

Modeling the Impact of Social Determinants of Health and Public Health Interventions

A THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

BY

RAGHAV AWASTHI
(Phd18203)

Computational Biology

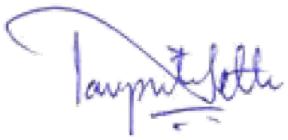
INDRAPRASTHA INSTITUTE OF INFORMATION TECHNOLOGY DELHI

NEW DELHI- 110020

June 3, 2023

THESIS CERTIFICATE

This is to certify that the thesis titled **Modeling the Impact of Social Determinants of Health (SDOH) and Public Health Interventions** submitted to **IIIT-D**, submitted by **Raghav Awasthi**, to Indraprastha Institute of Information Technology, Delhi, for the award of the degree of **Doctor of Philosophy**, is a bonafide record of the research work done by him under my supervision. The contents of this thesis, in full or in parts, have not been submitted to any other Institute or University for the award of any degree or diploma.



Dr. Tavpritesh Sethi
Thesis Supervisor
Associate Professor
Dept. of Computational Biology
IIIT Delhi, 110020

Place: New Delhi

Date: 09/11/2022

ACKNOWLEDGEMENTS

Not everything is provable

— Gödel's incompleteness theorem

With my Ph.D. thesis, I've completed a huge chapter in my life, and I'm grateful for all the kind individuals who helped me along the way. Every encounter gave me a new lesson and helped me become the person I am today.

I'd like to begin by expressing my sincere appreciation to my advisor, **Dr. Tavpritesh Sethi**, for his confidence, support, and constant encouragement to achieve greater heights. Through him, I have learned the importance of perseverance and self-confidence, which will serve me well throughout my life and profession. His compassion and generosity continue to inspire me to become a better person, while his technical skills and critical thinking motivate me to become a better researcher. I would love to express my sincere gratitude to the whole **BrainX** team, who have individually contributed to the advancement of my technical writing and critical thinking skills. In addition, I would like to thank my annual review and comprehensive exam committee members **Dr. Sriram K** and **Dr. Jaspreet Kaur Dhanjal** . I would want to convey my profound appreciation to **IIIT-Delhi**, which has served as my home for a number of years. The support offered by the Institute in terms of infrastructure, facilities, and helpful people continues to be among the finest I've encountered. I must also express my appreciation to **UGC** for helping me to get funds so that I may start my Ph.D. program without difficulty. A tremendous shout-out to all IIIT-Delhi administrators for being my only point of contact for administrative issues and for addressing them with a kind grin.

The previous several years would not have been possible without my family and friends' unwavering support. My parents **Sunil Kumar Awasthi** and **Greesha Awasthi** provided me the freedom to follow the route I wanted to and showed me that nothing is impossible if one has to believe in oneself. You taught me to accept everything with a pinch of fun, and that has been the guiding light on my roller coaster path.

And notable thanks to my wife and partner **Shreya Mishra**, who has been my person throughout this journey and continues to be so. I could not have asked for a greater life and research partner, brain-stormer, and support system. You have helped me imagine research and encouraged me anytime I felt apprehensive. Thanks for always believing in me. This would not have been possible without your help.

My sister, **Trapti Awasthi**, brother **Shivam** and brother-in-law, **Dr. Antriksh, Ankur Dwivedi, Harshit** and sister-in-law **Dr. Vartika Mishra**, who have supported me through thick and thin, and have been my confidence throughout the years. My niece **Anvi** and **Vedika** - their sweetness and kind words kept me going. My father-in-law **Dr. Avadesh Mishra**, mother-in-law **Dr. Shushma Mishra**, for supporting me during the Ph.D. journey.

The Parthya group **Aakash, Ajay, Anurag, Ankur, Aman, Harshit Awasthi, Harshit Tripathi, Piyush, and Pushpendra, Prabhash, Tejveer**, with whom I've spent an uncountable amount of time playing cricket, talking about cricket, and telling jokes. You were always there whenever I needs support. Please continue to be the same as you were in 2004.

Akash, My first school friend, taught me the value of friendship. He not only taught me how to improve my socio-behavioral habits but also assisted me in multiple activities like school courses, NCC parapdes, etc. He was my favorite sports companion. I'd want to take this time to thank Akash for covering the library fee so that I could acquire my 10th board exam admit card. This incident strengthened my trust in one another and helped us become close friends.

Ajay is a super spiritual person. He taught, "Why should we value emotions?" He taught me the lesson of silence and work. He is one of the most simple and humble people I know.

Aman is one of the most trustworthy people I know. I tried to instill from him the "no complain" mentality. He taught me that even the most challenging problems could be overcome by being calm.

Anurag, the group's leader, full of lofty aspirations, is always ready to switch us from tennis ball cricket to leather ball cricket. He inherited great leadership qualities from his family. I like his quality of rejoicing in the happiness of others. His contributions to the community deserve a standing ovation and widespread recognition. We all wish him the best of luck. May God grant him greater strength.

Harshit Awasthi, my brother and friend has always been my biggest supporter. We didn't communicate much over the phone or through text messages, but you are among the first individuals I think of when I am having a bad time.

Harshit Tripathi, has played an important role in my life; he is always easy to reach since I am extremely comfortable with him. I recall traveling with him to Kota and Delhi when I didn't have any close friends nearby. Without him, every game of cricket is really dreary. I admire his assistance with my Ph.D. admissions process.

Prabhash, known as "Mr. Perfectionist," has always motivated me to achieve higher career goals. I would like to thank him for providing all the motivation.

Pushpendra, always willing to help and to brim with positive energy, unconsciously taught me one of life's greatest lessons "Problems and obstacles may come and go, but our happiness should not depend on them because you are not driving yourself alone; many others are also involved with you." I would like to thank you for all your support and help.

Ankur and Piyush are two of our group's most intelligent members. I have gained

several academic skills from them.

Tejveer, the funniest person in my group, is always ready to have food. I would like to thank him for helping me many times. He is the one whom my parents recall most.

I would also like to express my appreciation to my labmates, **Dr. Aditya, Pradeep, Ridam, Akshaya, Jasmine, Ayushi, Alok**, On both the personal and the professional front, I gained invaluable insight from each and every one of you.

Aditya has been the most humble person I have come across. His perseverance and resilience toward work have always motivated me. He has always been a guiding figure in my life throughout my Ph.D. journey.

Ridam is like my younger brother, always ready to help. When it comes to empathy, everyone can take a lesson from him. Both Ridam and I are huge fans of Saurav Ganguli. Our talk about cricket is unforgettable. I have no doubt that he will reach any career objective he sets for himself since he is so dedicated to his work.

Pradeep is a very technical individual. Throughout my PhD., I considered Pradeep to be one of the most trustworthy and loyal labmates. He has excellent computer science abilities and is constantly eager to learn more.

I have had an opportunity to work with talented and enthusiastic bachelor's and master's students **Vibhor, Keerat, Vineet, Saif, Aniket, Arshita, Mehrab, Vaibhav, and many more** throughout my Ph.D. I'd want to express my gratitude to everyone.

During my Ph.D., I got an opportunity to work on different collaborations. I'd like to thank everyone who worked with me, including the **Swasti foundation, Google Research India, Facebook (now Meta), Pathcheck foundation, Stanford University, Adobe Research, and AIIMS Delhi**.

In addition, I would like to extend my gratitude to all of the experts who have volunteered their time to talk about SDOH and provide a message.

I am also obliged to my Phd freiends: **Neetesh Pandey,Indra Prakash Jha, Vivek Ruhela, Smriti Chawla, Sarita Poonia, Omkar Chandra, Madhu Sharma, Shiju, Shruti, Priyadarshini Rai, Samridhi, Chitrita Goswami, Sumit Patiyal, Nishant, Shubham, Anjali,**. Every single one of you spoke to me and understood all of my concerns like I was a kid.

Special appreciation to Pandey G and Indra sir for being my big brother and my rock of confidence, respectively. You were a tremendous help to me, both professionally and personally. I'm looking forward to having more fun with you.

If you hadn't been there for me through the good times and the bad, I never would have been able to complete my dissertation. My life has been impacted by each individual who has been named here, and their contributions have been essential to the completion of this work. I really hope that you are pleased with my dissertation.



ABSTRACT

KEYWORDS: Social Determinant of Health (SDOH) ; Antimicrobial Resistance; Mental Health; HIV; COVID-19; Homelessness; Machine Learning ; Bayesian Network; Causal Inference; Counterfactual Analysis, Reinforcement Learning

Social determinants of health (SDOH) are the non-medical factors that play a vital role in public health and form the basis of health policies. Building an effective public health policy is a complex endeavor that systematically understands multi-dimensional associations. The association of SDOH with public health outcomes is poorly understood due to the complex interplay of factors. Artificial Intelligence (AI) advancements have enabled models to make robust and explainable decisions in complex environments. Causal modeling and counterfactual analysis infer direct and indirect associations from observational data and rank policy indicators. Reinforcement learning (RL) is a paradigm for sequential decision-making under uncertainty that infers the policy after simulating sequential actions and accompanying rewards from the context. However, systematic application and utilization of these models are limited to public health settings. Through our work, we've made an integrative model that can measure complex interdependencies, bring together different kinds of knowledge about health system indicators, and use SDOH to guide public health interventions.

In our first contribution, we contribute to the discovery of potential interventions using an integrative machine learning framework incorporating structural causal models, counterfactual analysis, and predictive modeling to discover policy interventions. Using this framework, we found policy solutions for three use cases presented, i.e. (i)

antimicrobial resistance, (ii) mitigating the spread of HIV among women who work in the sex industry, and (iii) targeted interventions to improve mental health.

In our second contribution, we built a novel framework for optimizing potential interventions using reinforcement learning. Here we showcase a model to optimally allocate COVID-19 in the context of different SDOHs for states of India. This use case also aims to generalize the reinforcement learning framework for optimizing healthcare resource allocation.

Our final contribution is the real-world deployment of this policy-discovering and policy-optimizing AI model in response to COVID-19. Our machine learning models powered several web applications that provided forecasts of COVID-19 trajectory, mined emerging evidence from rapidly emerging COVID-19 literature, and for optimal allocation of COVID-19 vaccines.

LIST OF PAPERS BASED ON THESIS

1. Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.

2. Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.

3. Jha, I. P., Awasthi, R., Kumar, A., Kumar, V., Sethi, T. (2021). Learning the mental health impact of COVID-19 in the United States with explainable artificial intelligence: observational study. JMIR mental health, 8(4), e25097.

4. Awasthi, Raghav, et al. "Learning explainable interventions to mitigate hiv transmission in sex workers across five states in india." 34th ML4D Workshop at Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada (2020).

5. Pal, R., Chopra, H., Awasthi, R., Bandhey, H., Nagori, A., Gulati, A., ... Sethi, T. (2021). Predicting emerging themes in rapidly expanding COVID-19 literature with dynamic word embedding networks and machine learning. JMIR

6. Awasthi R, Devadiga A, Gupta M, Tewatia J, Sethi T Application of reinforcement learning in public health: a scoping review of literature JMIR Preprints. 24/07/2022:41381

7. Awasthi, Raghav, et al. "CovidNLP: A web application for distilling systemic implications of COVID-19 pandemic with natural language processing." MedRxiv (2020).

List of Abstracts

1. Sethi, T., Awasthi, R. (2020). Use of artificial intelligence based models for learning better policy for maternal and child health. European Journal of Public Health, 30(Supplement_5), ckaa165-291.

2. Awasthi, Raghav, et al. "Does Mental Health cause Homelessness or Homelessness causes Mental health?: A Bayesian Decision Network Analysis" 43rd Annual Meeting of the Society for Medical Decision Making, October 2021.

3. Ansari, M. S., Singh, V., Singh, P., Awasthi, R., Sethi, T. P. (2022, June). Spectral Pattern of Cough Sounds As an Aid to Diagnosis in Children with Common Respiratory Diseases. In PEDIATRIC PULMONOLOGY (Vol. 57, pp. S120-S120). 111 RIVER ST, HOBOKEN 07030-5774, NJ USA: WILEY.

Other Publications:

1. Mathur P, Mishra S, Awasthi R, Cywinski J, et al. (2022). Artificial Intelligence in Healthcare: 2021 Year in Review. DOI: 10.13140/RG.2.2.25350.24645/1

2. Mishra, S., Awasthi, R., Papay, F., Maheshawari, K., Cywinski, J. B., Khanna, A., Mathur, P. (2021). DiagnosisQA: A semi-automated pipeline for developing clinician validated diagnosis specific QA datasets.NeurIPS Data-Centric AI workshop 2021

3. Nagori, A., Kalia, A., Sharma, A., Singh, P., Bandhey, H., Ailavadi, P., Awasthi, R., Bhadra, W., Kaul, A., Lodha, R. and Sethi, T., Generalized Prediction of Shock in Intensive Care Units using Deep Learning.

4. Sethi, T., Kedia, S., Awasthi, R., Lodha, R., Ahuja, V. (2020). A Counterfactual Graphical Model Reveals Economic and Sociodemographic Variables as Key Determinants of Country-Wise COVID-19 Burden. medRxiv.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS **i**

ABSTRACT **vi**

LIST OF TABLES **xv**

LIST OF FIGURES **xix**

ABBREVIATIONS **xx**

1 Introduction **xxii**

1.1 Social Determinants of Health xxii

1.2 Incorporating Social Determinants of Health into Policy Modeling: Addressing Major Public Health Issues xxiii

1.3 Data Driven Models for Understanding the Impact of SDOH and Policy Development xxv

1.3.1 Compartmental Models xxv

1.3.2 Agent Based Models xxvii

1.3.3 Statistical Approaches for Modeling Correlations and Causation xxvii

1.3.4 Counterfactual Models xxix

1.3.5 Machine learning Approaches xxix

1.3.6 Supervised Machine Learning xxx

1.3.7 Unsupervised Machine Learning xxxi

1.3.8 Reinforcement Learning xxxii

1.4 The McNamara Fallacy in Data Driven Decision-Making: Importance of Considering SDOH xxxiii

1.5 Methodologies to Evaluate model-driven Policy xxxv

1.5.1 Matching xxxv

1.5.2 Difference-in-Differences xxxvi

1.5.3 Regression Discontinuity xxxvi

1.6 Objectives with Rationale of the Thesis: xxxvii
1.7 Thesis Organization and Contributions xxxvii

2 Discovering Potential Interventions: Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens xl

2.1 Preamble xl
2.2 Introduction xl
2.3 Methods xlii
2.3.1 Dataset xlii
2.3.2 Data Imputation and Discretisation xlv
2.3.3 WHO Declared Highly Important Pathogens Selection: . . . xlvi
2.3.4 Analysis xlvi
2.4 Results xlviii
2.4.1 Global Prevalence of Multiple Antibiotic Resistance of Critically Important Pathogens xlviii
2.4.2 Component Level Bayesian Network Revealed Strong Connection of Antibiotics with Independent WGI, Finance, GBD datasets xlix
2.4.3 Probabilistic Impact Quantification of Governance, Finance, Socioeconomic and Burden of Disease on Antibiotic Resistance xlix
2.4.4 Counterfactual Analysis Quantified Association between Ceftriaxone Resistance and Access to Immunisation and Obstetric care: li
2.4.5 Prediction of Antibiotic Susceptibility liii
2.4.6 Antibiotic Steward: An Interactive Web App to Discover AI-Driven Decision for AMR liv
2.4.7 Discussion and Conclusion lv

3 Discovering Potential Interventions: Learning Explainable Interventions to Mitigate HIV Transmission in Sex Workers Across Five States in India lix

3.1 Preamble lix
3.2 Introduction lix
3.3 Methods lx
3.3.1 Dataset lx
3.3.2 Structure Learning and Network Inference lxi
3.3.3 Discriminative Predictive Model for Condom Usage lxii
3.4 Results lxiii

3.4.1 Intervention Score and Vulnerability lxiii
3.4.2 Inferences from Bayesian Network lxiii

4 Discovering Potential Interventions: Mental Health in COVID-19: The Impact of Medical, Financial, Social, and Behavioral Factors lxviii

4.1 Preamble lxviii
4.2 Introduction lxviii
4.3 Methods lxix
4.4 Results lxx
4.4.1 Impact of Social Life and Work-Related Stressors lxxii
4.4.2 Impact of Symptoms and Comorbidities lxxii
4.4.3 Impact of Economic Factors lxxiv
4.5 Conclusion lxxiv

5 Optimizing Potential Interventions: Utilising SDOH for learning effective strategies for COVID-19 vaccine distribution using reinforcement learning lxxviii

5.1 Preamble lxxviii
5.2 Introduction lxxviii
5.3 Methods lxxx
5.3.1 Q Learning lxxxi
5.3.2 Deep Q-Network lxxxii
5.3.3 Actor-Critic using Kronecker-Factored Trust Region (ACKTR) lxxxii
5.3.4 Bandit Algorithms: Multi-armed bandits and Contextual Bandits problem lxxxiii
5.3.5 *VacSIM*: A Feed forward Approach lxxxvi
5.4 Results lxxxvii
5.4.1 Model Setup lxxxvii
5.4.2 Deep Reinforcement Learning Models lxxxix
5.4.3 Contextual Bandits Model xcii
5.5 Novel Approaches to Evaluate *VacSIM* Model xciii
5.5.1 Model Evaluation through Projection Scenario using SEIR models xciv
5.5.2 Model evaluation through learning the causal structure of simulated data obtained xciv

5.6 Discussion: xcvi
5.7 Limitations and future work xcviii
5.8 Conclusion c

6 Deployment:Web-Applications for Guiding Policy Interventions for Combating COVID-19 ci

6.1 Intoduction ci
6.2 CoVision: A Modeling and Simulation Suite for COVID Scenarios in India ci
6.3 CovidNLP: A Web Application for Distilling Systemic Implications of COVID-19 Pandemic with Natural Language Processing ciii
6.4 VacSIM: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning civ
6.5 EvidenceFlow:Predicting emerging themes in rapidly expanding COVID-19 literature with dynamic word embedding networks and machine learning cv
6.6 Conclusion cvi

7 Conclusion and Future work cviii

7.1 Conclusion cviii
7.2 Future Work cxii
7.2.1 Application of Reinforcement Learning in Public Health . . cxiii
7.2.2 Plug and play Resource Allocation Platform for Healthcare . cxiv
7.2.3 Integration of Social Media Data to Model Social Determinants of Health cxv

A Variables Description cxxxiii

LIST OF TABLES

1.1 Advantages and disadvantges of data-driven methods xxxiv

2.1 Age and temporal distribution of data. AMR Data spanned from 2004 to 2017 for 70 different countries. Adult (19 - 64) and elderly (65-84) years were the most predominant age groups and the majority of data represent samples collected between 2009 to 2017. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* xliv

2.2 Distribution of organisms in the random sample. Random sampling was conducted to avoid bias introduced in downstream models by the higher number of samples from high-income countries. A nearly equal number of samples were observed for each pathogen listed by the WHO global priority list of pathogens. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* xlvi

2.3 Inferences from pathogen-wise Bayesian network. We calculated the difference in probability of antibiotic resistance by changing levels of the variable, which was observed as a parent node of antibiotics in the Bayesian network. A significant effect of Government effectiveness, Access to Immunisation and Obstetric care, and Unsafe Wash sanitation on Antibiotic resistance were observed in all pathogens. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* lii

3.1 Prediction Model performance with 95% Confidence Interval. *Credits: Awasthi, Raghav, et al. "Learning explainable interventions to mitigate hiv transmission in sex workers across five states in india." 34th ML4D Workshop at Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada (2020).* lxvi

5.1 Features present in context space used in our models.*Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* lxxxix

5.2 Hyper-parameters used during Policy learning. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* xci

5.3 VacSIM model output on different bucket size. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* xciii

6.1 AI-powered Dashboards, which we developed across various stages of COVID-19. cii

LIST OF FIGURES

1.1 **Risk factors of premature mortality.** Among the various risk factors for premature mortality, health behaviours, physical environment, clinical treatment, and socioeconomic variables contribute 30%, 10%, 20%, and 40%, respectively. xxiii

2.1 **Flowchart showing the pipeline developed in this work.** Country-level world governance indicators, global finance indicators, and global burden of disease data were integrated with isolate-level AMR data sets. A Bayesian network structure-learning approach was used to identify the most actionable factors of AMR followed by prediction of antibiotic susceptibility at isolate level. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* xliii

2.2 **Bar plots showing the frequency distribution of samples in the data set.** (A) The top categories from 116 body tissue sources are shown. (B) Data were available from 70 middle and high-income countries from which the top five with their gender distributions are plotted here. (C) Sorted frequency distribution for facilities is shown here. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* xlv

2.3 **Principal component level association among AMR, WGI, GBD, and Finance.** A Bayesian network of principal components was learned from data to check the component-level association of AMR with Finance, GBD, and WGI. The first component of antibiotics (Antibiotics_Dim.1) was associated with the second component of GBD (Gbd_PC2) and WGI (Wgi_PC2), and the second component of antibiotics (Antibiotics_Dim.2) was associated with the first component of Finance (Finance_PC1). This indicates a strong association of AMR with WGI, Finance, and GBD. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* l

2.4 **Counterfactual analysis.** Bayesian network analysis showed that access to immunisation and obstetric care influenced the development of resistance in antibiotics, with the highest effect on ceftriaxone. To further understand this relationship, a counterfactual analysis, showed the change in ceftriaxone resistance in a situation where countries with lower care would counterfactually be switched to high care. For example, we observed that ceftriaxone resistance in Acinetobacter baumannii, in middle-income countries, could be reduced by 90%, if they had high access to immunisation and obstetric care. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* liii

2.5 **Antibiotic susceptibility prediction.** (A) AUROC of prediction models for Staphylococcus aureus. Error bar in the figure denotes the 95% CI of AUROC. (B) Feature importance was calculated using permutation-based methods, violin plot of the top-performing feature among all the given antibiotics is shown where "gbd val hep b prev(Hepatitis B prevalence)" was identified as a key feature. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.* lv

2.6 **Antibiotic steward Web Application.** Interactive web-app to visualise the Bayesian network and prediction model results for the pathogens E Coli, Acinetobacter baumannii, EnteroBacter Cloacae, Klebsiella pneumonia, Pseudomonas aeruginosa, Staphylococcus aureus. Bayesian network structure, area under the receiver operating characteristic curve (AUROC), and feature significance scores of prediction models are all shown graphically in the web-app. *Credits:https://antibioticsteward.tavlab.iiitd.edu.in/* lvi

3.2 **Network Structure.** Learned 101 Bootstrapped, Ensembled Bayesian Network. *Credits: Awasthi, Raghav, et al. "Learning explainable interventions to mitigate hiv transmission in sex workers across five states in india." 34th ML4D Workshop at Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada (2020).* . . lxv

4.1 **Methodology illustration.** Using survey data on mental health, a bootstrapped Bayesian network was trained. The inferences were derived using a joint probability calculation. High-impact factors on mental health were labelled as Interventional variables. lxx

