and alert clinicians to the evolving risks of patients within each hospital. Existing telehealth applications adopted by the NHS's Technology Enabled Care Services (TECS) program can be utilized for distant diagnosis and care for self-isolating individuals with evident symptoms. Implementation of AI-based models in hospitals can largely rely on existing hardware infrastructure to link input data streams from ICUs to software (powered by ML models) that displays patient risks to clinicians.

(c) Nation-level interface. Government officials and decision-makers can be supported with centralized applications that continuously update officials with the most recent epidemiological data on COVID-19, in addition to geographically-stratified aggregate risk assessments and hospital occupancy rates for the entire population. This can be directly linked to the NHS patient registration database, and would interface with the other applications to collect information on population risk. This can help officials make selective strategic decisions on which locations to lock down, and which hospitals need to be supplied with more resources.

In the UK, the National Health Service provides a centralized route to deployment for a whole country, while for example in the US, implementation may be most effective through large provider networks, HMOs or insurers.

The figure below provides a high-level illustration of how the proposed ML and AI solutions can help healthcare systems respond to COVID-19.



These techniques are proven, and must be implemented without delay

The methodologies outlined in this paper have been subject to peer review, come with confidence guarantees, and are sufficiently mature to be rapidly implemented at local, regional and national levels. The machine learning methods presented are interpretable and transparent, meaning users (policymakers, clinicians, scientists, patients, etc.) can safely and easily integrate them into their decision-making processes.

We offer these technologies, resources and knowledge freely in assistance of policymakers seeking to implement these techniques. For more information, please contact get in touch contact the lead author, Prof. Mihaela van der Schaar (mv472@cam.ac.uk).

References

- [1] <https://www.wired.co.uk/article/coronavirus-testing-uk>
- [2] <https://www.independent.co.uk/news/health/coronavirus-europe-intensive-care-beds-ventilators-a9403981.html>
- [3] Jiang, Fang, Liehua Deng, Liangqing Zhang, Yin Cai, Chi Wai Cheung, and Zhengyuan Xia. "Review of the clinical characteristics of coronavirus disease 2019 (COVID-19)." *Journal of General Internal Medicine* (2020): 1-5.
- [4] Wang, Minhong, Qianyi Zeng, Wenjun Chen, Jeff Pan, Honghan Wu, Cathie Sudlow, and Dave Robertson. "Building the Knowledge Graph for UK Health Data Science." ERA (2020).
- [5] <https://www.nytimes.com/2020/03/21/us/coronavirus-medical-rationing.html>
- [6] Alaa, A.M, and van der Schaar, M. "AutoPrognosis: Automated Clinical Prognostic Modeling via Bayesian Optimization with Structured Kernel Learning," *International Conference on Machine Learning (ICML)*, 2018.
- [7] Alaa, A.M., Ahuja, K., van der Schaar, M.: A micro-foundation of social capital in evolving social networks. *IEEE Trans. Netw. Sci. Eng.* 5(1), 14–31, 2018.
- [8] Zhou, Fei, Ting Yu, Ronghui Du, Guohui Fan, Ying Liu, Zhibo Liu, Jie Xiang et al. "Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study." *The Lancet* (2020).
- [9] Goldstein, Larry B., Gregory P. Samsa, David B. Matchar, and Ronnie D. Horner. "Charlson Index comorbidity adjustment for ischemic stroke outcome studies." *Stroke* 35, no. 8 (2004): 1941-1945.
- [10] Yoon, Jinsung, Ahmed Alaa, Scott Hu, and Mihaela Schaar. "ForecastICU: a prognostic decision support system for timely prediction of intensive care unit admission." In *International Conference on Machine Learning*, pp. 1680-1689. 2016.
- [11] Lee, Changhee, William R. Zame, Jinsung Yoon, and Mihaela van der Schaar. "Deephit: A deep learning approach to survival analysis with competing risks." In *Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*. 2018.
- [12] Yoon, Jinsung, William R. Zame, and Mihaela van der Schaar. "Estimating missing data in temporal data streams using multi-directional recurrent neural networks." *IEEE Transactions on Biomedical Engineering* 66.5 (2018): 1477-1490.
- [13] Yoon, Jinsung, James Jordon, and Mihaela Van Der Schaar. "Gain: Missing data imputation using generative adversarial nets." *International Conference on Machine Learning (ICML)*, 2018.
- [14] Yoon, Jinsung, James Jordon, and Mihaela van der Schaar. "INVASE: Instance-wise variable selection using neural networks." *International Conference on Learning Representations (ICLR)*, 2018.
- [15] Jordon, James, Jinsung Yoon, and Mihaela van der Schaar. "KnockoffGAN: Generating knockoffs for feature selection using generative adversarial networks." *International Conference on Learning Representations (ICLR)* 2018.
- [16] Yoon, Jinsung, James Jordon, and Mihaela van der Schaar. "ASAC: Active Sensing using Actor-Critic models." *Machine Learning for Healthcare*, 2019.
- [17] Yoon, Jinsung, William R. Zame, and Mihaela van der Schaar. "Deep sensing: Active sensing using multi-directional recurrent neural networks." *International Conference on Learning Representations (ICLR)* 2018.
- [18] C. Lee, J. Yoon, M. van der Schaar, "Dynamic-DeepHit: A Deep Learning Approach for Dynamic Survival Analysis with Competing Risks based on Longitudinal Data," *IEEE Transactions on Biomedical Engineering (TBME)*, 2019
- [19] Alaa, Ahmed M., and Mihaela van der Schaar. "Attentive state-space modeling of disease progression." *Advances in Neural Information Processing Systems*. 2019.
- [20] Qian, Zhaozhi, Ahmed M. Alaa, Alexis Bellot, Jem Rashbass, and Mihaela van der Schaar. "Learning Dynamic and Personalized Comorbidity Networks from Event Data using Deep Diffusion Processes." *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [21] Anderson, Roy M., Hans Heesterbeek, Don Klinkenberg, and T. Déirdre Hollingsworth. "How will country-based mitigation measures influence the course of the COVID-19 epidemic?." *The Lancet* 395, no. 10228 (2020): 931-934.
- [22] <https://news.sky.com/story/coronavirus-warning-over-ibuprofens-effect-on-covid-19-sparks-debate-11958827>
- [24] Alexis Bellot and Mihaela van der Schaar, "Boosting transfer learning with survival data from heterogeneous domains", *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.