4.2 **Sociodemographics of respondents who participated in the survey.** It can be seen that there was almost a similar representation from both genders. Age groups 25-75 years were predominantly captured in the survey. Most of the respondents had received a Bachelor's degree or above and were nearly equally distributed across geographies within the United States. HS: high school. *Credits: Jha, I. P., Awasthi, R., Kumar, A., Kumar, V., Sethi, T. (2021). Learning the mental health impact of COVID-19 in the United States with explainable artificial intelligence: observational study. JMIR mental health, 8(4), e25097.* lxxi

4.3 **Inferences from the Bayesian network.** The difference in inferred probability was calculated after conditioning the independent variables. A positive association implies a mental stress–inducing factor, whereas a negative association implies a mental stress reduction factor. The red circle shows the mean value, with green and blue showing confidence intervals. COPD: chronic obstructive pulmonary disease; SNAP: Supplemental Nutrition Assistance Program; TANF: Temporary Assistance for Needy Families. *Credits: Jha, I. P., Awasthi, R., Kumar, A., Kumar, V., Sethi, T. (2021). Learning the mental health impact of COVID-19 in the United States with explainable artificial intelligence: observational study. JMIR mental health, 8(4), e25097.* lxxiii

5.1 **VacSIM architecture.** A novel feed forwarded pipeline for policy learning of COVID-19 vaccine optimal distribution. The actions and rewards obtained from the Deep Reinforcement Learning models were fed forward into the training of Contextual Bandit algorithm so that faster optimal online decisions could be calculated for further dates taking in context, the ever changing demographics of the states. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* lxxxvii

5.2 **VacSIM model setup.** Flexible model setup for optimization of vaccine distribution policy in India using the VacSIM approach where components can be easily replaced with alternatives and adopted for diverse settings. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* lxxxviii

5.3 **Optimization.** Day-wise rewards at every iteration were recorded for all the states. The Mean and SD obtained over 26 day projections are shown as confidence bands. (Top) Smoothed and Increasing Reward Curve of DQN (Bottom) Smoothed and Increasing Reward Curve of ACKTR. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* xci

5.4 **Policy evaluation through SEIR model.** Additional projected infection cases prevented in next 45 days by following VacSIM driven approach instead of Naive approach. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* xcv

5.5 **Policy evaluation through causal model.** Ensemble averaged causal structure of the Bayesian network obtained from 501 bootstraps, using Hill Climbing optimizer for AIC (left) and BIC (right) as scoring functions. Vaccine Percentage obtained from the model was observed as a parent node of Susceptible cases, thus indicating the causality preserving nature of VacSIM simulations. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.* xcvi

6.1 **CoVision.** The interactive dashboard displays a state-by-state prediction of the COVID-19 trajectory using various models and scenarios. ciii

6.2 **CovidNLP.** The pipeline takes in WHO Database as input and creates text summarization on the article abstract. Word embeddings are generated from the vocabulary abstract to create word embeddings which serve as an input for tensorboard projection and word similarity. . . cv

6.3 **VacSIM.**The VacSIM dashboard displays vaccine allocation by the state daily in a geomap and compares models using SEIR-based projection (Chapter 4). The VacSIM model output was also evaluated with a baseline optimization model. VacSIM has a separate page for presenting model comparisons. cvi

6.4 **EvidenceFlow**.The Evidence flow Dashboard interactively presents the latent themes in the form of an unsupervised work embedding network. It also includes a word embedding projector and alluvial diagram functionality. cvii

7.1 **An Overview of RL's Public Health Applications.** a)In the 27 articles reviewed, we observed the application of RL in Public Health in Intervention modeling, Mitigation Policies, Personalised Treatment, and in general complex decision making. Five of the studies that were reviewed dealt with COVID-19. b)Year-wise number of papers selected for the reviews. *Credits: Awasthi R, Devadiga A, Gupta M, Tewatia J, Sethi T Application of reinforcement learning in public health: a scoping review of literature JMIR Preprints. 24/07/2022:41381* cxiv

ABBREVIATIONS

SDOH	Social Determinants of Health
FSW	Female sex worker
GBD	Global burden of disease
WGI	World governance indicators
BN	Bayesian Network
AIC	Akaike information criterion
BIC	Bayesian information Criterion
ML	Machine Learning
SVM	Support Vector Machine
RF	Random Forest
RL	Reinforcement Learning
DQN	Deep Q Network
ACKTR	Actor Critic using Kronecker-Factored Trust Region
ABM	Agent based Modeling
AMR	Antimicrobial Resistance
EM	Epidemiological Model
BMI	Body mass index
DTRs	Dynamic Treatment Regimes
SARSA	State–action–reward–state–action
T2DM	Type 2 Diabetes Mellitus
DTW	Dynamic Time Warping
SOTA	State-of-the-art
CCM	Classical Compartmental Model

IoT	Internet of Things
LUNA	Lung Nodule Analysis
CB	Contextual Bandits
CAD	Computer-Aided Diagnosis
CT Scan	Computed Tomography Scan
XCSR	Extended Classifier System with Continuous Real-coded Variables
CIAS	Chen Internet Addiction Scale
IF	Instantaneous Frequency
LIDC/IDRI	Lung Image Database Consortium/Image Database Resource Initiative
IF	Instantaneous Frequency
IoT	Internet of Things
RUL	Remaining useful life
SARSA	State–action–reward–state–action
CIAS	Chen Internet Addiction Scale

CHAPTER 1

Introduction

1.1 Social Determinants of Health

According to the *WHO* (Braveman *et al.* (2011), Braveman and Gottlieb (2014), Organization *et al.* (2008), Marmot and Wilkinson (2005)), social determinants of health (SDOH) are non-medical factors that influence health outcomes. These are the circumstances in which individuals are born, grow, work, and live while they mature. SDOH can also be considered in the context of aging, where the context is inclusive of cultural norms, economic strategies, development goals, political systems, and social policies. Given that health is a complex issue impacted by a variety of variables, the SDOH has been the most influential among a number of other relevant ones (Schroeder (2007)). SDOH, for example, are responsible with approx 60% of premature death, making them the most critical of all contributing factors, along with healthy behavior, clinical treatment, and the physical environment **Figure 1.1**.

SDOH are exemplified by the following:

- Housing, transit, and neighbourhoods that are safe
- Racism, prejudice, and violence
- Education, employment possibilities, and earnings
- Access to healthful meals and chances for physical exercise contaminated the air and water.
- Skills in language and literacy

The COVID-19 pandemic has clearly highlighted health disparities as key drivers of incidence rates and mortality(Bambra *et al.* (2020),Paremoer *et al.* (2021)). Disparities in health care are majorly outlined by the SDOH, which are related to the unfair

and avoidable disparities in health status that exist within and across countries. In all countries, social and economic factors have been shown to have a direct overbearing on population health(Fink *et al.* (2016)). This highlights the critical importance of SDOH in creating a comprehensive and sustainable healthcare system.

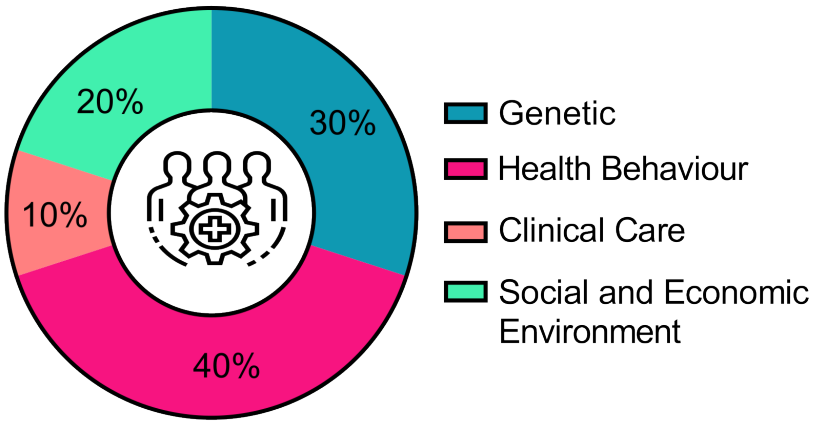


Figure 1.1: **Risk factors of premature mortality.** Among the various risk factors for premature mortality, health behaviours, physical environment, clinical treatment, and socioeconomic variables contribute 30%, 10%, 20%, and 40%, respectively.

1.2 Incorporating Social Determinants of Health into Policy Modeling: Addressing Major Public Health Issues

Developing policy models for public health issues is a critical challenge faced by policymakers and healthcare professionals globally, especially given that social determinants of health (SDOH) can account for up to 60% of health outcomes (Martins *et al.* (2021),Barten *et al.* (2007),Zhang *et al.* (2022)). Incorporating SDOH into policy modeling is essential to address major public health issues, as highlighted by the Global Sustainable Development Goals Report (GSDR) 2023 (Sachs *et al.* (2022)), which notes that the COVID-19 pandemic severely impacted global health in 2020 and 2021, caus-

ing an estimated 15 million deaths and disrupting essential health services in 92% of surveyed countries. According to WHO's Global Health Estimates (Arias *et al.* (2022)), despite the significant progress in controlling HIV transmission worldwide, in low-income countries, HIV/AIDS remains in the top 10 causes of death and accounted for more than 1.5 lakh death in 2020. Mental health disorders have also become a major public health problem for people worldwide, with 10% of the Indian population suffering from mental health conditions, the most common being depression and anxiety (Kallakuri *et al.* (2018)). Mental health illnesses are predicted to cost the world economy 3 trillion dollars by 2030 (Arias *et al.* (2022)). Antimicrobial resistance (AMR) is another major public health problem, predicted to account for 10 million deaths yearly by 2050 (Medina *et al.* (2020)). Drug-resistant bacteria have emerged due to antibiotic abuse and overuse in India and the world, responsible for thousands of annual fatalities. Resource allocation is a critical challenge in public health, especially in low- and middle-income countries with limited resources and a high disease burden. Social determinants of health (SDOH) can help develop an optimal healthcare resource allocation model. By considering SDOH, healthcare providers can allocate resources to priority areas, such as communities with a high prevalence of chronic diseases, and ensure equitable distribution to promote health equity and reduce disparities. Given the significance of social determinants of health (SDOHs) in policy, it is crucial to establish healthcare policies with a more data-driven approach. Utilizing data-driven models can provide a better understanding of the intricate interactions between these factors, leading to challenges in decision-making. The recognition of the need to use large, heterogeneous datasets in more comprehensive data-science frameworks to identify, quantify, and evaluate the effect of social determinants on health outcomes forms a primary tenet of this thesis work.

1.3 Data Driven Models for Understanding the Impact of SDOH and Policy Development

Modeling the impact of SDOH can involve many different approaches, ranging from classical mathematical models to state-of-the-art machine learning. First, we review some of the most commonly used approaches for developing policy models for decision-making in public health along with their standard use cases with advantages and disadvantages (Table 1.1). We also highlight the key studies that have outlined the application of such approaches to areas similar to this thesis (Hogan *et al.* (2021), Exworthy (2008), Galanis and Hanieh (2021), Carey and Crammond (2015)).

1.3.1 Compartmental Models

Classical mathematical modeling approaches such as compartmental modeling involves specifying a mathematical representation of the mechanism that drives the process of interest. These have traditionally been used in modeling contagion dynamics in infectious diseases and stages of illness development. (Nawaz *et al.*. Collignon *et al.* (2018). Cotta *et al.* (2020).Goldsztejn *et al.* (2020), Truszkowska *et al.* (2021)). More recently, these approaches have also been extended to be used for modeling the SDOHs assuming social behaviours may follow similar contagion dynamics (Galanis and Hanieh (2021), Abou-Ismail (2020), Brauer (2008)). This allows direct incorporation of factors such as income inequalities into the spread of epidemics modeled using differential equations. These works argue that health outcomes are essentially social phenomenon and leaving the SDOHs out of compartmental models limit their real world applicability. In these models, the overall population is assumed to consist of compartments, such as Susceptible, Infectious, or Recovered (S, I, R) with individuals moving across the compartments with rate-parameters that can be modeled from data. An SIR model is the simplest of all specifications, which form the basis of more complex compartmental models such as

SEIR, with the addition of an "Exposed" compartment representing the period between exposure and infection (incubation period. The order of the labels often shows the flow patterns between the compartments; for example, SEIS stands for susceptible, exposed, infected, and then susceptible again.

An example mathematical representation of an SEIR model is the following: We divide the entire population into blocks S,E,I,R i.e. $N = S+E+I+R$ and all four blocks can be modeled in by following equations

$$\begin{aligned} \frac{dS}{dt} &= -\beta SI \\ \frac{dE}{dt} &= \beta SI - \sigma E \\ \frac{dI}{dt} &= \sigma E - \gamma I \\ \frac{dR}{dt} &= \gamma I \end{aligned} \tag{1.1}$$

$$R_0 = \frac{\beta}{\gamma} \tag{1.2}$$

Whereas deterministic ordinary differential equations are the most commonly used method for these models, stochastic (random) frameworks, which are more realistic but substantially more difficult to assess can also be used. In compartmental models, such as the COVID-19 incidence model described in **Equation 1.1**, heterogeneity in the population is introduced by dividing the compartments into smaller subgroups (Kerr *et al.* (2021)), such as population density, economy, or population with pre-existing conditions, where each subgroup may have different transition rates between states. Notably, the compartmental model is developed at the population level assuming that the individuals in each compartment are homogeneous (top-down approach).

1.3.2 Agent Based Models

As oppose to compartmental models, agent-based models, also known as individual- or micro-simulation-based models, introduce heterogeneity at the individual level, whereby each individual may be assigned unique characteristics, risk profiles and a personal pattern of contacts with other individuals within the population (bottom-up approach)(Tracy *et al.* (2018), Badham *et al.* (2018). Therefore, these are higher-resolution versions of mechanistic models used to define and simulate outcomes based on behavior at the individual level. (Janssen and Ostrom (2006), Maglio and Mabry (2011)). ABM considers each individual in a group to be unique and may contain or represent information about their interactions with other agents.

1.3.3 Statistical Approaches for Modeling Correlations and Causation

The well known dictum: "correlation is not causation" cannot be overemphasized for modeling data-driven policy. Correlation is the relationship between two elements (or variables) in which one does not always cause the other. However, in the context of policy models, researchers are typically focused on learning causal influences that may enable interventions. A causal influence is said occurs when one variable or circumstance causes another(Pearl and Mackenzie (2018),Pearl (2009)). However, while it is simple to uncover evidence of a link between two things, it is usually challenging to demonstrate causality. occasionally, two things might have a link without one being the cause of the other. Even if there is a causal link between variables, it might be difficult to determine which one influences the other. For instance, there may be a link between people's mental health and physical health. Still, it is unclear which variable drives the other: does a good mental state enhance physical health, or does excellent physical health improve people's mental health(Biddle *et al.* (2019))? Some methods

of study may provide proof of causal linkages between two variables, whilst others can just identify correlations. Randomized controlled trials, for instance, may give strong evidence of causal linkages.

However, conducting randomized controlled trials are both expensive and difficult, hence data-driven approaches to determine causality from observational data are gaining importance (Petersen and van der Laan (2014), Cheng *et al.* (2020), Kohl *et al.* (2017)). In this thesis, we have extenxively used the Bayesian network framework for discovering potentially causal relationships. A Bayesian network is a class of graphical causal models that used to represent causal interactions(Pearl and Mackenzie (2018) Pearl (2009),Pearl (2014)). Since it is a directed edge graph, parent-child interactions may be read as cause-and-effect under strict restrictions. A directed acyclic graph (DAG) G is defined as a triple $N = \{X, G, P\}$ over a collection of random variables X. It encodes the underlying joint probability distribution P, which may be factored as the product of the probabilities of each node v given its parents $par(v)$. A Bayesian network explicitly models the confounding (fork motifs), mediation (chain motif) and inter-causal (collider motif) effects.

$$P(X) = \prod_{v \in V} P(x_v | x_{pa(v)}) \tag{1.3}$$

where v corresponds to the random variable *v* and *par(v) are the set of parent variables.*

1.3.4 Counterfactual Models

The counterfactual analysis measures the hypothetical 'what-if' question without real-world intervention. For example what would be the COVID-19 death rate if we were able to successfully increase non pharmaceutical interventions (NPIs) or reduce vaccine inequity. Let us suppose, 0 denotes population of men, 1 denotes population of women. Variable Y_j denotes 'mortality rate' that is a function of Non-pharmaceutical Interventions such as lockdown. The conditional distribution function of mortality rate for men and women defined as $F_{Y_0|X_0}(y|x)$ and $F_{Y_1|X_1}(y|x)$ respectively. It describes the random assignment of NPI to patients with given traits x. Now, let $F_{Y\langle 1|1\rangle}$ and $F_{Y\langle 0|0\rangle}$ signify the observed distribution function of NPIs for women and men, and let $F_{Y\langle 0|1\rangle}$ represent the distribution function of NPIs that would have been obtained for women if they received the men's NPIs

$$F_{Y\langle 0|1\rangle}(y) = \int_{\chi_1} F_{Y_0|X_0}(y|x) dF_{X_1}(x) \tag{1.4}$$

Equation 1.4 represent the counterfactual distribution since it does not result from any observed population distribution. The objective of the counterfactual analysis model is to learn the counterfactual distribution, which can be further used to infer the counterfactual estimates **Equation 1.5**.

$$F_{Y\langle 1|1\rangle} - F_{Y\langle 0|0\rangle} = [F_{Y\langle 1|1\rangle} - F_{Y\langle 0|1\rangle}] + [F_{Y\langle 0|1\rangle} - F_{Y\langle 0|0\rangle}] \tag{1.5}$$

1.3.5 Machine learning Approaches

Machine learning is a sub-field of artificial intelligence that takes a data-driven approach "learn" models that can fit complex data with nonlinear properties and interactions. The use of machine learning models for public health data is increasing in the research

community. ML models have the potential to significantly advance our knowledge of health and our ability to intervene. ML approaches present new opportunities to learn interventions from real-world data that include health, behavioural, and social sciences factors. Although the use of ML for modeling socioeconomic determinants of health is in infancy, there are encouraging studies showing its potential impact. (Irvin *et al.* (2020), Kino *et al.* (2021),Segar *et al.* (2022) Mooney and Pejaver (2018),Lewis (2022) Sethi *et al.* (2019)). Machine Learning models are commonly categorised into 3 approaches 1)Supervised 2)Unsupervised, and 3)Reinforcement Learning which are briefly introduced next.

1.3.6 Supervised Machine Learning

Supervised machine learning (Singh *et al.* (2016))approach is used when the data has the availability of labels, e.g. healthy/diseased, which are then used to train a model with the objective of predicting labels for fresh and unseen data. The training and fine-tuning of the model relies upon the adjusting model weights on a training set to minimize a user-specified loss function and to achieve a balance between the bias and variance tradeoff of the model achieved. The training set consists of inputs and the labelled outputs, enabling the model to acquire knowledge as it sees more data. Supervised machine learning further comes in two flavors: (i) $\boldsymbol{classification}$, i.e. when the labels for the dependent variable are discrete classes, and (ii) $\boldsymbol{regression}$, i.e. when the dependent variable to be modeled is continous in nature. Common classification methods include linear classifiers, support vector machines (SVMs), decision trees, k-nearest neighbour algorithms, and random forests. Common regression approaches include linear regression, polynomial regression and machine learning based methods such as support vector regression, random forest regression and markov models. Supervised machine learning has been used extensively to develop predictive models for diagnosis, outcomes, prognosis and treatment (Ranapurwala *et al.* (2019)). Commonly

used models have been developed to predict positive or negative health outcomes, such as breast cancer and 5- or 10-year cardiovascular disease risks (Quante *et al.* (2012), Gail *et al.* (2007)) on factors such as age, gender, smoking status, blood pressure, and cholesterol levels. Other interesting use cases have targeted assessing the likelihood that a person would have an opioid overdose within a certain time frame (Wilson *et al.* (1998)). Using such predictive algorithms, a physician, counsellor, health education expert, social worker, policymaker, or other professional may identify persons who are at a greater chance of experiencing a certain undesirable event within a specified time period. Further, such predictive models may also be used to educate target audiences so that people can comprehend their specific risks based on their unique demographics, behavioural traits, or environmental / social situations.

1.3.7 Unsupervised Machine Learning

Unsupervised learning (Gentleman and Carey (2008)) approaches are applied when the data lacks labels, hence the model does not receive any supervision in the form of labels. The underlying motivation for unsupervised machine learning models is usually to discover sub-classes or clusters which may be meaningful for decision making in the clinical or public health settings. Another motivation for discovering clusters is to develop tailored interventions such as precision medicine approaches or in policy modeling. In the public health and SDOH settings, unsupervised machine learning has previously been used to cluster health-related behavior patterns and demographics, which sought to identify various subgroups of a population based on the health behaviours of alcohol consumption, nutrition, and smoking, while also taking into account the influence of the parameters sex, age, and education. Also, unsupervised machine learning was used to evaluate the transferability of public health interventions. Cluster analysis was used to identify groupings of nations that had the best chance of successfully transferring a certain intervention(Rabel *et al.* (2019)).

1.3.8 Reinforcement Learning

Reinforcement learning (Sutton and Barto (2018), Kaelbling *et al.* (1996)) is a subset of machine learning which makes a sequence of decisions based on the environment. Policy optimization relies heavily on these interactions between an "agent" and its "environment", the latter may often be murky and ever-changing. Therefore, the agent learns how to accomplish a task in a potentially complex and uncertain environments. This setup resembles a game and where a player gets rewards or penalties for their activities and the goal is to accomplish the objectives while maximizing overall reward (Chen (2021), Iwata (2016)). Although reinforcement learning offers the most advanced form of machine learning applied to healthcare settings, it is difficult to implement and evaluate in the real world settings. However, the need for the development of intelligent public health systems where navigating complex decision making becomes possible (Chelouati *et al.* (2021)) has led to a rapid increase in the implementation of RL in improving public health over the last few years. These applications include fitness(Kordonouri and Riddell (2019)) , glycemic control and diabetes treatment policy. In the HeartSteps fitness application, a reinforcement learning algorithms was used (Ho *et al.* (2018)) to boost individual-level fitness regimens through exercise reminder messages. This algorithm encouraged physical exercise through context-tailored activity suggestions five times per day (Hassouni *et al.* (2018)). Similarly, mobile phone apps combined with a reinforcement learning algorithm have been developed for glycemic control in large groups of diabetic individuals (Javad *et al.* (2019)). Another model for diabetes detection model used the well-known RL Method, Q learning, using the PIMA Indian Women diabetes dataset. To identify patients with Type 2 Diabetes Mellitus, the model develops an off-policy-based RL and instructs the learning agent to choose the best policy for three criteria (glucose level, subject age, and BMI) (Biswas *et al.* (2021)). While COVID-19 was affecting the health of millions, the RL was trained so that artificial agents might look for COVID-19 control strategies that would benefit

the general public. Furthermore, the RL agent suggested a combination of travel restrictions and lockdowns, as well as higher-intensity methods than those employed by governments (Arango and Pelov (2020)). RL helped to enhance mitigation policies that reduce the financial impact while not overburdening hospital capacity.(Kompella *et al.* (2020), Kwak *et al.* (2021*a*) Kwak *et al.* (2021*a*), Martinho *et al.* (2021), Liu *et al.* (2021)),Liu *et al.* (2021)).

Despite the growing utilisation Rl in public health, there are significant knowledge gaps regarding the use of RL algorithms in policy modelling.

1.4 The McNamara Fallacy in Data Driven Decision-Making: Importance of Considering SDOH

Decisions entirely governed by data and quantitative approaches might produce cognitive biases. The McNamara fallacy is the cognitive bias generated when decision-makers only consider quantitative data and measurements while ignoring vital qualitative data. The McNamara fallacy can be countered by the social determinants of health (SDOH), which offer a framework for comprehending the intricate interactions between social, economic, and environmental factors that affect health outcomes. Decision-makers can guarantee that health interventions are culturally and socially acceptable and equitable by considering SDOH and the larger social and political context in which they are implemented. For example, Vaccination policy solely depends on cost-effectiveness outcomes. It may need to take into account the concerns and beliefs of the communities that are being targeted, leading to poor uptake and effectiveness of the intervention. On the other side considering, by considering SDOH, decision-makers can identify and address the underlying social and cultural factors that may be influencing vaccine uptakes, such as mistrust of the healthcare system or language barriers, and design interventions that are more responsive to the needs and preferences of the communities they are serv-

Method	Advantages	Disadvantages
Compartmental Models	- Allows for the incorporation of factors such as income inequalities into the spread of epidemics and modeling of social determinants of health (SDOHs). - Provides a mathematical representation of the mechanism driving the process of interest.	- Assumes homogeneity within compartments, limiting real-world applicability. - May require complex differential equations to model more complex scenarios.
Agent-Based Models	- Introduces heterogeneity at the individual level, allowing for unique characteristics, risk profiles, and personal patterns of contacts. - Provides higher-resolution versions of models to simulate outcomes based on individual behavior.	- Requires detailed data on individual characteristics and interactions, which may be difficult to obtain. - Computationally intensive due to the large number of individual agents and their interactions.
Statistical Approaches	- Focuses on uncovering associations to enable interventions.	- Difficult to establish causality in observational data.
Bayesian Network	- Bayesian network framework can represent causal interactions and discover potentially causal relationships.	- Relies on assumptions.
Counterfactual Models	- Analyzes hypothetical "what-if" questions without real-world interventions. - Helps understand the potential impact of interventions or changes in specific variables.	- Requires counterfactual distribution estimation, which may not result from observed data. - Relies on assumptions and may not fully capture the complexity of real-world scenarios.
Supervised Machine Learning	- Uses labeled data to train models for predicting labels of unseen data. - Can be used for classification and regression tasks.	- Requires availability of labeled data, which may be limited or expensive to obtain. - Performance heavily relies on the quality and representativeness of the training data.
Unsupervised Machine Learning	- Discovers sub-classes or clusters in data for decision-making purposes. - Used for identifying patterns or groupings within populations.	- Relies on data without labels, which may limit the interpretability of the results. - Results may vary depending on the choice of clustering algorithm and parameters.
Reinforcement Learning	- Enables decision-making in complex and uncertain environments. - Used for optimizing policies and interventions in public health.	- Implementation and evaluation can be challenging in real-world settings. - Requires extensive computational resources and expertise.

Table 1.1: Advantages and disadvantges of data-driven methods

ing. In the next section, we provide background for evaluating policy modeling methods that might help assess bias-free policy evaluation.

1.5 Methodologies to Evaluate model-driven Policy

Pre-deployment evaluation of model-guided public health policies (PHP) is essential. Randomized control trials (RCTs) are useful for establishing the societal impact of a policy, but they are expensive and difficult to conduct (Deaton and Cartwright (2018)). As observational data are more abundant and less expensive, numerous alternative, observational data-driven approaches for policy evaluation have been proposed (Abadie *et al.* (2010), Campbell and Stanley (2015), Fougère and Jacquemet (2021) Athey and Imbens (2017), Stuart (2010)). These strategies can be broadly divided into the three groups that are explained below.

1.5.1 Matching

The primary objective of the matching exercise is to create cohorts that resemble clinical trials in which recipients and non-recipients share as much similarity as possible (Stuart (2010)). Matching works by selecting subjects that eliminates disparities in the distributions of observable features and is based on two assumptions- (i) conditional independence, which assumes that the assigned invervention status is conditionally independent from the potential outcomes and (ii) overlap which implies that each person has a non-zero probability (defined as strictly between 0 and 1) of becoming an intervention beneficiary a priori, regardless of their characteristics or whether they are genuinely (i.e. a posteriori) a beneficiary of the intervention. When these two premises are fulfilled, the approach consists of averaging the disparities between each beneficiary's outcome and the average of the non-beneficiaries outcomes who seems to be "near" in relation to observable attributes (gender, age, level of education, etc.) The

distance formula (Mahalanobis or the Euclidean) or even a propensity score, which is simply defined as the choice of the beneficiary under consideration's "closest neighbors" can be made using the probability of becoming a beneficiary of the intervention provided the observable factors describing the individual. Finally, this probability can be estimated using a probit or logit model.

1.5.2 Difference-in-Differences

The difference-in-differences (DID) approach is a quasi-experimental technique for evaluating the changes in outcomes over time between a treatment group and a control group.. DID is used to estimate the impact of a specific intervention or treatment (such as a health policy, large-scale scheme or a law) by comparing the changes in outcomes over time between the intervention group and the control group.(Dimick and Ryan (2014)). DID first calculate the before-after difference in the treatment group's outcomes (D1). In comparing the same group to itself, D1 takes into account variables in that group that remain consistent throughout time. Then, to capture time-varying factors, difference-in-differences (D2) takes the before-after difference in the control group, which was exposed to the same set of environmental conditions as the treatment group. Therefore, by deducting the second difference from D1, D2 filters all time-varying components from D1. The effect estimate is known as the difference-in-differences.

1.5.3 Regression Discontinuity

Regression Discontinuity (RDD) is a crucial method for divulging the causal impact of policies. The RDD is a quasi-experimental mechanism to evaluate the effect of interventions based on an eligibility threshold. (Venkataramani *et al.* (2016)). The applicability of RDD in policy evaluation is very useful because clinical decisions are often driven

by decision rules premised on specific thresholds (Venkataramani *et al.* (2016)). For instance, precise laboratory measures and policy eligibility requirements provide eligibility thresholds for certain treatments or health services. This threshold might be age (for Treatment). RDD predicts local average treatment effects close to the cutoff point, when treatment and comparison units are most comparable. The data points above and below the cutoff resemble one another more and more as they approach the cutoff (Imbens (2004)).

Therefore, with this background on a myriad of modelling and evaluation approaches, we next define the objectives and rationale of this thesis focused on learning and evaluation of policy from observational data.

1.6 Objectives with Rationale of the Thesis:

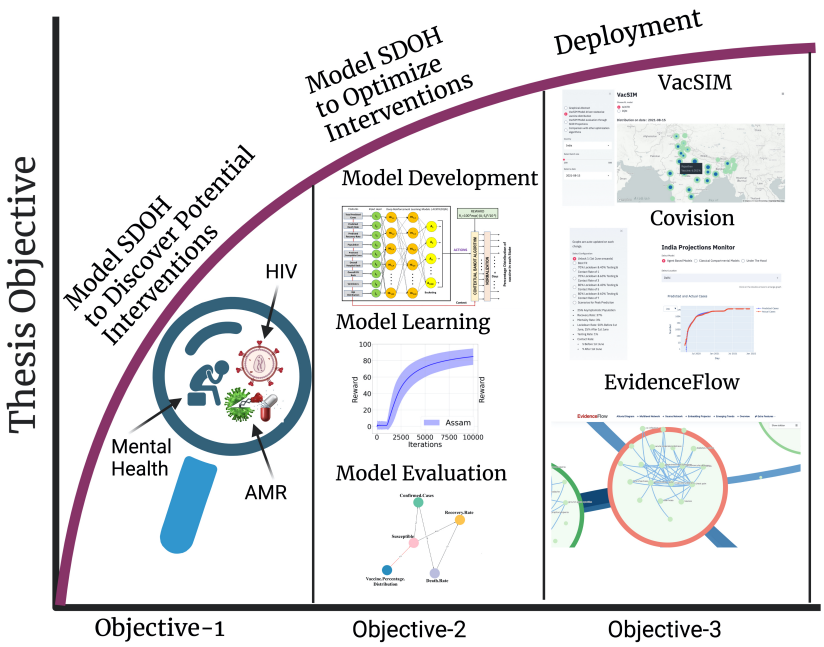
We propose integrative modeling of public health challenges to include SDOHs with the following objectives:-

To,

1. discover policy influences from data, i.e. identify and quantify SDOH influences using an explainable graphical model approach.
2. optimize interventions by providing a reinforcement learning based framework for allocating vital healthcare resources efficiently.
3. deploy and test the real-world applications of above objectives for combating COVID-19.

1.7 Thesis Organization and Contributions

Following a comprehensive background on the problem statement and thesis objectives, this section provides an overview of the remaining thesis's organizational structure and the noteworthy contributions to attaining the stated objectives.



.

.

Figure 1.2: **Organizational structure and contributions of the thesis.** The thesis comprises seven distinct chapters, including the introduction chapter, each chapter presenting a case study published in peer-reviewed journals or conferences which employs machine learning and artificial intelligence to tackle public health issues using SDOH. The common thread is using innovative AI solutions to guide public health decision-making using SDOH

.

1. **The second chapter** uses graphical model techniques to estimate the impact of health system factors on antimicrobial resistance in priority pathogens. The study identifies key factors associated with antimicrobial resistance by analyzing data from more than 6 lakh isolates. This information can be used to develop policy and public health interventions to mitigate antimicrobial resistance.

2. **The third chapter** employs graphical model to identify key factors contributing to HIV transmission in female sex workers across five states in India. The study's findings can guide public health policies and interventions to address HIV transmission among female sex workers.

3. **The fourth chapter** investigates the impact of medical, financial, social, and behavioral factors on mental health during the COVID-19 pandemic. Using graphical techniques to analyze data collected from online surveys, the study identifies factors contributing to poor mental health outcomes. This information can be used to develop interventions and policies to mitigate the pandemic's mental health impact.

4. **The fifth chapter** focuses on optimizing COVID-19 vaccine distribution using reinforcement learning. The study utilizes data on social determinants of health to develop an AI algorithm that identifies the most effective vaccine distribution

strategies. This information can inform policy and public health interventions to ensure equitable vaccine distribution.

5. **The sixth chapter** involves deploying web applications to guide policy interventions to combat COVID-19. We utilized machine learning techniques to develop web-based tools to inform policy and interventions.

6. **The seventh chapter** presents the thesis's conclusion and discusses the research's future scope.

CHAPTER 2

Discovering Potential Interventions: Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens

2.1 Preamble

Parts of this chapter have been published in:

"*Awasthi, R., Rakholia, V., Agrawal, S., Dhingra, L. S., Nagori, A., Kaur, H., & Sethi, T. (2022). Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens. Journal of Global Antimicrobial Resistance, 30, 133-142*".

2.2 Introduction

WHO has alerted that the worldwide number of deaths attributable to antimicrobial resistance (AMR) may increase from 700,000 to 10 million by 2050 if un-checked [1]. Sulphonamides, the first effective antibiotic group, were introduced in 1937, and the first resistance was reported in late 1930 (Davies and Davies (2010)). The story of the emergence of reports of antimicrobial resistance has been similar to other classes of antibiotics. Even though we have witnessed a significant improvement in global health, millions out there do not have adequate access to health services. Factors such as rising incomes, incessant infectious diseases, poor or marginalized populations having limited access to primary health care, and consuming antibiotics without any prescription have

[1] https://apps.who.int/iris/handle/10665/112642

been compounding the problem of antimicrobial resistance in low-income and middle-income countries. Ensuring proper sanitation, better governance, increased focus on public health care, and regulation of the private health care sector is a must to combat the spread of antimicrobial resistance (Chandler (2019)). While its disastrous effect on health outcomes can be understood by an estimated million deaths, the cost of treating resistant infections is estimated to reach the US $100 trillion dollars by the year 2050 (Smith *et al.* (2012)) (Naylor *et al.* (2018)). It is essential to develop evidence-based policies for AMR mitigation. The role of government (policymakers) is multifold in improving antibiotic stewardship policies, increasing antibiotic use surveillance, and financing these along with funding the development of new drugs (Laxminarayan *et al.* (2013),Collignon *et al.* (2018),Maragakis *et al.* (2008)). There arises a need to formulate governance frameworks to help the policymakers in designing, monitoring national action plans for tackling antimicrobial resistance at all levels: local, regional, national and global. Tackling antibiotic resistance is a high priority for WHO. A global action plan on antimicrobial resistance Chua *et al.* (2021), including was endorsed at the World Health Assembly in May 2015. The global action plan aims to ensure the prevention and treatment of infectious diseases with safe and effective medicines. WHO has been leading multiple initiatives to address antimicrobial resistance like World Antimicrobial Awareness Week, The Global Antimicrobial Resistance Surveillance System (GLASS) [2], Global Antibiotic Research and Development Partnership [3](GARDP), Interagency Coordination Group on Antimicrobial Resistance (IACG) [4] with the motivation to improve awareness and understanding of antimicrobial resistance, to strengthen surveillance and research, to reduce the incidence of infection, to optimize the use of antimicrobial medicines and to ensure sustainable investment in countering antimicrobial resistance. WHO published a priority pathogen list in 2017. This list consisted of 12 families of bacteria that pose the greatest threat to human health. The priority

[2](http://www.who.int/glass/en/)
[3](http://www.gardp.org/)
[4](http://www.who.int/antimicrobial-resistance/interagency-coordination-group/en/)

pathogens were divided into three categories; critical, high, and medium priority depending upon the extent of prevalence of antimicrobial resistance and the threat they posed. The critical pathogens include Acinetobacter, Pseudomonas, and various members of the Enterobacteriaceae family, including Klebsiella, E.coli, Serratia, and Proteus. The list was meant to be a 'new tool to ensure R&D responds to urgent public health needs. The need for models that can predict and understand antimicrobial resistance has never been more. To our knowledge, our study presents, for the first time, an exhaustive analysis of antimicrobial resistance in critical pathogens and the factors affecting their prevalence. Here, we quantified the relationships between a wide range of factors, from governance, healthcare access, finance, population density, to the type of sample and gender and Global AMR.

2.3 Methods

In this study, we developed a machine learning framework **Figure 2.1** to understand the robust connection amongst AMR, disease burden, governance, finance, and different socioeconomic indicators. We then used our results to better understand and predict the susceptibility of antibiotics amongst critical pathogens.

2.3.1 Dataset

The AMR base data was extracted from the AMR surveillance competition [5] which had 633820 isolates. The characteristics for the same have been displayed in **(Table 2.1, Figure 2.2)**. We then integrated the AMR data with the World Governance Indicators (WGI) data, the Global Burden of Disease Study (GBD) data, and the Finance data. WGI data covers six dimensions of governance for over 200 countries over the period

[5](https://sagebionetworks.org/research-projects/wellcome-trust- partners-with-sage-bionetworks-to-launch-wellcome-data-re-use-prizes, https://amr.theodi.org/, https://atlas-surveillance.com,)

1996-2018. The six dimensions of governance are:1) Control of Corruption, 2) Voice and Accountability, 3) Political Stability and Absence of Violence/Terrorism, 4) Government Effectiveness, 5) Regulatory Quality, 6) Rule of Law. These were then merged with the AMR data set by country and year, thus giving us the resultant AMR WGI data set. The GBD covariates data set consists of 334 covariates data files for the period 1980-2017. Types of covariates used include socioeconomic, demographic, Access to Immunisation and Obstetric care, climate, and food consumption data., these were then merged with the AMR data by country, age, and gender, taking only the value of covariates, thus giving us the resultant AMR WGI GBD data set. We reshaped the finance data such as columns country, year and then merged with the above- mentioned AMR WGI GBD data set by country and year. We removed variables with more than 10% missing data, resulting in a dataset that spans from 2004 to 2017.

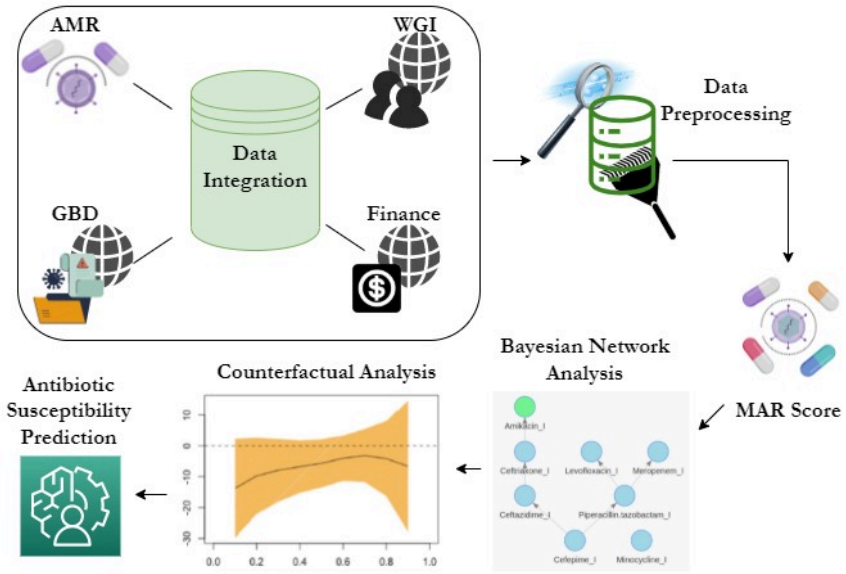


Figure 2.1: **Flowchart showing the pipeline developed in this work.** Country-level world governance indicators, global finance indicators, and global burden of disease data were integrated with isolate-level AMR data sets. A Bayesian network structure-learning approach was used to identify the most actionable factors of AMR followed by prediction of antibiotic susceptibility at isolate level. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

	Female	Male	Overall
	278128 (43.88%)	348136 (54.93%)	633820
Age.Group			
0 to 2 Years	15329 (42.83%)	20012 (55.91%)	35793
13 to 18 Years	6348 (47.06%)	7043 (52.21%)	13490
19 to 64 Years	130162 (44.04%)	163201 (55.22%)	295537
3 to 12 Years	12613 (46.88%)	14052 (52.22%)	26907
65 to 84 Years	86561 (41.63%)	119801 (57.62%)	207922
85 and Over	23406 (54.29%)	19364 (44.91%)	43114
Year			
2004	9433 (46.93%)	10655 (53.01%)	20101
2005	10473 (48.04%)	11274 (51.71%)	21801
2006	13967 (46.65%)	15876 (53.03%)	29940
2007	18264 (45.70%)	21409 (53.57%)	39964
2008	16303 (44.33%)	19925 (54.18%)	36773
2009	18575 (44.55%)	22514 (54.00%)	41692
2010	14476 (44.59%)	17341 (53.42%)	32462
2011	11622 (44.66%)	13931 (53.53%)	26023
2012	22432 (42.16%)	29477 (55.40%)	53206
2013	29962 (42.73%)	38850 (55.40%)	70125
2014	30492 (43.23%)	39482 (55.98%)	70529
2015	27992 (43.21%)	36104 (55.73%)	64785
2016	29744 (42.74%)	39290 (56.45%)	69598
2017	24393 (42.93%)	32008 (56.33%)	56821

Table 2.1: Age and temporal distribution of data. AMR Data spanned from 2004 to 2017 for 70 different countries. Adult (19 - 64) and elderly (65-84) years were the most predominant age groups and the majority of data represent samples collected between 2009 to 2017. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

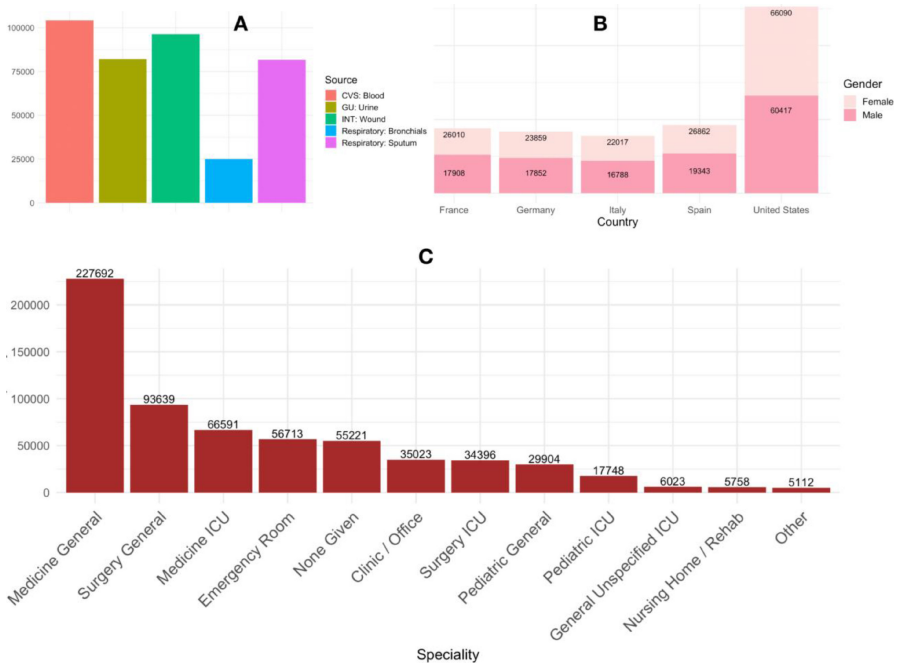


Figure 2.2: **Bar plots showing the frequency distribution of samples in the data set.** (A) The top categories from 116 body tissue sources are shown. (B) Data were available from 70 middle and high-income countries from which the top five with their gender distributions are plotted here. (C) Sorted frequency distribution for facilities is shown here. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

2.3.2 Data Imputation and Discretisation

Missing microorganism phenotypes and susceptibility were replaced with 'Not_Given' and "Not_Tested" respectively. Residual missingness in the merged AMR, WGI, GBD, and Finance datasets was addressed through imputation using a state-of-the-art Random Forest approach (Stekhoven and Bühlmann (2012)). Continuous data were then discretized into three disjoint intervals i.e. Low, Medium, and High using a KNN-based algorithm(Kim and Kim (2015)).

Priority category of Pathogens	Pathogens	Samples in High Income Country	Samples in Middle Income Country
Critical	Pseudomonas aeruginosa	11373	13266
Critical	Klebsiella pneumoniae	11901	15437
Critical	Escherichia coli	14864	18125
Critical	Acinetobacter baumannii	4468	5816
Critical	Enterobacter cloacae	7462	6860
High	Staphylococcus aureus	21456	27314

Table 2.2: Distribution of organisms in the random sample. Random sampling was conducted to avoid bias introduced in downstream models by the higher number of samples from high-income countries. A nearly equal number of samples were observed for each pathogen listed by the WHO global priority list of pathogens. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

2.3.3 WHO Declared Highly Important Pathogens Selection:

WHO declared critically important and highly important pathogens were selected [6]. These include Pseudomonas aeruginosa, Klebsiella pneumoniae, Escherichia coli, Acinetobacter baumannii, Enterobacter cloacae, Staphylococcus aureus. These were then segregated into those isolated from high, middle, and low-income countries respectively **Table 2.2**.

2.3.4 Analysis

A generative machine learning approach and what-if analysis were carried out to capture global patterns. These were followed by a discriminative model for the isolate-level prediction model

Bayesian Network Analysis

Bayesian networks learn the latent structure in complex data, represent it as compact graphical representations, and allow for inferential and intervention modelling. Interpretability and explainability are the key challenges in artificial intelligence–based decision models. Causal networks (Pearl (1998)) are one of the most promis- ing machine

[6]https://www.who.int/news/item/27-02-2017-who-publishes-list-of-bacteria-for-which-new-antibiotics-are-urgently-needed

learning approaches that combine expressive power with explainability. We learned Bayesian networks directly from complex multivariate data for explainable intervention modelling . In order to avoid overfitting, we took a two-level approach. The first-level Bayesian network was constructed on principal components of WGI, GBD, Finance, and AMR data sets. Specifically, a principal component analysis (PCA) was carried out for numerical data sets and multiple correspondence analysis (MCA) for cate- gorical data sets. The top two components were used for learning the associations via a Bayesian network. The second level used the variables from Markov blankets of all antibiotics calculated. Network inference was carried out using a junction tree algorithm for exact inference implemented in wiseR (Maheshwari *et al.* (2021)). To identify the global determinants of AMR, the WHO declared critically important pathogens were selected for the construction of Bayesian networks for each pathogen. This was followed by a counterfactual what-if analysis.

Counterfactual Analysis

Counterfactual analysis (Ramachandran (1997)) was performed using the interventions identified from the BN analysis to answer the hypothetical 'what if' questions. Counterfactual queries were formed such as 'What would be the resistance of antibiotics if a country with poor access to immunisation and obstetric care gained an advanced access to immunisation and obstetric care, provided all other covariates remained the same'. These scenarios were carried out using the Counterfactual R package(Chen *et al.* (2017)).

Prediction of Antibiotic Susceptibility using Supervised Machine Learning

The Markov blanket of pathogen-specific Bayesian network included variables from GBD, WGI, and finance indicators. These along with demographic and clinical information of isolates were used to build a model for the prediction of antibiotic sus-

ceptibility. Data were partitioned into training (80%) and testing (20%) sets, and the class imbalance was corrected using the Synthetic Minority Oversampling Technique. Supervised machine learning. Models were constructed using random forest (RF), support vector machine, and logistic regression, and Naïve Bayes approach were used for predicting the susceptibility of antibiotics using the Scikitlearn library (Pedregosa *et al.* (2011)) in Python. Feature importance was calculated using a permutation-based feature importance method.

2.4 Results

2.4.1 Global Prevalence of Multiple Antibiotic Resistance of Critically Important Pathogens

To understand the spread pattern of Antibiotic resistance of different pathogens, we calculated the Multiple Antimicrobial Resistance (MAR) score country-wise. This revealed that every pathogens had a different pattern of resistance spread. For Acinetobacter baumannii, MAR score was found in the range [0, 0.86] with the highest value in Vietnam (0.86). For E coli, the MAR score was found in the range [0.032, 0.45] with the highest values in Indonesia (0.45), India (0.40) and the lowest value in Norway (0.032). For Klebsiella pneumoniae, the MAR score was found to be in the range of [0.09, 0.64] with the highest value in Serbia (0.64) and the lowest value in Japan (0.09). For Pseudomonas aeruginosa, the MAR score was observed to be in the range [0.016, 0.75] with the highest value in El Salvador (0.75) and the lowest value in Norway (0.016).

2.4.2 Component Level Bayesian Network Revealed Strong Connection of Antibiotics with Independent WGI, Finance, GBD datasets

On a one-time network **Figure 2.3** we found a strong association of governance, finance, and GBD principal components with Antibiotics components.

2.4.3 Probabilistic Impact Quantification of Governance, Finance, Socioeconomic and Burden of Disease on Antibiotic Resistance

For WHO declared critical pathogens we learned a robust bootstrapped network. From the Network Inference governance of the country and Access to Immunisation and Obstetric care, we found a strong influence of AMR in all pathogens **Table 2.3**. We found the probability of cefepime resistance in E coli pathogens to be 21% higher in a country where voice and accountability (Refer Supplementary Definitions) was low as compared to a country where voice and accountability were high. The probability of piperacillin resistance in Acinetobacter baumannii pathogens is 50% lower in a country where there is a high level of government effectiveness. The probability of cefepime resistance is 12.98% lower in the pathogens Pseudomonas aeruginosa and 13% lower in the pathogens Staphylococcus aureus in a country where there is a high level of government effectiveness as compared to a country having a low level of government effectiveness. Similarly, the probability of amikacin resistance was 12.17% lower in the pathogens Pseudomonas aeruginosa and 11.6% lower in the pathogens Staphylococcus aureus in a country where there was a high level of government effectiveness as compared to a country having a low level of government effectiveness. We also found that good Access to Immunisation and Obstetric care in a country significantly decreases

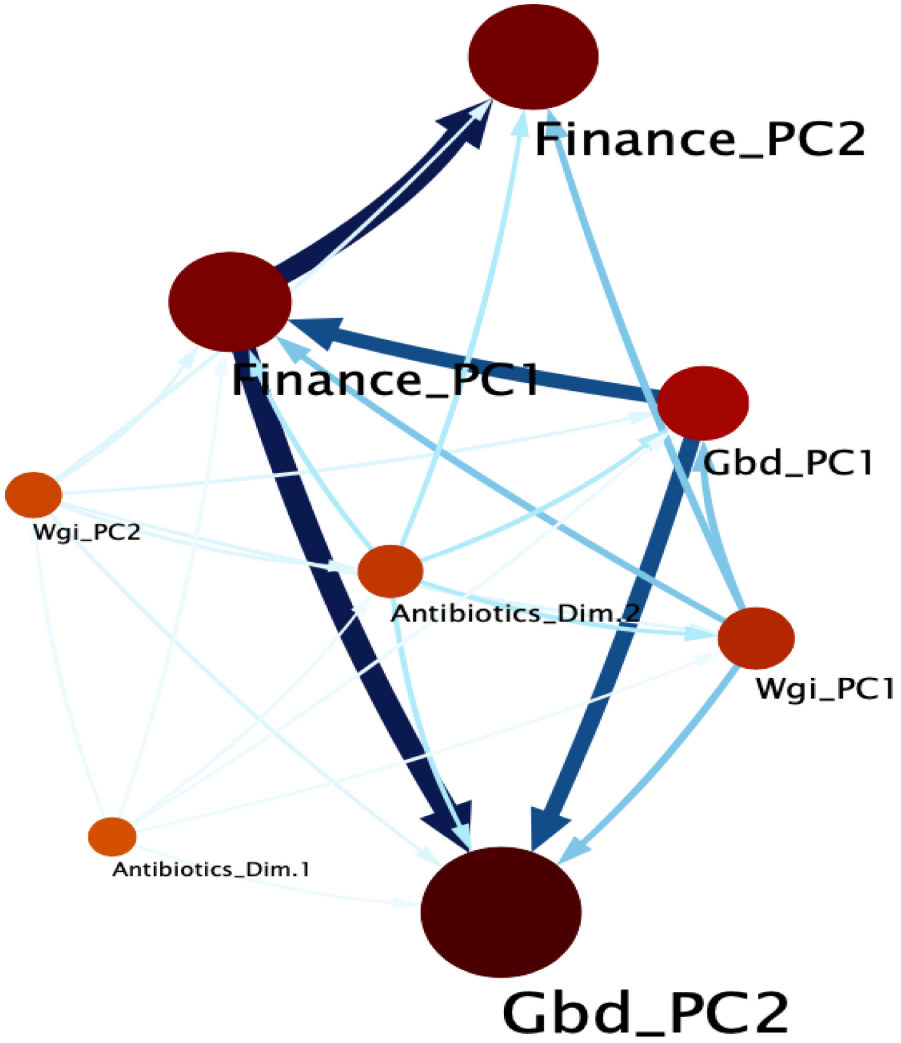


Figure 2.3: **Principal component level association among AMR, WGI, GBD, and Finance.** A Bayesian network of principal components was learned from data to check the component-level association of AMR with Finance, GBD, and WGI. The first component of antibiotics (Antibiotics_Dim.1) was associated with the second component of GBD (Gbd_PC2) and WGI (Wgi_PC2), and the second component of antibiotics (Antibiotics_Dim.2) was associated with the first component of Finance (Finance_PC1). This indicates a strong association of AMR with WGI, Finance, and GBD. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

ceftriaxone and meropenem resistance in many pathogens of bacteria. The probability of ceftriaxone resistance in countries with good Access to Immunisation and Obstetric care was observed to be 14.74% lower in E Coli, 36% lower in Acinetobacter baumannii, 12.47% lower in Enterobacter Cloacae, 15% lower in Klebsiella pneumoniae while the probability of meropenem resistance in Klebsiella pneumoniae was observed to be

3.43% higher in countries with good Access to Immunisation and Obstetric care. Improvements in water sanitation and hygiene (WASH) are critical elements of preventing infections and reducing the spread of antimicrobial resistance (AMR) as identified in the Global Action Plan to combat AMR. In our analysis, we found out that the probability of cefepime resistance in Klebsiella pneumonia was 19.87% higher in countries where the level of Unsafe Wash Sanitation was high. Our analysis also quantified the effect of fruit consumption on cefepime resistance in Klebsiella pneumoniae. High fruit consumption decreased the chance of cefepime resistance by 13.04%

2.4.4 Counterfactual Analysis Quantified Association between Ceftriaxone Resistance and Access to Immunisation and Obstetric care:

From Bayesian network analysis, it was observed that the Access to Immunisation and Obstetric care has impacted the Ceftriaxone resistance with a high magnitude consistently for all pathogens. So, we performed a counterfactual analysis on the ceftriaxone resistance in the Middle and High-income countries. For the pathogen Enterobacter Cloacae, the counterfactual effect of Access to Immunisation and Obstetric care (what would be the ceftriaxone resistance if a poor Access to Immunisation and Obstetric care country gained an advanced Access to Immunisation and Obstetric care, provided all other covariates remained the same) on ceftriaxone resistance was observed **Figure 2.4**. It was consistent in the Middle and High-income countries with median quantile effect in Middle-income countries being 0.007 (95% CI [-0.05, 0.0.7]), and in High-income countries being 0.02 (95% CI [0.003, 0.04]). For Acinetobacter baumannii and E Coli, we found the counterfactual effect of Access to Immunisation and Obstetric care on ceftriaxone in Middle-income countries to be much higher than that in high-income countries. The median quantile effect in the Middle-income countries was observed to be -0.9 (95% CI [-0.17,-0.02]) in the case of Acinetobacter baumannii and -0.03 (95%CI

pathogens	Variables (Parent)	Antibiotics (Child)	Pr(Antibiotics=Resistant|Variable=Low)	Pr(Antibiotics=Resistant|Variable=High)
E Coli	Mean_Temp_Long_Term	Cefepime	0.15	0.32
	Voice and Accountability	Cefepime	0.33	0.12
	Fruits_G_Adj	Cefepime	0.2	0.2
	Access to Immunisation and Obstetric care Capped	Ceftriaxone	0.3302	0.1828
	Access to Immunisation and Obstetric care Capped	Ceftazidime	0.1357	0.189
	Dengue_Outbreak	Ceftazidime	0.0978	0.1652
	Access to Immunisation and Obstetric care Capped	Meropenem	0.0139	0.006
Acinetobacter baumannii	Mean_Temp_Long_Term	Piperacillin Tazobactam	0.5	0.7489
	Government_Effectiveness	Piperacillin Tazobactam	0.79	0.29
	Hospital_Bed_per1000	Meropenem	0.66	0.47
	Access to Immunisation and Obstetric care Capped	Meropenem	0.38	0.67
	Access to Immunisation and Obstetric care Capped	Ceftriaxone	0.7	0.34
EnteroBacter Cloacae	Rota Coverage Prop	Meropenem	0.0238	0.0319
	Access to Immunisation and Obstetric care Capped	Meropenem	0.05	0.0262
	Sev Scaler Diarrhea	Cefepime	0.09	0.2218
	Access to Immunisation and Obstetric care Capped	Ceftriaxone	0.4107	0.286
	Access to Immunisation and Obstetric care Capped	Minocycline	0.1255	0.0508
Klebsiella pneumonia	Access to Immunisation and Obstetric care Capped	Ceftriaxone	0.39	0.24
	GFDD_EI_04	Ceftriaxone	0.2856	0.1
	Fruits_G_ADJ	Cefepime	0.4111	0.2807
	Mean_Temp_Long_term	Cefepime	0.2865	0.3339
	Unsafe_Wash_Sanitation	Cefepime	0.2134	0.4119
	Access to Immunisation and Obstetric care Capped	Ceftazidime	0.1984	0.3516
	Dengue_Outbreak	Ceftazidime	0.2186	0.25
	pop_Dens_under_150_Psqkm_Pct	Ceftazidime	0.226	0.3005
	Fruits_G_ADJ	Piperacillin Tazobactam	0.229	0.2195
.	Access to Immunisation and Obstetric care Capped	Piperacillin Tazobactam	0.1996	0.2046
	Access to Immunisation and Obstetric care Capped	Amoxicillin Clavulanate	0.2889	0.2697
	Rule of Law	Amoxicillin Clavulanate	0.3396	0.119
.	pop_Dens_under_150_Psqkm_Pct	Amoxicillin Clavulanate	0.2205	0.3291
	Access to Immunisation and Obstetric care Capped	Meropenem	0.0587	0.093
Pseudomonas aeruginosa	Access to Immunisation and Obstetric care Capped	Meropenem	0.1826	0.2413
	Hospital Beds per 1000	Meropenem	0.2865	0.2381
	Access to Immunisation and Obstetric care Capped	Cefepime	0.2151	0.1353
	Government Effectiveness	Cefepime	0.2	0.0702
	Government Effectiveness	Amikacin	0.1417	0.02
Staphylococcus aureus	Government Effectiveness	Cefepime	0.2	0.07
	Government Effectiveness	Amikacin	0.1417	0.0257
	Access to Immunisation and Obstetric care Capped	Cefepime	0.2151	0.1353
	Access to Immunisation and Obstetric care Capped	Meropenem	0.1826	0.2413
	Hospital Beds Per 1000	Meropenem	0.2865	0.2381

Table 2.3: Inferences from pathogen-wise Bayesian network. We calculated the difference in probability of antibiotic resistance by changing levels of the variable, which was observed as a parent node of antibiotics in the Bayesian network. A significant effect of Government effectiveness, Access to Immunisation and Obstetric care, and Unsafe Wash sanitation on Antibiotic resistance were observed in all pathogens. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

[-0.16, 0.10]) in the case of E Coli. The median quantile effect in the High-income countries was observed to be 0.05 (95% CI [-0.03,013]) in the case of Acinetobacter baumannii and -0.004 (95% CI [-0.03, 0.02]) in the case of E Coli. Although we performed the same counterfactual analysis for meropenem resistance, which was found to have a strong association with Access to Immunisation and Obstetric care, we did not find any monotonic trend in this case.

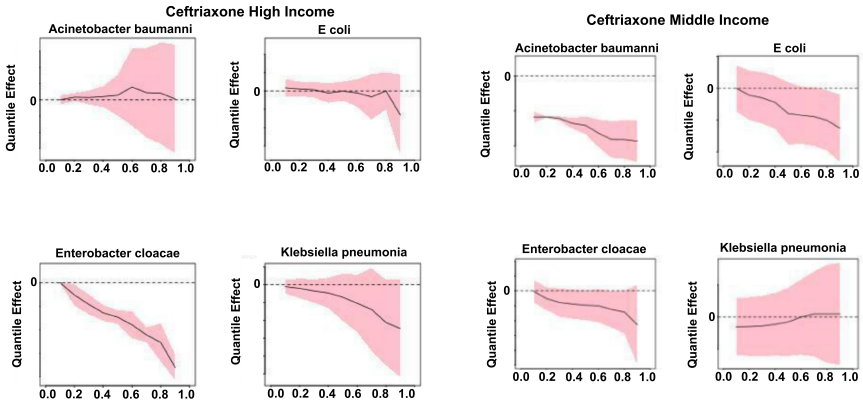


Figure 2.4: **Counterfactual analysis.** Bayesian network analysis showed that access to immunisation and obstetric care influenced the development of resistance in antibiotics, with the highest effect on ceftriaxone. To further understand this relationship, a counterfactual analysis, showed the change in ceftriaxone resistance in a situation where countries with lower care would counterfactually be switched to high care. For example, we observed that ceftriaxone resistance in Acinetobacter baumannii, in middle-income countries, could be reduced by 90%, if they had high access to immunisation and obstetric care. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

2.4.5 Prediction of Antibiotic Susceptibility

Overall, the Random Forest Model was the best performing model. Amongst all the pathogens, interestingly our prediction models have performed best for Staphylococcus aureus **Figure 2.5** which is known to be the most dangerous amongst the common staphylococcal bacteria that often cause skin infections. For Staphylococcus aureus, our model gave the best prediction for ceftaroline and oxacillin with the highest AU-

ROC, 0.94 and 0.89 respectively (from the RF model). We also found that our model performed well for cefepime and ceftazidime susceptibility prediction for Klebsiella pneumoniae with AUROC 0.88, 0.92 respectively (from the RF model). From our prediction models, we also found that meropenem susceptibility is highly predictable for Escherichia coli (AUROC=0.93) compared to other pathogens like Pseudomonas aeruginosa (AUROC=0.75), Klebsiella pneumoniae (AUROC=0.81), Acinetobacter baumannii (AUROC=0.75). On further analysis of the feature importance of prediction model (RF) for meropenem susceptibility, we found that Echerichia coli and, interestingly, Access to Immunisation and Obstetric care of a country to be the most important feature. Although our model has not shown high performance for pathogens like Pseudomonas aeruginosa, overall our models have shown decent and actionable performance.

2.4.6 Antibiotic Steward: An Interactive Web App to Discover AI-Driven Decision for AMR

In order to make our findings accessible, a web-based application **Figure 2.6** for exploring Bayesian networks and prediction models for the pathogens E. coli, Acinetobacter baumannii, EnteroBacter Cloacae, Klebsiella pneumonia, Pseudomonas aeruginosa, Staphylococcus aureus was developed. The visualization includes the Bayesian network structure, AUROC, and feature significance scores of prediction models. By selecting a pathogen and an antibiotic of interest, stakeholders (policymakers, physicians, etc.) may identify the most significant indicators which can be used to help with AMR policies. This web app focuses on using AI models that can be explained as a way to improve antibiotic stewardship.

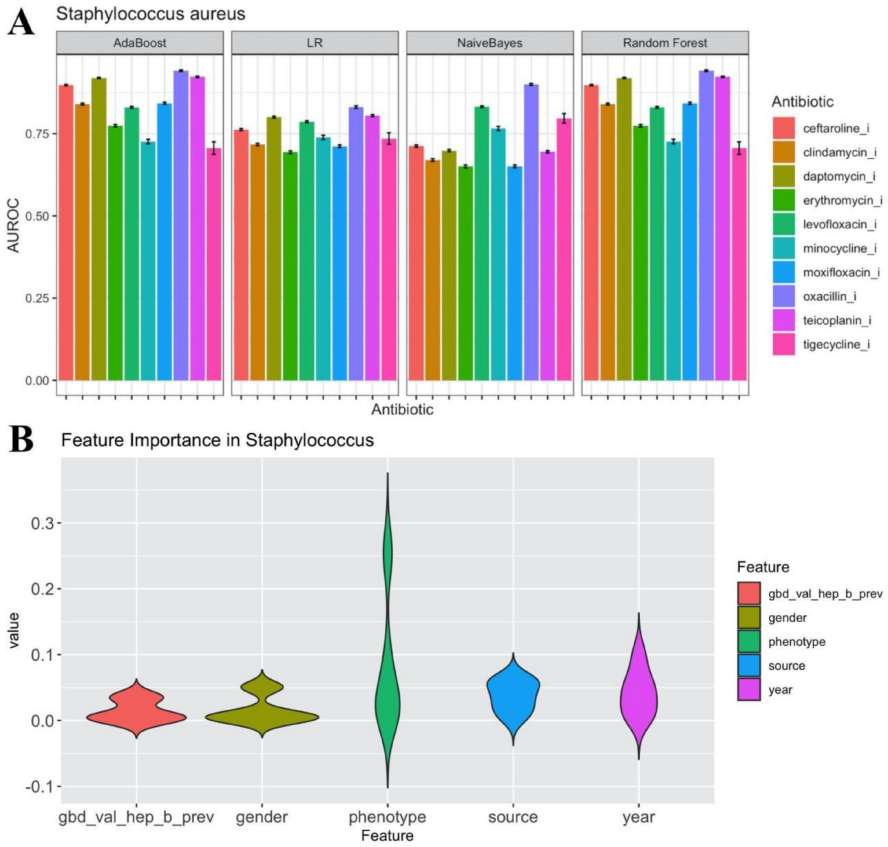


Figure 2.5: **Antibiotic susceptibility prediction.** (A) AUROC of prediction models for Staphylococcus aureus. Error bar in the figure denotes the 95% CI of AUROC. (B) Feature importance was calculated using permutation-based methods, violin plot of the top-performing feature among all the given antibiotics is shown where "gbd val hep b prev(Hepatitis B prevalence)" was identified as a key feature. *Credits: Awasthi, Raghav, et al. "Estimating the impact of health systems factors on antimicrobial resistance in priority pathogens." Journal of Global Antimicrobial Resistance 30 (2022): 133-142.*

2.4.7 Discussion and Conclusion

Antibiotic resistance is rising to alarmingly high levels in all parts of the world. New resistance mechanisms are emerging and spreading globally, threatening our ability to treat common infectious diseases. The COVID-19 pandemic has shown us the havoc that a hitherto unknown pathogen can cause when we don't have antimicrobials that work effectively. The pandemic has also led to rampant use of antimicrobials, even in individuals who did not require them. The virus has also increased hospital admissions, which increases the risk of acquiring and spreading multi-resistant organisms



Figure 2.6: **Antibiotic steward Web Application.** Interactive web-app to visualise the Bayesian network and prediction model results for the pathogens E Coli, Acinetobacter baumannii, EnteroBacter Cloacae, Klebsiella pneumonia, Pseudomonas aeruginosa, Staphylococcus aureus. Bayesian network structure, area under the receiver operating characteristic curve (AUROC), and feature significance scores of prediction models are all shown graphically in the web-app. *Credits:https://antibioticsteward.tavlab.iiitd.edu.in/*

from hospitals, which in turn prompts the use of more antimicrobials. A growing list of infections such as pneumonia, tuberculosis, blood poisoning, gonorrhea, and foodborne diseases are becoming harder and sometimes impossible to treat as antibiotics become less effective. The emergence and spread of the resistance are worse in those regions where antimicrobial medicines can be bought for human or animal use without a prescription. Similarly, in countries without standard treatment guidelines, antibiotics are often over-prescribed by health workers and veterinarians and overused and misused by the public. Certain economic and governance factors such as Political Stability and Absence of Violence/Terrorism, Control of Corruption, Voice and Accountability, Government Effectiveness, Regulatory Quality, and Rule of Law play a key role in designing such frameworks which can help curtail the spread of antimicrobial resistance. By consolidating four different datasets (the AMR, GBD, WGI, and Finance data sets) into one, this study aims to find the actionable global determinants. Taking into consideration the prevalence of MAR score across the globe, our findings report that MAR

score is non-uniform across the globe for critically important pathogens. This finding calls upon devising policies that are effective at each country level. Using explainable and expressive machine learning techniques like Bayesian Network Analysis, we found a strong connection of AMR with the WGI, Finance, and GBD datasets which motivated us to proceed with Counterfactual analysis of Ceftriaxone and meropenem. As per the results, the counterfactual effect of Access to Immunisation and Obstetric care on Ceftriaxone in Middle-income countries turned out to be much higher than that in High-income countries. Our analysis showed that most of the antibiotics were more likely to encounter resistant bacteria in countries with lower Access to Immunisation and Obstetric care capped', a measure of Access to Immunisation and Obstetric care. In GBD datasets, this was estimated using a principal component analysis of antenatal clinics, DTP3 immunization, measles immunization, in-facility delivery, and skilled birth attendance. The analysis revealed interesting relationships between antibiotics such as meropenem, ceftazidime, and Access to Immunisation and Obstetric care. The resistance to ceftazidime in E.coli and Klebsiella, and the resistance to meropenem in Acinetobacter, staphylococcus, Klebsiella, Pseudomonas, was more likely in countries with higher access to health care. These may be explained by the fact that these antibiotics are typically administered in ICU settings, hence may have lower administration and resistance in countries with lower access. Further, a previous study showed that the ratio of resistant and drug- sensitive infections significantly correlated with hospital size. In countries with lower access to health care, the proportion of small hospitals is lesser than in countries with higher access to healthcare. This might also explain the rise in antimicrobial resistance in countries with low Access to Immunisation and Obstetric care. Low-income countries often face the problem of over-prescribing of antibiotics, which in turn is believed to contribute to the emergence of antimicrobial resistance. An interesting counterpoint to this hypothesis is the near absence of a correlation between WASH criteria (which have a positive correlation to income of the country) and antimicrobial resistance prevalence. In a country with poor sanitation and a lack

of standard treatment guidelines, we would expect a widespread community spread of antimicrobial-resistant organisms and a positive correlation between WASH conditions and the prevalence of AMR. On a similar note is the positive correlation between coverage of the rotavirus vaccine and the probability of resistance to meropenem.

CHAPTER 3

Discovering Potential Interventions: Learning Explainable Interventions to Mitigate HIV Transmission in Sex Workers Across Five States in India

3.1 Preamble

Parts of this chapter have been published in:

"*Awasthi, R., Patel, P., Joshi, V., Karkal, S., Sethi, T. (2020). Learning explainable interventions to mitigate hiv transmission in sex workers across five states in india. 34th ML4D Workshop at Conference on Neural Information Processing Systems*".

3.2 Introduction

Female sex workers (FSWs) are one of the most vulnerable sections of society and their challenges are multiplied manifold in developing countries. Cultural factors leading to power asymmetry, economic vulnerability, social stigma and lack of access to good healthcare increases their vulnerability for practicing unsafe sex and lack of knowledge of their HIV status. The interventions that could mitigate the spread of infection are well known, e.g., condom usage and testing (Das and Horton (2015)). However, the factors that govern these protective behaviors are complex, multi-dimensional and under-studied. With a high burden of HIV and limited resources, developing countries could benefit the most from data-driven evidence and machine learning approaches that

can prioritize interventions (Wilder *et al.* (2021)). In India, community-led organizations (COs) working towards HIV prevention are increasingly using digital tools and collecting massive data, thus opening the possibility of understanding these complex influences. The Avahan program[1] has been working in a phased manner in five Indian states since 2003, with a strong emphasis on improving condom usage behaviors. Work in Avahan III showed the importance of engagement with the COs for improving the self-confidence and self-efficacy to procure condoms (Mahapatra *et al.* (2020)). In this study, we decided to further delineate the actionable factors through learning the representation of this complex dataset using structure learning approaches, followed by predictive modeling. The goal of the study was to prioritize the set of actions that would have the strongest influence on higher condom use behavior. We believe that this is the first study combining a unique dataset, expertise in community engagement and a principled machine-learning approach for transparent discovery and prioritization of interventions for mitigating the spread of HIV in female sex workers.

3.3 Methods

3.3.1 Dataset

Survey based data on FSW was collected by professionals social workers under the Avahan III program (https://allianceindia.org/portfolio/avahan/) in Andhra Pradesh, Karnataka, Maharashtra, Tamil Nadu and Telangana states of India. A pretested, structured questionnaire Outcome Monitoring Survey tool was used to collect information directly on a digital Tablet. In addition to the basic demographic and socio-economic characteristics of the respondents, the detailed questionnaire also captured a variety of other factors including CO membership, sexual behavior, condom use with clients and intimate partners, history of STI (Sexually transmitted infections) during last six months,

[1] https://allianceindia.org/portfolio/avahan/

HIV testing behavior, coverage with gender based violence, social protection (civic identities and social schemes), financial security services and products (Savings Account, Savings Products, insurance and other financial schemes), reproductive health status as well as financial and food crisis or insufficiency experienced by respondents. A total of 222 variables from 11016 female sex workers were recorded and considered in this analysis during the period July – September, 2017 (3 years after the programme was initiated).

3.3.2 Structure Learning and Network Inference

Explainability and interpretability are key challenges in *AI* based decision models, especially in the public health settings. The expressive power of Bayesian networks combined with the advances in learning structure directly make them well suited to address complex problems that need transparency. Although the structure learned from observational data and edge-directions **do not** indicate *causality*, yet these can be used as a *probabilistic reasoning system* (Pearl (1998)). . An ensemble averaged directed acyclic graph (DAG) structure was learned directly from the data **Equation 1.3**.

The structure was learned as a search problem over a constrained space of DAGs and evaluated with the Akaike Information Criterion (*AIC*) as a metric. Since the DAG representation for a set of conditional independencies is not unique, and only up to Markov Equivalence (Pearl (2009)) an ensemble of such structures learned from bootstrapped samples of data was constructed using a majority voting criterion (Koller and Friedman (2009)). For conducting inferences and queries, the DAG was parametrized with the posterior probability distributions for all variables.$x_{v_j} \in X$ conditionalized upon an evidence $e = \{e_1, e_2, ..e_m\}$ from a set of variables $X(e)$. Where, the prior marginal distribution $P(x_{v_j})$ is given by the following set of equations

The likelihood was then defined as follows.

$$\begin{aligned} L(x_{v_j}|e) &= P(e|x_{v_j}) \\ &= \sum_{x \in X \setminus \{x_{v_j}\}} \prod_{i \neq j} P(x_{v_i}|x_{pa(v_i)}) \prod_{x \in X(e))} E_x \end{aligned} \tag{3.1}$$

For each $x_{v_j} \notin X(e)$, where E_x is the evidence function for $x \in X(e)$ and v_i is the node representing x_{v_i}. Finally the posterior was inferred through application of the Bayes rule and is always proportional to the estimated Likelihood

$$P(x_{v_j}|e) \propto L(x_{v_j}|e)P(x_{v_j}) \tag{3.2}$$

where the proportionality constant $P(e)$ can be computed from $P(X, e)$ by summation over X and $P(x_{v_j})$ may be obtained by inference over empty set of evidence.

3.3.3 Discriminative Predictive Model for Condom Usage

Next, a discriminative approach to confirm whether the discovered influences from the generative approach had enough predictive power to be useful in the real world settings was taken. Data were partitioned into training (80%) and testing (20%) sets with the class imbalance was corrected using the Synthetic Minority Oversampling Technique(SMOTE). Condom usage was defined as the dependent variable and the top predictors were identified. Different supervised machine learning models like RandomForest (RF), Support Vector Machine (SVMs), Logistic Regression, XgBoost and AdaBoost were then used to predict condom usage.

3.4 Results

3.4.1 Intervention Score and Vulnerability

A Combined Vulnerability Reduction Intervention Coverage Score (CVRICS) was defined and calculated for all female sex workers in our database **Figure 3.1**. CVRICS is a combination of variables like social protection intervention score, financial security intervention score and number of incidents of crisis faced by the FSW within the last 6 months and its values range between 0 to 30. These include access to entitlements and social protection schemes, access to formal financial products and incidents of violence and reporting of the same as experienced in the last 6 months before the survey. It was observed that although the observed maximum CVRICS score was 21, more than 6000 female sex workers had a score of 7-11 indicating that a majority of FSWs were vulnerable either financially, socially or physically.

3.4.2 Inferences from Bayesian Network

From our 101 bootstrapped ensemble-averaged DAG we found that financial empowerment was a key association for condom use. The learned conditional probabilities showed that having at least one financial investment such as RD (Recurrent Deposit), FD (Fixed Deposit) and Chit Fund increased the probability of using condoms with clients by as much as 6%. We also observed Depression symptoms to be key influencers on a female sex worker's confidence to buy a condom on her own **Figure 3.2**. Our model showed that the probability of a FSW buying a condom on her own decreases by 14% if she reported symptoms of depression as measured using CESD short scale. We also identified the determinants for HIV testing and our DAG revealed legal education and financial stability as two key influencers for HIV testing. The conditional probability of a FSW getting tested in every 6 months, increased by 16% if she had been

A

No. of incidents of crisis faced by the KPs in last six months	Reported of last incident of crisis to the CO / Authority	SSJ intervention coverage Score
> 6	No	1
> 6	Yes	2
1 - 6	No	3
1 - 6	Yes	4
0	NA	5

B

No. of Civic Identities, the KPs possess	No. of Social Protection Schemes, the KPs have received benefits of	SP intervention coverage Score
≤ 1	0	1
> 1	0	2
≥ 1	1	3
≥ 1	2	4
≥ 1	> 2	5

C

Savings Account in Bank or Post Office	Saving products (RD, FD, PPF, NSC, SHG, Co- op, Shares, Mutual Fund) Insurance (Life, Health, Accident) investments (Gold, Land Plot, Business)	FS intervention coverage Score
No	0	1
Yes	0	2
No / Yes	≥ 1 product of any one category	3
No / Yes	≥ 1 product of any two categories	4
No / Yes	≥ 1 product of all three categories	5

D **Definition:**
Combined Vulnerability Reduction Intervention Coverage Score = SSJICS + FSICS + SPICS

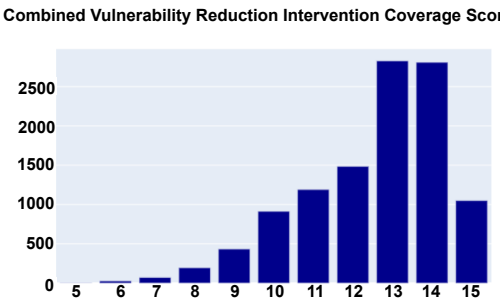


Figure 3.1: **Intervention coverage scores.** A)Safety, Security and Justice Intervention score(SSJICS) B)Social Protection Intervention Coverage Score (SPICS) C)Financial Security Intervention Coverage Score(FSICS) D)Combined Vulnerability Reduction Intervention Coverage Score(CVRICS)

trained on legal education. It was also observed that creation of a personalized financial plan increased the probability of HIV testing among the sex workers by 12%. Although it is intuitively expected that these factors could play a role, our model revealed these to be top influences and also quantified their marginal impact. Government schemes were also observed to strongly influence many key FSW related variables. Access to at least one government scheme increases a high level of self-efficacy by 13%, which in-turn drives condom usage behaviors. Since depression symptoms were influencing condom buying behavior, we observed an association between unwanted pregnancy and depression, although a cause and effect relationship cannot be established. How-

ever, this highlights the role of mental vulnerability and unsafe practices, and the need to design holistic interventions. Finally, although the DAG approach identifies potential confounding transparently and does not require confounders to be specified up-front, the structure was evaluated by field experts for possible biases.

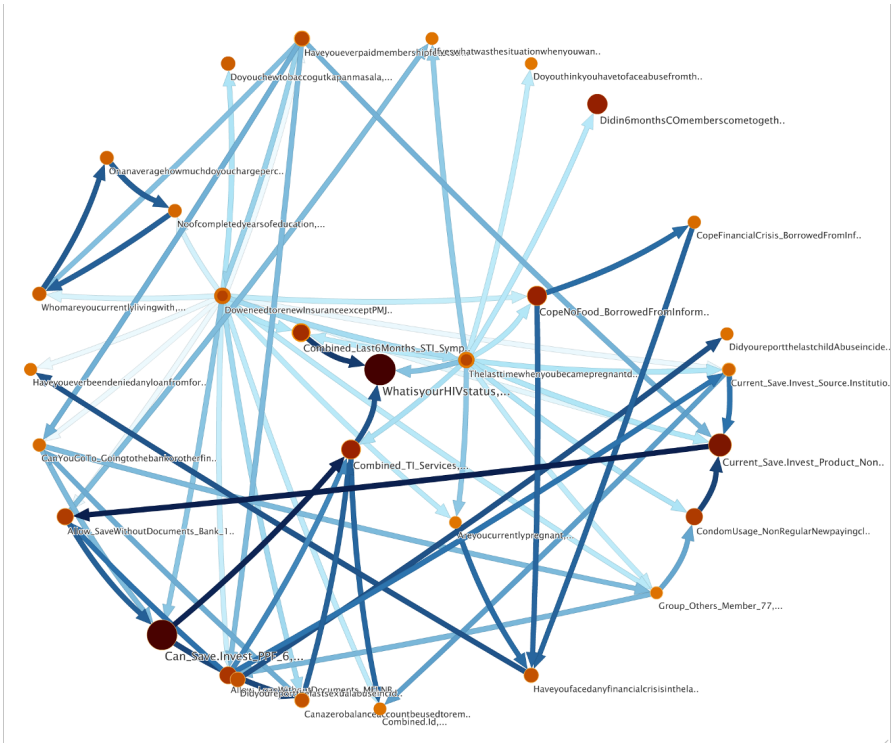


Figure 3.2: **Network Structure.** Learned 101 Bootstrapped, Ensembled Bayesian Network. *Credits: Awasthi, Raghav, et al. "Learning explainable interventions to mitigate hiv transmission in sex workers across five states in india." 34th ML4D Workshop at Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada (2020).*

Predictive model for Condom Usage

The marginal probability of FSWs using condoms with regular paying clients and non regular paying clients was observed as 88.6%, 80.07% respectively **Table 3.1**. Hence there was certainly a scope of improvement to push this beyond the 90 - 95% goal, but this also led to class imbalance, which was corrected using SMOTE algorithm. Discriminative models were trained on training set and their performance was evaluated on test set **Table2.1**. The best performing models were Random Forest and XgBoost

	Model	Accuracy	F1 Score	Sensitivity	Specificity	AU-ROC	AU-PRC
Condom usage with Non-Regular Paying Clients	Random Forest	0.96 [0.95,0.97]	0.98 [0.97,0.98]	0.76 [0.64, 0.86]	0.97 [0.96,0.98]	0.86 [0.81,0.91]	0.991 [0.987,0.995]
	SVM	0.94 [0.92,0.95]	0.967[0.961,0.975]	0.635[0.52,0.74]	0.95 [0.94,0.96]	0.79 [0.75,0.84]	0.987 [0.982,0.992]
	Logistic Regression	0.93 [0.92,0.96]	0.96 [0.95,0.97]	0.66 [0.54,0.76]	0.95 [0.94,0.96]	0.81 [0.75,0.85]	0.98 [0.98,0.99]
	AdaBoost	0.93 [0.92,0.95]	0.96 [0.95,0.97]	0.71 [0.61,0.81]	0.95 [0.93,0.96]	0.83 [0.75,0.89]	0.99 [0.98,0.99]
	XgBoost	0.96 [0.95,0.97]	0.979 [0.972,0.984]	0.687 [0.56,0.82]	0.97 [0.96,0.98]	0.83 [0.77,0.88]	0.99 [0.98,0.993]
Condom usage with Regular paying Clients	Random Forest	0.97[0.96,0.98]	0.988 [0.982,0.99]	0.82 [0.71,0.91]	0.98 [0.97,0.99]	0.90 [0.84,0.94]	0.993 [0.991,0.999]
	SVM	0.94 [0.93,0.96]	0.972 [0.96,0.979]	0.77 [0.61,0.84]	0.96 [0.94,0.97]	0.86 [0.82,0.91]	0.991 [0.98,0.994]
	Logistic Regression	0.94 [0.92,0.95]	0.96 [0.95,0.97]	0.77 [0.67,0.87]	0.95 [0.93,0.96]	0.86 [0.80,0.91]	0.99 [0.98,0.99]
	AdaBoost	0.94 [0.92,0.95]	0.96 [0.95,0.97]	0.85 [0.74,0.92]	0.94 [0.93,0.95]	0.89 [0.85,0.93]	0.993 [0.99,0.996]
	XgBoost	0.97 [0.96,0.98]	0.985 [0.98,0.99]	0.87 [0.78,0.95]	097 [0.96,0.98]	0.89 [0.89,0.96]	0.995 [0.992,0.997]

Table 3.1: Prediction Model performance with 95% Confidence Interval. *Credits: Awasthi, Raghav, et al. "Learning explainable interventions to mitigate hiv transmission in sex workers across five states in india." 34th ML4D Workshop at Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada (2020).*

for non-regular paying clients and regular paying clients respectively with accuracy 0.96 in both the cases. The sensitivity of the model was only 0.71 and 0.69 in the best performing models, but the intent of interventions is not usually to screen for positive condom usage class, but the opposite. The specificity of detection of positive class was upwards of 97% in Random Forest and XgBoost models. Finally, in both the generative DAG and discriminative modeling approaches, the presence of financial variables as either top influencers or predictors indicate the potential impact of intervening with increased financial empowerment and literacy in order to instill confidence and self-efficacy in the FSWs. However, this study has several limitations. Firstly, as indicated earlier, it is nearly impossible to derive causal influences from observational data and our study does not claim to do that. However, the usefulness of DAGs as probabilistic reasoning systems is well known and little applied in public health settings. This is especially important in sensitive studies where black-box models can lead to serious ethical concerns. Secondly, since the study deals with the highly sensitive issue, the responses of FSWs may have some degree of incorrectness. However, we believe that the FSWs are not explicitly aware of the structure of the influences, hence it is difficult to confabulate the relationships that were discovered in a data-driven manner.

Conclusion

Sexual vulnerability is a highly sensitive issue, especially in the developing world, and female sex workers are the most vulnerable. A transparent and explainable approach with expert evaluation is critical to modeling sensitive issues. We demonstrate the potential of such an approach for learning to improve safety practices and mitigate the risk of HIV transmission in a vulnerable population in India; our work provides a framework for learning and replicating interventions in any setting across the globe. Having at least one financial investment, such as an RD (recurring deposit), FD (fixed deposit), or Chit Fund, increased the likelihood of using condoms with customers by as much as 6%, which was found as a crucial association for condom usage. Depression symptoms significantly influenced the FSW's confidence to purchase a condom independently. Two major determinants of HIV testing were found to be legal education and financial stability, with legal education raising the FSW's conditional likelihood of getting tested every six months by 16%. Our study offers insightful information that may be used to create comprehensive interventions that will enhance safe behaviors and reduce the risk of HIV transmission in vulnerable groups.

CHAPTER 4

Discovering Potential Interventions: Mental Health in COVID-19: The Impact of Medical, Financial, Social, and Behavioral Factors

4.1 Preamble

Parts of this chapter have been published in:

"*Jha, I. P., Awasthi, R., Kumar, A., Kumar, V., Sethi, T. (2021). Learning the mental health impact of COVID-19 in the United States with explainable artificial intelligence: observational study. JMIR mental health, 8(4), e25097*".

4.2 Introduction

People's mental health has been negatively affected by the COVID-19 pandemic because of the disruptions it has caused in their daily lives and communities. Though new cases and fatalities caused by the novel coronavirus may eventually decline, the increased need for mental health and substance use services is likely to persist long term, and this must be taken into account as policymakers discuss additional actions to alleviate the burdens of the COVID-19 pandemic. The Lancet Psychiatry(Pereira-Sanchez *et al.* (2020)) highlighted the needs of vulnerable groups during this time, including those with severe mental illness, learning difficulties, and neurodevelopmental disorders, as well as socially excluded groups such as prisoners, the homeless, and refugees.

Calls to action for engaging more early-career psychiatrists (Chen *et al.* (2020)), using technology such as telepsychiatry, and stressing the high susceptibility of frontline medical workers themselves (Holmes (2020)) have highlighted the magnitude of the problem. Further, interventions are expected to have a gender-specific impact, with women more likely to be exposed to additional stressors related to informal care, already existing economic disparity, and school closures. Similarly, age and comorbidity status may have a direct impact on susceptibility to mental health challenges due to their relationship with COVID-19 morbidity and mortality. Indeed, it has been established that emotional distress is ubiquitous in affected populations—a finding certain to be echoed in populations affected by the COVID-19 pandemic (Pfefferbaum and North (2020)). Finally, the role of social media (Seabrook *et al.* (2016)) is complex, with some research indicating an association between social media exposure and a higher prevalence of mental health problems (Gao *et al.* (2020)). This study is focused on learning a ranked list of factors that could indicate a predisposition to a mental disorder during the COVID-19 pandemic.

4.3 Methods

We extracted the data of 17,764 adults[1] from two weekly surveys (April 20-26 and May 4-10, 2020) of the US adult household population nationwide for 18 regional areas including 10 states (California, Colorado, Florida, Louisiana, Minnesota, Missouri, Montana, New York, Oregon, Texas) and 8 metropolitan statistical areas (Atlanta, Baltimore, Birmingham, Chicago, Cleveland, Columbus, Phoenix, Pittsburgh). Two rounds of data collection were available at the time of this analysis, and both rounds of data until May 25, 2020, were included in this analysis. To summarize, the data set comprised variables on physical health, mental health, insurance-related policy, economic security, and social dynamics. **Figure 4.2** shows the sociodemographic characteristics

[1] https://www.covid-impact.org/about-the-survey-questionnaire

of respondents participating in the survey. Next, we used data-driven Bayesian Network (BN) structure learning on the selected data because mental health variables may have complicated connections with possible confounding factors, mediation, and intercausal dependence. Using bootstrapping and ensemble averaging of edge directions, the structure of the learnt BN was made resilient. The hill climbing optimizer with the Akaike information criterion–based score was used to choose the optimal probabilistic graphical model that adequately described the data. Bootstrapped learning and majority voting were performed on 101 BNs. Using the belief propagation approach, exact inference was taught to measure the strength of learnt connections.

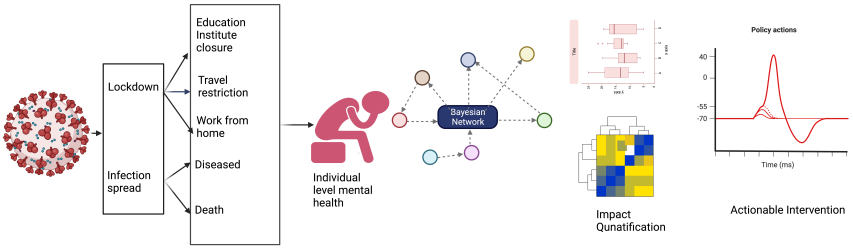


Figure 4.1: **Methodology illustration.** Using survey data on mental health, a bootstrapped Bayesian network was trained. The inferences were derived using a joint probability calculation. High-impact factors on mental health were labelled as Interventional variables.

4.4 Results

The data-driven BN structure learning exercise revealed interesting and actionable findings. From the learned structure, attribute soc5a (felt nervous, anxious, or on edge in the last 7 days) was found to be the parent variable for other mental health indicators in almost 100% of the bootstrapped networks, represented as the strength of the edges. Being a driver variable in the structure, attribute soc5a was taken as the primary dependent variable for downstream modeling analysis.

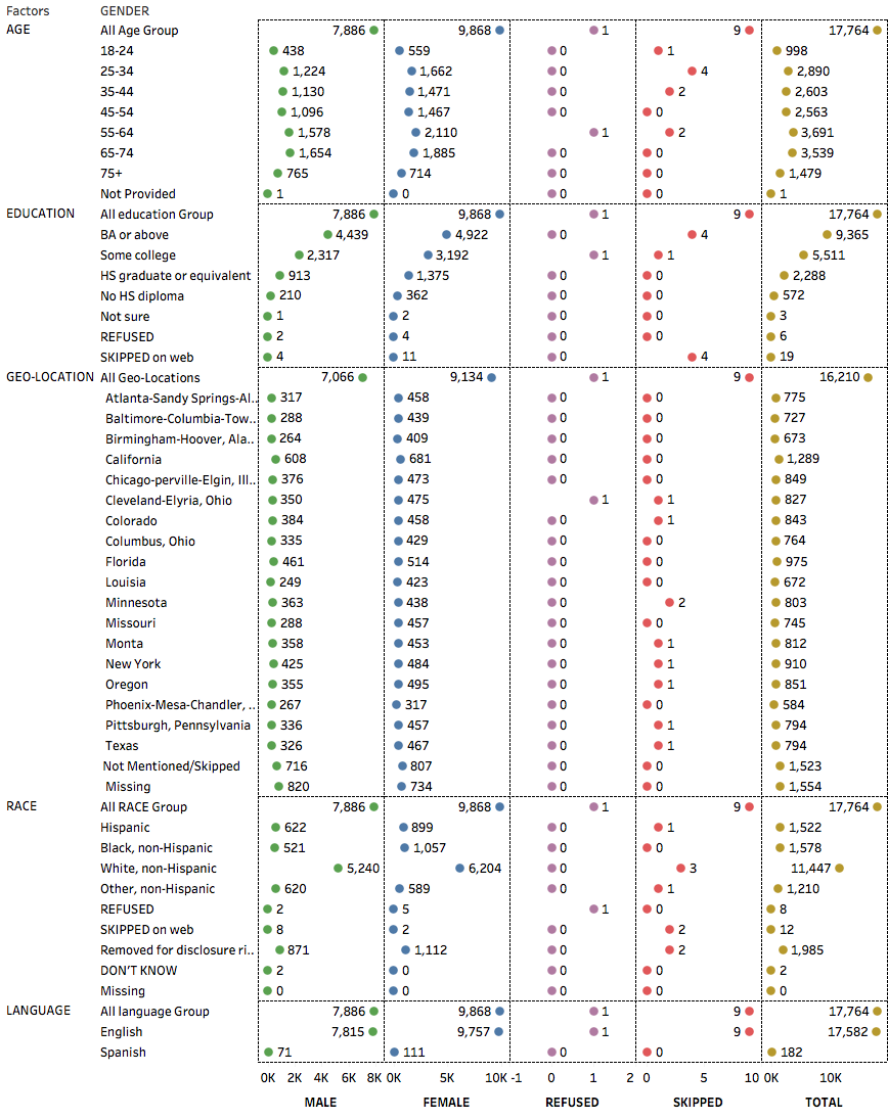


Figure 4.2: **Sociodemographics of respondents who participated in the survey.** It can be seen that there was almost a similar representation from both genders. Age groups 25-75 years were predominantly captured in the survey. Most of the respondents had received a Bachelor's degree or above and were nearly equally distributed across geographies within the United States. HS: high school. *Credits: Jha, I. P., Awasthi, R., Kumar, A., Kumar, V., Sethi, T. (2021). Learning the mental health impact of COVID-19 in the United States with explainable artificial intelligence: observational study. JMIR mental health, 8(4), e25097.*

4.4.1 Impact of Social Life and Work-Related Stressors

Our analysis using network inference via the exact inference algorithm showed a clear impact of in-person social communication on the reduction of anxiety levels. A strong (>5% with CI 1% on both sides) and (>6.5% with CI 1.5% on both sides) monotonic increase between control of anxiety and frequency of speaking with neighbors (attribute soc2a, attribute soc2b) were observed. This effect was weaker (1.5% with a wide confidence interval) with digital communication with friends and family conducted over phone, text, email, or other internet media (attribute soc3a, attribute soc3b). This finding underscores the importance of social communication while maintaining the appropriate measures such as masks and social distancing to maintain mental health during such isolating times. We also observed that the presence of kids in the house reduces the probability of depression by >11% with CI 2% on both sides. Furthermore, the exact inference upon the network revealed an increase in the conditional probability of anxiety (attribute soc5a) arising from canceled or postponed work (>4% with CI 1.4%), canceled or postponed school (7% with CI 1.5%), working from home (>5% with CI 1.3%), and studying from home (>7% with CI 1.8%). Interestingly, although 83% of all volunteers chose to wear the mask, 77% avoided restaurants, and 83% avoided public and crowded places, these measures were not found to be associated with a significant change in anxiety levels as inferred from our model. These inferences are summarized in **Figure 4.3**.

4.4.2 Impact of Symptoms and Comorbidities

We also investigated the relationship between mental stress and COVID-19 symptoms indicators. The World Health Organization recommends contacting health service providers if any COVID-19 symptoms (attributes phys1a to phys1q) are experienced within the last 7 days. Our network did not indicate any significant impact of these responses

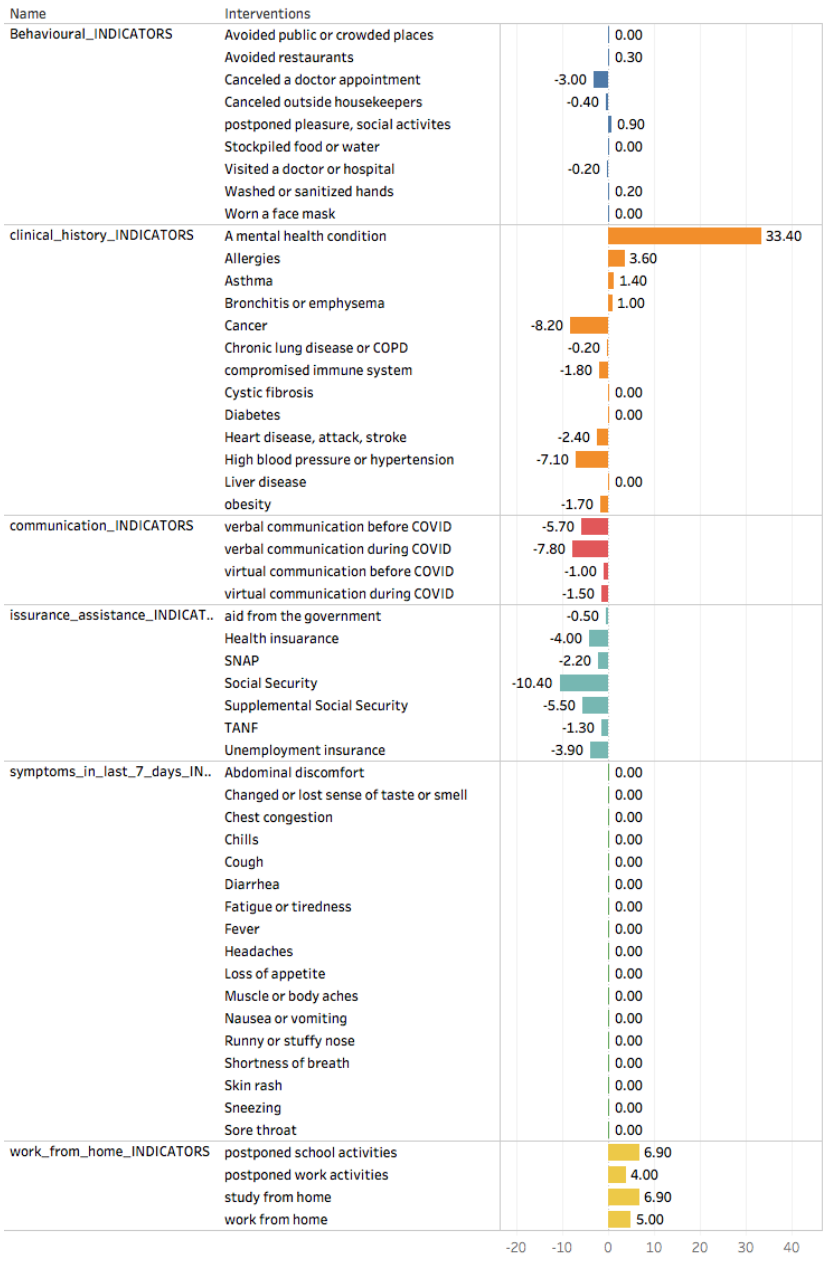


Figure 4.3: **Inferences from the Bayesian network.** The difference in inferred probability was calculated after conditioning the independent variables. A positive association implies a mental stress–inducing factor, whereas a negative association implies a mental stress reduction factor. The red circle shows the mean value, with green and blue showing confidence intervals. COPD: chronic obstructive pulmonary disease; SNAP: Supplemental Nutrition Assistance Program; TANF: Temporary Assistance for Needy Families. *Credits: Jha, I. P., Awasthi, R., Kumar, A., Kumar, V., Sethi, T. (2021). Learning the mental health impact of COVID-19 in the United States with explainable artificial intelligence: observational study. JMIR mental health, 8(4), e25097.*

on mental health (attribute soc5a), the conditional probability of which remained unchanged (62.2%) across the responses. Although medical conditions (attributes phys3a

to phys3m) are known to increase the risk of serious illness from COVID-19, our model showed that having cancer (attribute phys3k) and hypertension (attribute phys3b) had a reverse impact on anxiety levels. Those with cancer had approximately 8.3% (with 2% CI) higher conditional probability of having less than 1 anxiety-ridden day in a week (>7% effect for hypertension with CI 1.5%). Additionally, cystic fibrosis (attribute phys3i) and liver disease (attribute phys3j) had wide confidence intervals with nonsignificant differences in mean values **Figure 4.3**.

4.4.3 Impact of Economic Factors

Receiving income assistance through Social Security improved the conditional probability of less than 1 day of anxiety in a week by 10.4% (with CI 1.5%) as compared with the segment of people who did not apply or receive it. Just applying for income assistance led to a 4% improvement **Figure 4.3**. Supplemental Social Security (5.5% with CI 4%) and health insurance (5% with CI 2%) also led to similar results. In addition to this, older adults (>60 years) found health insurance more relaxing than younger people. COVID-19 has also severely affected the financial condition of individuals, which may also lead to mental stress.

4.5 Conclusion

Mental health is a public health concern. Mood disorders and suicide-related outcomes have increased substantially over the last decade among all age groups and genders (Beautrais (2000) Brådvik (2018)). The rapid spread of COVID-19 forced governments worldwide to close public gathering places, schools, colleges, restaurants, and industries. Social isolation, digital communication, and working and schooling from home have become the new normal, and many jobs have been lost. Collectively, this has triggered a high level of anxiety, stress, and depression globally. We did not find studies

that have used models to not only predict but also explain the subtle effects of life situations on mental health. An explainable probabilistic graphical modeling approach with bootstraps and exact inference allowed us to capture many of these effects in a robust manner. Our study revealed that individuals with a prior diagnosis of any mental illness are the most vulnerable for mental illness during the COVID-19 pandemic, which recommends building national-level policies to regularly track their mental status and treat them accordingly. Most importantly, our results reiterate the economic underpinnings of a collective mental health response. Income assistance via Social Security or Supplemental Social Security had a demonstrable effect on the alleviation of anxiety as inferred from our model, which provides the first scientific evidence, to the best of our knowledge, proving the utility of such efforts. The extent of such measures' effect may be captured in such modeling studies conducted in various parts of the world, with widely varying assistance structures during this time.

Our findings from the United States can also stimulate further cultural and social research in other geographies with similar or different social structures. For example, the effects of in-person communication, as opposed to digital connectedness, may be different in countries where community living and joint families are still commonplace, such as India. Digital connectedness was not as effective as talking to a neighbor, at least in the United States, highlighting that these have fundamentally different influences on mental health and need to be further explored in systematic studies. We conjecture that such differences may arise from the evolutionary mechanisms that have shaped human societies to live and share in close physical connectedness. Such an effect has been previously shown in primates kept in isolation who display depressive symptoms (Hennessy *et al.* (2014)). Similarly, parenting and its association with neuropeptide hormones may partially explain ((Kohl *et al.*, 2017)) our results that the presence of kids reduces anxiety levels. Interestingly, the COVID-19 pandemic has created a unique natural experiment on the collective mental health response of individuals to a health

emergency.

The life cycle of such a response may need to be further studied as the world goes through various phases of the pandemic until its resolution. However, our study indicates that the mental health impact is observable within a span of a few months, especially on young individuals. Further research will be needed, ideally in a longitudinal setting, where the same individuals can be surveyed again to understand the dynamics of the collective mental health response.

Our results also highlight that modern technological development in virtual communication is not able to replace natural socializing. Hence, it becomes imperative to design better and more empathetic technological tools that may shape a society and prevent isolation and alienation even while maintaining physical distancing and preventive measures for limiting spread. Personalization and contextualization of such measures will also be important, as our results indicate that persons with previous mental health conditions may be disproportionately affected. Our study has a few limitations. Establishing causal inference in cross-sectional data is nearly impossible, and we acknowledge the possibility of confounders. However, this was precisely the reason we chose the structure learning approach, as some of the confounding influences can be transparently discovered and explained. The ensemble voted structure over the sufficiently large number of bootstrapped structures is expected to be robust, as a set of 101 BNs was found to be sufficiently large enough for this study to address the challenge of incorrect learning. Our approach is best suited as a probabilistic reasoning model to explain mental health determinants and to make predictions, a useful outcome in COVID-19–induced mental health morbidity. We could not explain why anxiety levels may be lower in persons with pre-existing cancer or hypertension. This may be a result of reduced work environment–related stress or more contact with family members at home. However, the current data set is not suited to address this at a finer level of explainability. In addition, we could not comment upon the temporality and persistence

of these effects. Our results are currently limited to only one geography (ie, the United States). However, the relatively large sample size and multiethnic involvement in the survey makes the model representative for most of the ethnicities and influences across the United States; hence, it is likely to hold true in the United States.

CHAPTER 5

Optimizing Potential Interventions: Utilising SDOH for learning effective strategies for COVID-19 vaccine distribution using reinforcement learning

5.1 Preamble

Parts of this chapter have been published in:

"*Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060*".

5.2 Introduction

Vaccines have played a crucial role in combating infectious diseases for 200 years (Greenwood (2014)). The successful eradication of smallpox, a transmittable disease responsible for causing huge casualties, happened due to the gradual and effective widespread use of vaccines [1]. Avoidance of millions of deaths and side effects caused by diseases like polio and tetanus was because of vaccination efforts which prove their impact and significance for global health. Despite this impact, if such a vital resource is not optimally distributed and allocated to the regions and communities where its need is the most, this can lead to huge repercussions that could aggravate the situation and can increase the risk of escape variants in mixed and partially vaccinated communities (Rella *et al.* (2021)). All countries across the globe have eagerly waited for the

[1] https://www.who.int/news-room/feature-stories/detail/smallpox-vaccines

launch of an effective vaccine against *SARS-CoV-2*. In India, the vaccine allocation process started on a first-come-first-serve basis, from January 2021 with the launch of the CoWIN [2] portal by the Government of India. As potential candidates continue to introduce their products in the market, there is an urgent need for optimal distribution strategies that would mitigate the pandemic at the fastest rate possible (White and Angus (2020), Khamsi *et al.* (2020)). At the current rate of vaccination estimated to be 2.2 million doses per day, India may only be able to cover 30% of its population by the end of this year, triggering concerns about a third wave hitting India, especially its rural parts, in the absence of herd immunity[3]. Here we summarize the key factors that will need to be considered for effective mitigation:

- **Scarcity of supply:** Despite large scale production efforts, it is expected that the vaccine will still be a scarce resource as compared to the number of people who would need it. In addition to the vaccine itself, there may also be scarcity in the components leading to its delivery, e.g. syringes. Many researchers estimated that 60-70% of the population would require vaccination to achieve herd immunity. This implies that India would need 1.89 billion syringes to deliver the vaccine en-masse [4].
- **Need for Equitable distribution and Digital inclusion:** A truly equitable distribution will not just be defined by the population or incidence of new cases alone, although these will be strong factors. Other factors ensuring equity of distribution include quantum of exposure, e.g. to the healthcare workforce that needs to be protected. In this paper, we assume that the exposure is proportional to the number of cases itself, although the proposed methodology allows more nuanced models to be constructed. There may also be unseen factors such as vaccine hoarding and political influence, which are not accounted for in this work. Another issue that has recently come into light is regarding the increased inaccessibility of vaccines in India, which are largely booked through the online CoWIN portal, owing to the significant digital divide in the country's population.
- **Transparent, measurable and effective policy:** The design of policy would need to be guided by data, prior to, during and after the vaccine administration to a section of the population. Since the viral dynamics are rapidly evolving, the policy should allow changes to be transparent and effects to be measurable in order to ensure maximum efficacy of the scarce vaccine resource. On the larger scale of states and nations, this would imply continuous monitoring of incidence rates vis-a-vis a policy action undertaken.

Although the aforementioned factors seem straightforward, the resulting dynamics

[2] https://www.cowin.gov.in/
[3] https://api.covid19india.org/
[4] https://www.nature.com/articles/d41586-021-00728-2

that have emerged during the actual roll-out of the vaccine are far too complex for optimal decision making by humans. The daunting nature of such decision-making can be easily imagined for large and diverse countries such as India, especially where health is a state subject. Artificial intelligence for learning data-driven policies is expected to aid such decision making as there would be limited means to identify optimal actions in the real world. A 'near real-time' evaluation as per the demographic layout of states and consequent initiation of a rapid response to contain the spread of COVID-19 (Foster *et al.*) will be required. Furthermore, these policies will need to be contextualized to the many variables governing demand or 'need' for the vaccine distribution to be fair and equitable (Deo *et al.* (2020)). Ground testing of these scenarios is not an option, and countries will have to face this challenge.

Reinforcement Learning (RL) algorithms have shown significant progress for decision making in medicine and public health (Liu *et al.* (2020) Saria (2018) Raghu *et al.* Komorowski *et al.* (2018)). Here, we introduce *VacSIM*, a novel feed-forward reinforcement learning approach for learning effective policy combined with near real-time optimization of vaccine distribution and demonstrate its potential benefit through the example of five states presenting contrasting contexts in India. Since real world experimentation was out of the question, the change in projected cases obtained via a standard epidemiological model was used to compare the *VacSIM* policy with a naive approach of incidence-based allocation. Finally, our novel model is open-sourced and can be easily deployed by policymakers and researchers, thus can be used in any part of the world, by anyone, to make the process of distribution more transparent.

5.3 Methods

To get an optimal distribution of vaccine, we proposed a novel pipeline where we combined *deep reinforcement learning* models with a *contextual bandit* model in a feed-forward way which makes the optimization process more robust, context-specific and

easily deployable in real-time. In the following sections, we are adding some background information on Reinforcement Learning Algorithms in order to provide a sense of our pipeline.

5.3.1 Q Learning

Policy learning algorithms can be broadly categorized into two types: off-policy and on-policy. Off-policy learning algorithms are preferred as these allow the agent to learn the optimal policy regardless of the agent's actions (under the current policy). This not only facilitates faster learning (convergence) but also avoids learning of a sub-optimal policy by the agent as it can explore continuously. One of the most popular off-policy algorithms in *RL* is Q-learning. Q-Learning uses the Bellman Equation for iterative updating:

$$Q^*(s,a) = E[r + \gamma max_{a'} Q^*(s',a')] \tag{5.1}$$

where r is the immediate reward and γ is the discount factor. Given that an agent follows a policy π, then Q_π(s,a) denotes the expected return or reward obtained by the agent if it chooses action a at state s and then keeps on following the policy π. $Q(s,a)$ estimates how favourable it is to choose action a at state s according to the current policy. The optimal Q function, $Q^*(s,a)$ denotes the maximum return or reward that can be achieved by first choosing action a at state s and then following the optimal policy. As evident from equation (1), $Q^*(s,a)$ is the sum of the immediate reward and the maximum reward possible thereafter from state s' discounted by γ (same as γ discounted maximum Q value for state s'). However, Q-Learning becomes impractical for environments where there may be a large (potentially infinite) set of actions to choose from. An example of a large environment with continuous action space (Komorowski *et al.*

(2018)) is optimization of number or syringes or medication dosage.

5.3.2 Deep Q-Network

Deep Q-Networks *(DQNs)* are useful in the presence of large environments with elaborate state and action spaces. Here Q-learning is extended to *DQN* by using a deep neural network used to map input states to (action, Q-value) pairs(Hester *et al.* (2018), van Hasselt *et al.* (2015)). Most real-world scenarios such as healthcare resource allocation have complex state-action spaces, and thus need advanced approaches such as *DQNs*. A *DQN* minimizes the loss at each iteration i as given below:

$$L_i(\theta_i) = E_{s,a,r,s' \sim \rho(.)}[((y_i - Q(s,a;\theta_i))^2] \tag{5.2}$$

where $y_i = r + \gamma \, max_{a'} \, Q(s',a';\theta_{i-1})$ is the Temporal Difference (TD) target; y_i - Q is the TD error and ρ represents behaviour distribution.

In our work, we employ *DQN* to better parameterize real-world conditions and make our model more efficient in terms of time and memory requirements.

5.3.3 Actor-Critic using Kronecker-Factored Trust Region (ACKTR)

Q-learning and DQN are both value-based algorithms, where the goal is to learn a single deterministic action from a discrete set of actions by finding the maximum Q value. There are no trainable parameters in value-based methods that control the probabilities of action. Thus, a value-based method cannot solve an environment where the optimal policy is stochastic requiring specific probabilities. To overcome this challenge, we need to learn a stochastic map from state to action. In Policy Gradients (Williams (1992), Schulman *et al.* (2018)), we usually use a neural network (or other function

approximators) to directly model the action probabilities $\pi_\theta(a|s)$. Each time the agent interacts with the environment, we tweak the parameters θ of the neural network so that "good" actions (actions giving maximum reward) will be sampled more likely in the future. We repeat this process until the policy network converges to the optimal policy π^*.

Formally, the objective of Policy Gradients is to maximize the total future expected rewards $E[R_t]$, where R_t represents the sum of future discounted rewards. Whereas, in Q-learning and DQN the "policy" followed the action that maximized the Q value at each step, actor-critic methods combine Policy Gradient methods with a learned value function. With actor-critic methods (Mnih *et al.* (2016)), we learn two different functions: the policy (or "actor"), and the value (or the "critic"). In this paper, we used the ACKTR algorithm (on-policy algorithm) (Wu *et al.* (2017)) with the Policy Gradient method which follows an actor-critic architecture: one deep neural network to estimate the policy and another network to estimate the advantage function (a measure of how good or bad a certain action is at a given state). ACKTR uses the Kronecker-factored approximation (Ly *et al.* (2017), Martens (2020)) to optimize both the actor and critic.

5.3.4 Bandit Algorithms: Multi-armed bandits and Contextual Bandits problem

Reinforcement algorithms being computationally expensive and time-consuming, cannot be used in real-time decision making. Contextual Bandits (Collier and Llorens (2018)) algorithms on the other hand, are space and time-efficient. These are extensions of multi-armed bandits algorithms. Multi-armed bandits can simply be understood as a gambler with a fixed amount of money taking a sequence of actions on multiple slot machines while trying to optimize the overall reward. The K machines have reward probabilities, $\theta_1, \ldots, \theta_K$. $\langle$Action(A),Reward(R)$\rangle$ tuples describe a multi-armed bandit, where at each time step *t*, an agent takes an action a on one slot machine and receives

a reward r. The goal is to maximize the cumulative reward. This also corresponds to minimization of the potential loss by not picking the sub-optimal action, in case optimal actions with best reward are known, e.g. vaccination of a high-risk individual. The optimal reward probability θ^* of the optimal action a^* is:

$$\theta^* = Q(a_*) = \max_{a \in S}\{Q(a)\} = \max_{1<i<K}\{\theta_i\} \tag{5.3}$$

where $Q(a)$ is the expected reward.

The Loss function or *Regret* (Riquelme *et al.* (2018)) is a conceptual function to understand and optimize the performance of the Bandits problem. Since we don't know if an action played was the most 'reward-fetching', rewards against all the actions that can be played are sampled, and the difference between the action chosen and the action against the maximum reward is defined as 'regret'. Therefore, minimizing regret achieves the goal of maximizing reward. The total regret we might have by not selecting the optimal action up to the time step *T*:

$$L_t = E[\sum_{t=1}^{t=T}(\theta^* - Q(a)] \tag{5.4}$$

Thus, a multi-armed bandit algorithm returns the best set of actions, provided the context is fixed. However, a multi-armed bandit algorithm does not account information from the changing context, a common occurrence in real-world scenarios such as the evolving pandemic. Therefore the agents may end up playing the same action multiple times even though the context may have changed, thus getting stuck at a sub-optimal condition. To circumvent this problem, Contextual Bandits algorithm is proposed which contextualizes the choice of the bandit to its current environment. Contextual Bandits, represented as a tuple $\langle$Action(A),Context(C),Reward(R)$\rangle$, play an action based on its

current context, given a corresponding reward, hence are more relevant to real world environments such as the vaccine distribution problem addressed in this work. Given, for time t = 1...n, a set of contexts C and a set of possible actions A and reward/payoffs R are defined. At a particular instant, based on a context $c_t \in C$, an action $a_t \in A$ is chosen and a reward or a payoff

$$R = E[r \mid c, a] \tag{5.5}$$

is obtained. For an optimal action $a^* \in A$ such that the expectation of reward against this action is maximum,

$$r^* = \max_{a_t \in A} E(r \mid c, a_t) \tag{5.6}$$

the regret can be expressed as

$$Z = [r^* - E(r \mid c, a_t)] \tag{5.7}$$

and *cumulative regret* can be expressed as

$$Z^* = \sum_{t=1}^{n} Z \tag{5.8}$$

The Contextual Bandits model was implemented using the python package of Vowpal Wabbit [5].

[5]https://vowpalwabbit.org/

5.3.5 *VacSIM*: A Feed forward Approach

In **VacSIM**, we concatenated the Deep Reinforcement Learning models (i.e. ACKTR or DQN) and Contextual Bandits model in a feed-forward manner to create an off-policy learning framework. This was incorporated to select the optimal policy through the Contextual Bandits approach from the ones generated by the Deep Reinforcement Learning models, as shown in **Figure 5.1**. This was done to address the following challenges that need to be tackled in real world problems such as optimal vaccine distribution:

- **Solving in real-time:** A vaccine distribution problem is expected to be fast-paced. Thus, an overwhelming amount of brainstorming with constrained resources and time would be required to develop an effective policy for the near future. This calls for the development of a prompt and an easily reproducible setup.
- **Lack of ground truth:** This is one of the key challenges in this paper. Since the roll-out of the vaccine will not give us the liberty of testing various hypotheses, a lack of ground truth data is generally expected. This is analogous to zero-shot learning problems, thus precluding a simple supervised learning-based approach.
- **Absence of evaluation with certainty:** Lack of ground testing naturally implies nil on-ground evaluation. In that case, it often becomes challenging to employ evaluation metrics that offer a significant amount of confidence in results. In order to solve this problem, we rely upon the simulated decrease in the number of susceptible people as the vaccine is administered.
- **Model scaling:** We have ensured that the learning process of the models simulates the relationship between different objects in the real world environment accurately, and at the same time can be scaled down in response to computational efficiency and resource utilization challenges. This is done by choosing the right set of assumptions that reflect the real world scenario.

We pipelined the Deep Reinforcement Learning models with a Contextual Bandits approach where recommendations for vaccine distribution policy (generated by the RL models) were used as training data for the Contextual Bandits. These models were trained separately using the same reward function defined in equation (6).

Reinforcement Learning approaches replicate human decision-making processes. However, the absence of evaluation makes them less trustworthy, especially when real lives

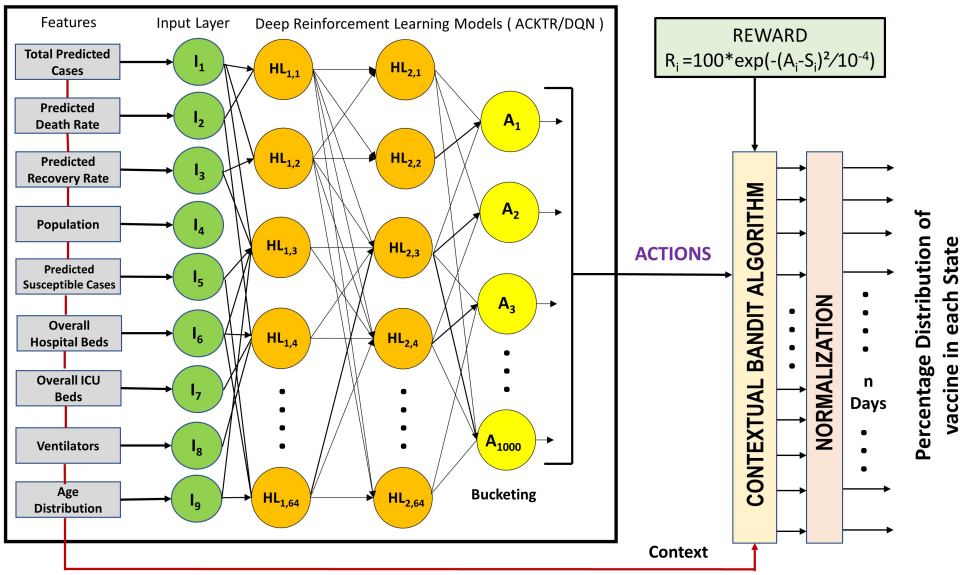


Figure 5.1: **VacSIM architecture.** A novel feed forwarded pipeline for policy learning of COVID-19 vaccine optimal distribution. The actions and rewards obtained from the Deep Reinforcement Learning models were fed forward into the training of Contextual Bandit algorithm so that faster optimal online decisions could be calculated for further dates taking in context, the ever changing demographics of the states. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

may be at stake. To this end, we also presented two different novel model evaluation techniques in this paper:

1. Estimating future impact of vaccination by projecting COVID-19 cases, using SEIR models (Li and Muldowney (1995)).
2. Learning Bayesian Network to verify the underlying structure of COVID-19 trajectory indicators.

5.4 Results

5.4.1 Model Setup

The State-wise time series data of COVID-19 characteristics from the website of the Ministry of Health and Family Welfare of India [1] were used in the experiments as de-

[1] https://mohfw.gov.in/

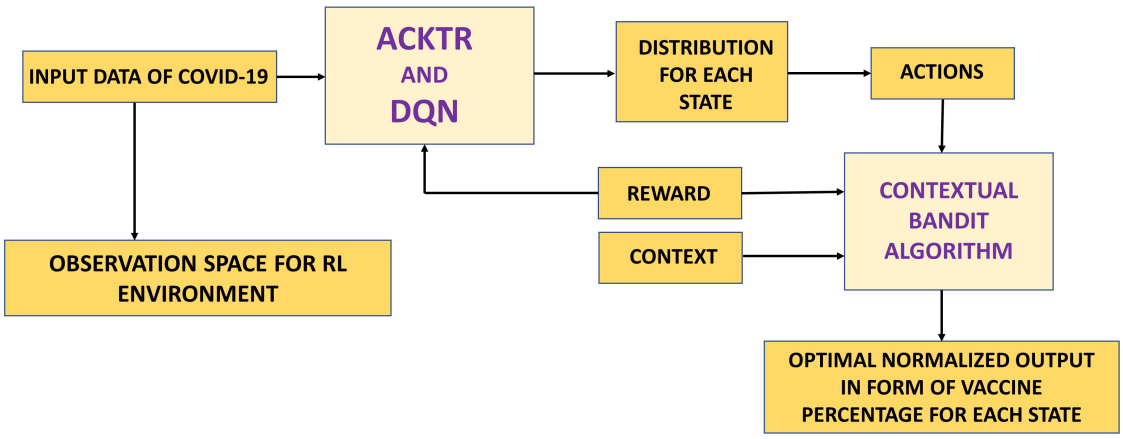


Figure 5.2: **VacSIM model setup.** Flexible model setup for optimization of vaccine distribution policy in India using the VacSIM approach where components can be easily replaced with alternatives and adopted for diverse settings. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

scribed in **Figure 5.2**. In order to determine vaccine allocation in the near future, we require future scenarios of COVID-19. Thus, in order to train VacSIM, we projected Susceptible, Confirmed, Recovered cases from 1 December 2020 to 26 December 2020 using SEIR model (refer supplementary for model details). The five States chosen for this study, i.e., Assam, Delhi, Jharkhand, Maharashtra and Nagaland are representative of possible scenarios for the vaccine distribution problem, i.e., high incidence (Maharashtra and Delhi), moderate incidence (Assam) and low incidence (Jharkhand and Nagaland). In choosing the five different States, we hope to generalize our predictions to other States across the spectrum while minimizing the bias introduced into the learning by a widely variant *COVID-19* incidence across the country.

In order to avail vaccine for the highly vulnerable population and to combat COVID-19 infection spread, we selected the context comprising features listed in **Table 5.1**.

S.No.	Feature	Description
1	Predicted Death Rate	The percentage ratio of the predicted deaths in the State to the total predicted cases in that State calculated using projections obtained from a fitted standard *Compartmental model*, i.e. a Susceptible, Exposed, Infected and Recovered (*SEIR*) model.
2	Predicted Recovery Rate	The percentage ratio of the predicted recoveries in the State to the total predicted cases in that State using projections obtained from the SEIR model.
3	Population of a State	We extracted the population for each State from the 2011 census data conducted by the government of India [6].
4	Predicted Susceptible Cases of a State	We estimated the susceptible population as the difference between the population of a particular State and the total number of predicted infected cases of that State.
5	Hospital Facilities in a State	We used overall Hospital Beds, overall ICU Beds and Ventilators data in our models[7].
6	Age Distribution of a State	In order to prioritize vulnerable population, we considered people with age greater than 50.

Table 5.1: Features present in context space used in our models.*Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

The implementation of the *VacSIM* pipeline is detailed henceforth.

5.4.2 Deep Reinforcement Learning Models

Open-AI([8]) stable-baselines framework was used to construct a novel and relevant environment suited to our problem statement for ACKTR and DQN to learn in this environment.

Input

The observation space(s) was considered as a grid over context vectors with the latter describing the context vector of states. These include Population, Age Distribution, Ventilators, Overall ICU Beds, Overall Hospital Beds, Predicted Susceptible Cases, Total Predicted Cases, Predicted Death Rate, and Predicted Recovery Rate at a given time. The predicted rates were estimated by using standard SEIR projections.

The Action space describes the quantum of vaccines distributed to each of the states across India, represented as discrete buckets. For example, a bucket size of 1000 implies a batch of 1000 vaccines rolled out to a particular state of India. Formally, the action

[8]https://openai.com/

space a was defined as, a= $\{b_0, b_1, b_2, b_3, ..., b_i, ..., b_{n-1}\}$

$$b_i = i * \frac{Total\ Available\ Vaccine}{n} \tag{5.9}$$

where *n* is defined as *bucket size* and considered as a hyper parameter.

Training

Following are the assumptions used while building the environment:

- The nature of the vaccine is purely preventative and not curative, i.e., it introduces immunity when administered to a susceptible person against future COVID infection but plays no role in fighting against existing infections (if any).
- The vaccine has 100% efficacy, i.e. all people who are vaccinated will necessarily develop immunity in 45 daysPolack *et al.* (2020) against future *SARS-CoV-2* infection. This assumption is easily modifiable and is expected to have little consequence in deciding the optimal allocation unless data reveal differential immunity response of certain populations within India. However, we leave scope for it to be modified as per the situation by taking this as a hyper-parameter for *VacSIM*.

Reward Function: The reward function was designed to maximize the decrease in the susceptible population count with the minimum amount of vaccine, considering it to be a limited resource.

$$R_i = 100 * exp(-(A_i - S_i)^2/10^{-4}) \tag{5.10}$$

where R_i is the reward given by the environment; S_i is the fraction of susceptible population in a particular State and A_i is the fraction of vaccine distributed to the State by the model.

The Explore-Exploit Trade-off: We tap into the explore-vs-exploit dilemma, allowing the model to reassess its approach with ample opportunities and accordingly, redirect

the learning process. We set the exploration rate at 10% for DQN. However, this too is flexible and can be treated as a hyper-parameter. Here we reason that the local context is most strongly influenced by its own past values, hence we trained our model independently for the selected 5 States from 1 December to 26 December, in order to maintain spatial independence. 10000 and 2000 iterations in **Figure 5.3** were used to learn DQN and ACKTR respectively. Iterations for training were decided based on the stability of the reward function curve. Where Iterations refers to one complete exploration of environment by a model agent

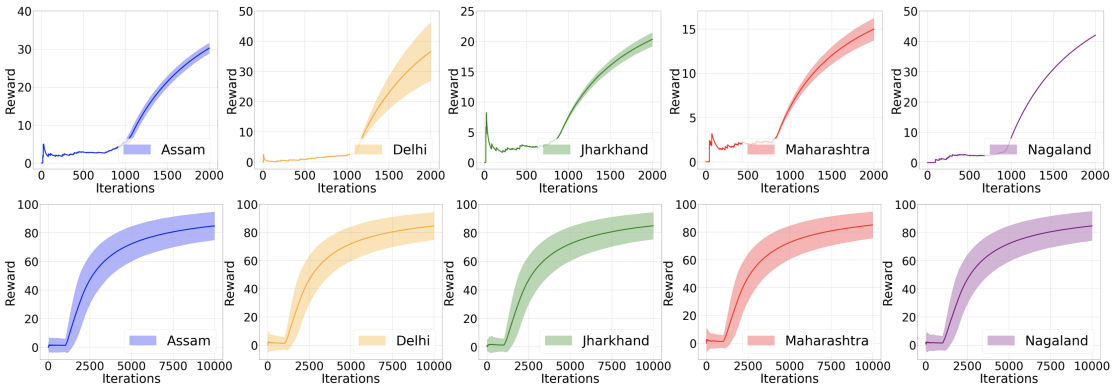


Figure 5.3: **Optimization.** Day-wise rewards at every iteration were recorded for all the states. The Mean and SD obtained over 26 day projections are shown as confidence bands. (Top) Smoothed and Increasing Reward Curve of DQN (Bottom) Smoothed and Increasing Reward Curve of ACKTR. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

Hyper-parameters: A complete list of the hyper-parameters is given in **Table 2**.

S.No.	Hyperparameter Name	Hyperparameter Value
1	Batch Size: number of vials in one round of distribution.	1000000
2	Exploration Rate of DQN	10%
3	Vaccine efficacy	100%
4	Number of days to reach full efficacy	45
5	Bucket size	1000
6	Number of recipients per day	5

Table 5.2: Hyper-parameters used during Policy learning. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

Output

The output of the Deep Reinforcement Learning models was a distribution set dictating the share of each recipient in the batch of vaccines dispatched at the time of delivery.

5.4.3 Contextual Bandits Model

The output of the Deep Reinforcement Learning model is spanned over 26 days (1 December-26 December 2020). The distribution sets so obtained were scaled to bucket sizes ranging from 100 to 1000 with an increment step of 100. The actions of the five States associated with each bucket size were normalized to get the percentage distribution ratio for all States. Normalized here refers to the percentage ratio of a given distribution set of a State for a given bucket size and the sum of the distribution sets of the five States over that bucket size for a given date.

Training

We considered 1 December-26 December data along with distribution sets and the corresponding set of rewards obtained from the Deep Reinforcement Learning models as training data for Contextual Bandits.

Number of actions and features

The action space and the features in the Context space are the same as those in the Deep Reinforcement Learning model.

Testing

Using the Context, actions and the corresponding set of rewards as the training dataset, we tested the model day-wise for a period of five days (27 December-31 December)

with each day having 10 possible bucket sizes, i.e. 100 to 1000 for each State as output.

Output

The unadjusted actions (which were not normalized) obtained after testing the model were first adjusted bucket-wise for each State for a particular date by taking the percentage ratio of the unadjusted action of that State and the sum of unadjusted actions of all five States.

We consider that the vaccine is distributed on 31 December 2020. The consequent distribution of vaccines among the five States for each bucket size was then evaluated. Our models learn well for the bucket range 200-500. The vaccine distribution for the same in case of the VacSIM policies is shown in **Table 5.3**. It was observed that both ACKTR+CB and DQN+CB yielded similar results, which is conceptually expected as both were aimed at reducing the overall susceptible population as optimally as possible.

Model	***Bucket Size***	***Assam***	***Delhi***	***Jharkhand***	***Maharashtra***	***Nagaland***
ACKTR	200	15.5	8	17.5	58	1
+	300	15.5116	7.9208	17.4917	57.7558	1.3201
CB	400	15.5172	8.1281	17.734	57.3892	1.2315
	500	15.5206	8.055	17.6817	57.3674	1.3752
DQN	200	15.8974	8.2051	16.9231	58.4615	0.5128
+	300	15.9864	8.1633	17.0068	58.1633	0.6803
CB	400	16.1616	8.3333	17.1717	57.5758	0.7576
	500	16.129	8.2661	17.1371	57.6613	0.8065

Table 5.3: VacSIM model output on different bucket size. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

5.5 Novel Approaches to Evaluate *VacSIM* Model

We have utilised two alternative ways for evaluating reinforcement learning model based vaccination allocation outputs. Vaccine allocation was conducted in a spatiotemporal way therefore our approaches are able to capture spatiotemporal outcomes.

5.5.1 Model Evaluation through Projection Scenario using SEIR models

This approach may be used to compare the long-term impacts of policies. We compared the learned distribution from our model to the Naive baseline distribution policy, which was defined as percent of vaccination supplied to a State Equals percent of Infected individuals within the State. The vaccine's distribution was simulated on December 31, under the *VacSIM* and Naive policies, using 100,000 doses and 5 States. The SEIR Model was used to determine the number of resulting illnesses for 45 days following the distribution of the vaccine. For both models (ACKTR/DQN), the daily total cases across all 5 States were summed up for each bucket size (200,300,400,500), and their differences were then determined using total cases for all 5 States came up after Naive distribution as indicated in **Figure 5.4**. Our findings show that the *ACKTR* based strategy additionally eliminates a total of 8845, 7787, 4703, and 5103 predicted contaminated instances over the course of the following 45 days, with bucket sizes of 200, 300, 400, and 500, respectively. A total of 9039, 6686, 5355, and 4698 projected contaminated instances are reduced by *DQN* based policy as well, with bucket sizes 200, 300, 400, and 500, respectively.

5.5.2 Model evaluation through learning the causal structure of simulated data obtained

This strategy may be used in the absence of a ground truth policy. This method demonstrates the causal association between model outputs and the influence variable, confirming the validity of model outputs. The reduction of mortality and morbidity is the main objective of vaccination delivery. Given that our model solely relied on simulations in the absence of real-time vaccination data, we evaluated whether the data generated by such a method obey the cause-and-effect linkages as anticipated in the real

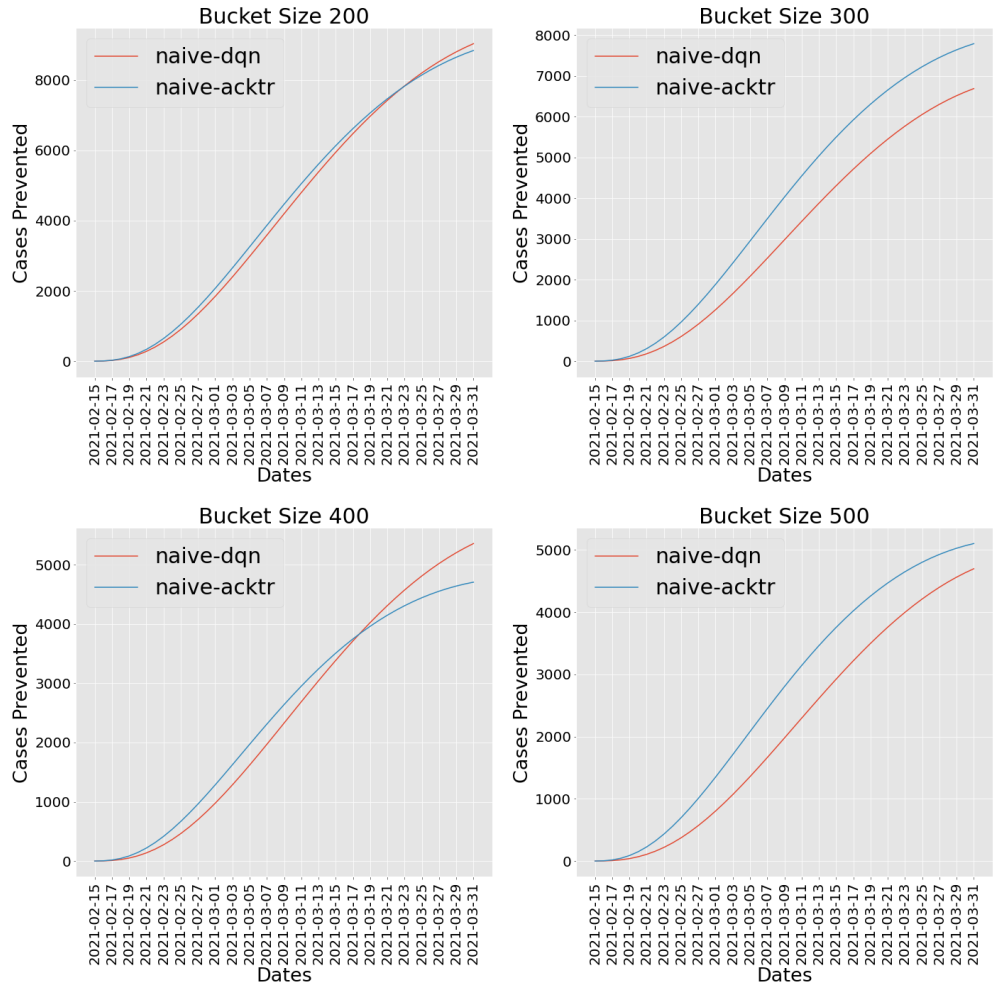


Figure 5.4: **Policy evaluation through SEIR model.** Additional projected infection cases prevented in next 45 days by following VacSIM driven approach instead of Naive approach. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

world data. A special type of probabilistic graphical model that is helpful for modeling causal links is a Bayesian network.

To learn the BN structure **figure 5.5**, state-wise time series data from our model were employed, including data on mortality, recovery, susceptible and sick people, and the quantity of vaccination. In order to prevent the vaccine percentage from being the child node of the *COVID-19* trajectory indicators, blacklisting among nodes was done (Infected People, Susceptible People, Death, and Recovery). The resulting structure demonstrates a causal relationship between the dose of the vaccination (parent node)

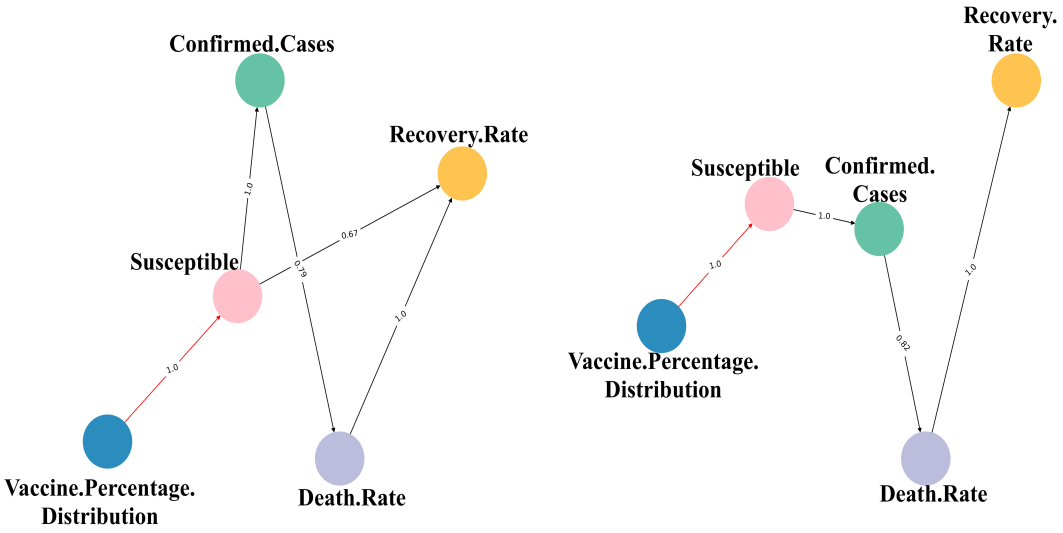


Figure 5.5: **Policy evaluation through causal model.** Ensemble averaged causal structure of the Bayesian network obtained from 501 bootstraps, using Hill Climbing optimizer for AIC (left) and BIC (right) as scoring functions. Vaccine Percentage obtained from the model was observed as a parent node of Susceptible cases, thus indicating the causality preserving nature of VacSIM simulations. *Credits: Awasthi, Raghav, et al. "Vacsim: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning." Intelligence-Based Medicine (2022): 100060.*

and the number of susceptible individuals (child node), demonstrating the accuracy of the *VacSIM* model from a technical standpoint.

5.6 Discussion:

Researchers worldwide have been working round-the-clock to find a vaccine against *SARS-CoV-2*, the virus responsible for the *COVID-19* pandemic. Now that it is available, the distribution of the vaccines has numerous logistical challenges (supply chain, legal agreements, quality control and application to name a few) which might slow down the distribution process. To circumvent these challenges, a strong decision support system is needed (Avorn and Kesselheim (2020)). Reinforcement Learning (RL) based decision models (Kwak *et al.* (2021*b*)) have been built during COVID-19. In this paper, we have developed a novel distribution policy model, *VacSIM* using Reinforce-

ment Learning. We have pipelined an Actor-Critic using Kronecker-Factored Trust Region (ACKTR) model/ Deep Q-Networks (DQN) model and Contextual Bandit model in a feedforward way such that the output (Action and Rewards) of ACKTR/DQN model are fed into the Contextual Bandit model in order to provide a sensible Context comprising of actions and rewards. Contextual Bandits then optimize the policy considering demographic metrics such as the population of the State, Overall Hospital Beds, Overall ICU Beds, Ventilators and time series-based characteristics of the *COVID-19* spread (susceptible population, recovery rate, death rate, total infected cases) as Context. While distributing the vaccine, identifying the part of the population who needs it the most is a challenging task, and in our case, we addressed it by the usage of the aforementioned Context. Rather than using the present-day count of infected and susceptible people, we have used SEIR-based projection, which makes our predicted policy more robust and trustworthy. The percentage of vaccine distribution was the final output of VacSIM. From VacSIM, Maharashtra, the most vulnerable state for COVID-19, has received the maximum amount of vaccine, which showed the model could understand the context. Finally we proposed two approaches to evaluate such models 1) SEIR-based Simulation-based approach, 2) Bayesian Network-based Causality driven approach. In simulation-based approaches, we have projected the COVID-19 cases after the vaccination. Our results indicate that the VacSIM additionally reduces the 9039 cases compared to the naive-based approach. And from the Bayesian network, we found that vaccine distribution is most causally associated with susceptibility, which again showed that VacSIM output followed the design of model set-up and variability of context. Our model design and evaluation are entirely novel. Any other resource allocation algorithm like Genetic Algorithm, Simulated annealing, Hill climbing is reward independent. A comparison of policy with these algorithms are shown in our dashboard (http://vacsim.tavlab.iiitd.edu.in:8000/). VacSIM outperformed these models when compared using the SEIR based simulation approach. It is necessary to update models with new influences on a regular basis in order to capture varying trends occur-

ring with emerging variants of COVID-19. In order to do so for different settings, our model is extendable and context independent, therefore will require re-parameterization to adapt to the new setting and learn optimal policy accordingly. Finally, we have provided open-source code for our pipeline which will enable testing of our claims by other researchers.

5.7 Limitations and future work

VacSIM may have some limitations shared by all Reinforcement Learning models, i.e. the transparency of their learning process and the explainability of their decisions. The development of *VacSIM* has been carried out while observing the pandemic in the past few months. However, the dynamic nature of pandemic may require a change in actions, thus calling for common sense working alongside artificial intelligence. A third limitation is the potential dependency between COVID-19 trajectory variables across states as these may not be completely independent especially after the lifting of lockdown. However, parametrizing these dependencies is a challenge and we reason that the local context of the states is most strongly influenced by its own past values. As mentioned before, the availability of authentic data that is granular in terms of features covered is a major challenge in any kind of modelling work associated with the pandemic. By virtue of the same, we have given utmost priority to achieving herd immunity and decreasing the susceptible population as much as possible, given that proper data capturing of those features is available. The same is reflected through our choice of reward function in this work. For future work, if a more reliable and extensive data source is available describing the information of different cohorts of the Indian population on a more atomic level (such as information regarding the comorbidities of a certain region), then such relevant features should be taken into account while designing the reward function. Similarly, designing reward functions with features which are important for the efficient distribution of scarce resources like ventilators, drugs on the availability of

extensive reliable data sources can be done through our framework depending on the observation and action space.

It should also be noted that while we assume certain vaccine characteristics, such as efficacy and gap between both doses, these are hyperparameters that can be changed depending on the availability of exact reliable data. To date, there does not exist a single well-proven statistic for the dose-wise efficacy of Covaxin and Covishield (vaccines available in India currently). Even the gap between both doses has been changed repeatedly by the Government, the latest being 12-16 weeks [9] for Covishield. In the emergence and spread of various mutations of COVID-19 virus, and responding variants of vaccines entering the market, such assumptions have been taken to make the pipeline as generalised as possible, so that when required, different pieces of information can be plugged into our framework and adjusted to identify the efficiency for different resources and context. Given these limitations, we have attempted to carefully design the VacSIM pipeline and associated reward function by making realistic assumptions which can be helpful in testing such a framework before moving to on-ground deployment.

To truly facilitate our ultimate goal of on-ground deployment of VacSIM, Plug-n-Play(PnP) pipeline can be created that gives policymakers the flexibility to provide varying input contexts, such as those collected on the district level. This platform would be helpful in testing our framework's efficiency in a variety of settings, across many data points, observation spaces and even reward functions. In conclusion, we believe that artificial intelligence has a role to play in the optimal distribution of resources such as vaccines, syringes, drugs, personal protective equipment (PPEs), ICUs etc. The second wave of COVID-19 in India has also highlighted the importance of the same and our open source VacSIM framework is a potential step towards enabling use of artificial intelligence by researchers and policymakers.

[9] https://www.bmj.com/content/372/bmj.n18

5.8 Conclusion

We provide a novel, open-source and extensible solution vaccine allocation problem that policymakers and researchers may refer to while making decisions. This feedforward network can be adapted and used for optimal allocation of various essential resources in varying scenarios and context as the Reinforcement Learning part of VacSIM can always be retrained to adapt to spatial and temporal variations by accommodating the latest values of variables like recovery rate and age distribution. Since there is some evidence that COVID-19 infections tend to show oscillatory patterns on shorter timescales (Bukhari *et al.* (2020)), the key goal of our approach to inform an agile response to the problem of vaccine distribution has been fulfilled through this work. Although the present results are in five states, *VacSIM* will be extended to all the states of India.

CHAPTER 6

Deployment:Web-Applications for Guiding Policy Interventions for Combating COVID-19

6.1 Intoduction

COVID-19 has posed several hurdles throughout every stage (Stone (2020), De Camargo (2022)). In the early stage, there was a severe lack of information; clinicians were uninformed about the optimal course of treatment, and the government was uncertain about the infection trajectory. Once the vaccine was discovered, optimum distribution became challenging. To combat these challenges, real-time data played a vital role and revolutionized public health policy modeling by using data more comprehensively. AI models were used to simulate the pandemic decisions. Interactive dashboards were utilized to present AI model outputs. Dashboards/Web-Applications are an interactive way to show data graphically. DashboardsWeb-Applications can quickly display real-time (or almost real-time) data dynamics.

Below **Table 6.1** describe several dashboards developed at various stages of COVID-19 with the collaboration of a multidisciplinary team.

6.2 CoVision: A Modeling and Simulation Suite for COVID Scenarios in India

COVID-19 infection trajectory projection is a challenging task due to biological and health system factors(Jewell *et al.* (2020) Rahmandad *et al.* (2022) Watson *et al.* (2021)

Dashboard Title	Background AI Models	Purpose	Link of Dashboard
CoVision	Agent based models, Compartmental Epidemiological models	Projection of COVID-19 trajectory	http://covision.tavlab.iiitd.edu.in
CovidNLP	Text Summariser (NLP)	COVID-19 Literature summariser	http://covidnlp.tavlab.iiitd.edu.in/
EvidenceFlow	Word Embedding, Link Prediction, Machine Learning	Latent Respreseta-tion and prediction ofCOVID-19 themes	https://evidenceflow.tavlab.iiitd.edu.in/.
VacSIM	Deep Q Network, ACKTR, Contextual Bandits(RL)	Vaccine Allocatio of COVID-19	http://vacsim.tavlab.iiit.edu.in:8000/

Table 6.1: AI-powered Dashboards, which we developed across various stages of COVID-19.

Roda *et al.* (2020)). Around the world, many models have been developed for the same purpose, but relatively few studies take real-world assumptions into account. Here, we developed a suite of models to guide the development of policies under different scenarios. These had been deployed to create an interactive dashboard called COVision, which includes the Agent-based Models (ABM) and classical compartmental models (CCM).

Our dashboard allows the simulation of scenarios by changing the strength of lockdown, basic reproduction number(R_0), asymptomatic spread, testing rate, contact rate, recovery rate, incubation period, leakage in lockdown, etc. We optimized ABM and CCMs and evaluated them on multiple error metrics. Out of these models in our suite, ABM could capture the data better than CCMs. Our evaluation suggests that ABM models were able to capture the dynamic nature of the epidemic for a longer duration of time, while CCMs performed inefficiently. We computed R_0 using CCMs, which were found to be decreasing with lockdown duration, indicating the effectiveness of policies in different states of India. Models have been deployed on a dashboard hosted at [1], which allows users to simulate outcomes under different parameters and will enable

[1]http://covision.tavlab.iiitd.edu.in

the policymakers to make informed decisions and efficient monitoring of the covid19 pandemic in India **Figure 6.1**.

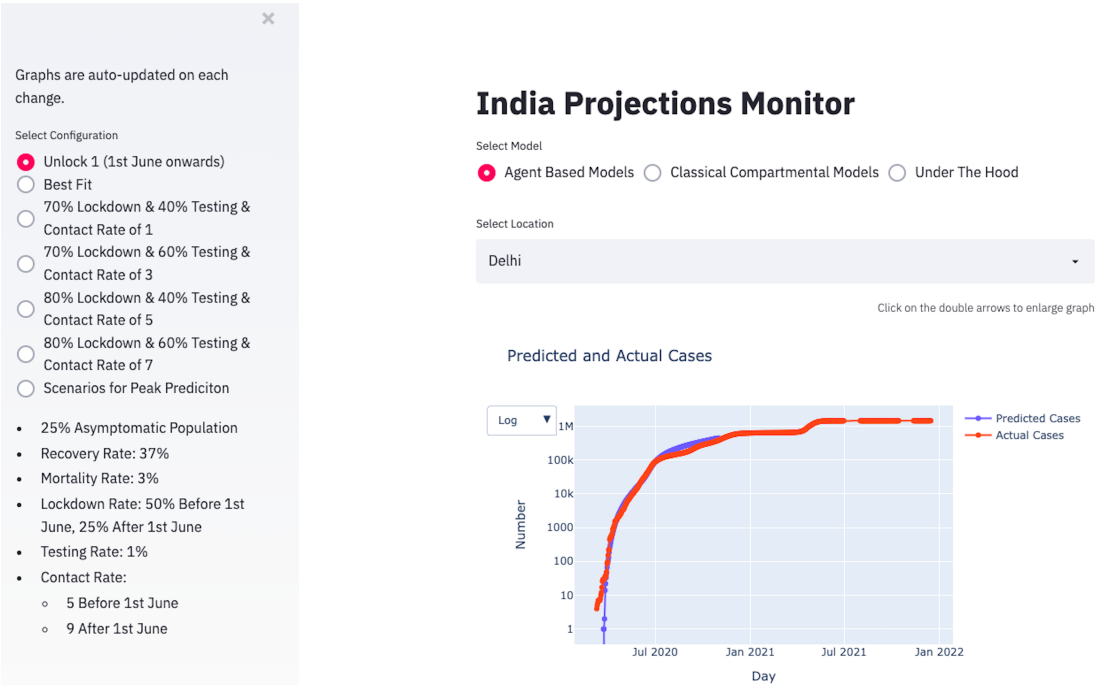


Figure 6.1: **CoVision.** The interactive dashboard displays a state-by-state prediction of the COVID-19 trajectory using various models and scenarios.

6.3 CovidNLP: A Web Application for Distilling Systemic Implications of COVID-19 Pandemic with Natural Language Processing

During the early phase of COVID-19, a flood of contradictory COVID-19 research resulted in a lack of viral understanding (Leng *et al.* (2021) Tagliabue *et al.* (2020)). Millions of research papers on COVID-19 have been published, and since the disease has reached pandemic proportions, doctors and researchers are attempting to synthesize this information to advance clinical care and research. In this work, we address this deficiency for a specific user group, notably physicians, researchers, and policymakers, by employing natural language processing to create a CovidNLP dashboard to accel-

erate information discovery. The World Health Organization (WHO) has compiled a database of peer-reviewed and curated research publications on several topics, such as epidemiology, clinical characteristics, diagnosis, therapy, social issues, and economics. All articles in the WHO Database were summarised using an extractive summarizer, followed by an investigation of the feature space using word embeddings, which were then utilized to show the summarised associations of COVID-19 as discovered in the text. Clinicians, researchers, and policymakers will find not only the direct consequences of COVID-19 but also the systemic ramifications, such as the projected increase in TB and cancer mortality during the export shutdown due to pharmaceuticals' inaccessibility. These examples highlight the usefulness of huge literature mining with natural language processing for quick knowledge updates. This may aid users in comprehending, synthesizing, and taking preventative action based on the existing peer-reviewed information about COVID-19. Our models will be continually updated with new literature, and we have made our resource CovidNLP **Figure 6.2** user-friendly and accessible to the public at [2].

6.4 VacSIM: Learning effective strategies for COVID-19 vaccine distribution using reinforcement learning

Once the COVID-19 vaccine was introduced, optimally distributing COVID-19 resources became a formidable obstacle(Wouters *et al.* (2021), Kahn *et al.* (2020), National Academies of Sciences *et al.* (2020) Islam *et al.* (2021)). Because vaccines were anticipated to be scarce, an appropriate allocation method, particularly in nations with access imbalances and temporal separation of hotspots, might be an efficient means of preventing the spread of the illness. To optimize the distribution of the COVID-19

[2]http://covidnlp.tavlab.iiitd.edu.in/

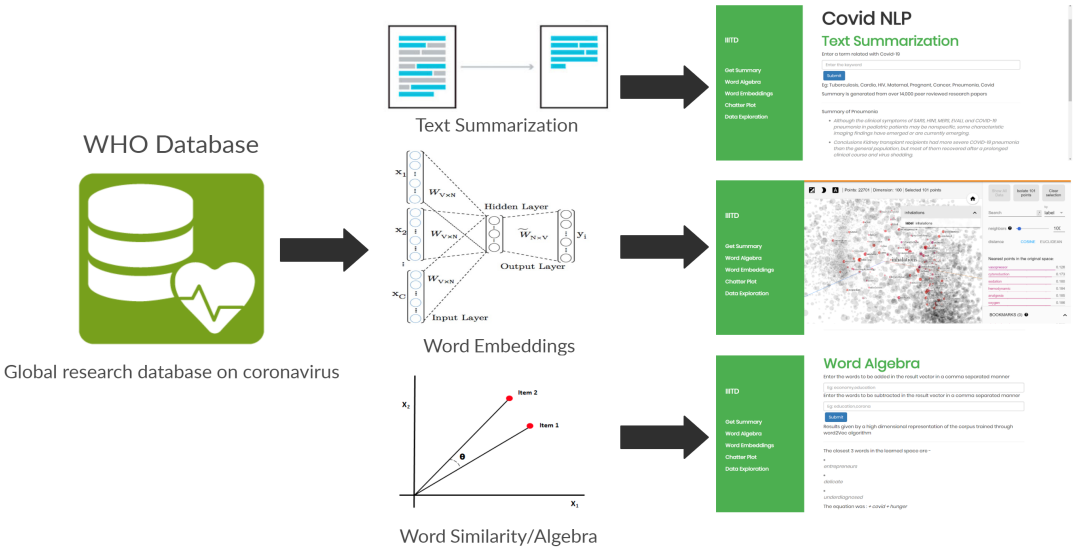


Figure 6.2: **CovidNLP.** The pipeline takes in WHO Database as input and creates text summarization on the article abstract. Word embeddings are generated from the vocabulary abstract to create word embeddings which serve as an input for tensorboard projection and word similarity.

vaccine, we propose a unique pipeline, VacSIM, that combines Deep Reinforcement Learning models with a Contextual Bandits method. In addition, we have compared the VacSIM **Figure 6.3** pipeline against other mathematical models, such as the Genetic Algorithm, Simulated Annealing, and Hill climbing. We have open-sourced our dashboard. Users may access [3] to investigate the vaccination distribution throughout India's various states.

6.5 EvidenceFlow:Predicting emerging themes in rapidly expanding COVID-19 literature with dynamic word embedding networks and machine learning

For formulating solutions to global threats such as COVID-19, peer-reviewed literature is the cornerstone(Chen *et al.* (2021)). Publication-based knowledge and interpretation

[3] http://vacsim.tavlab.iiit.edu.in:8000/

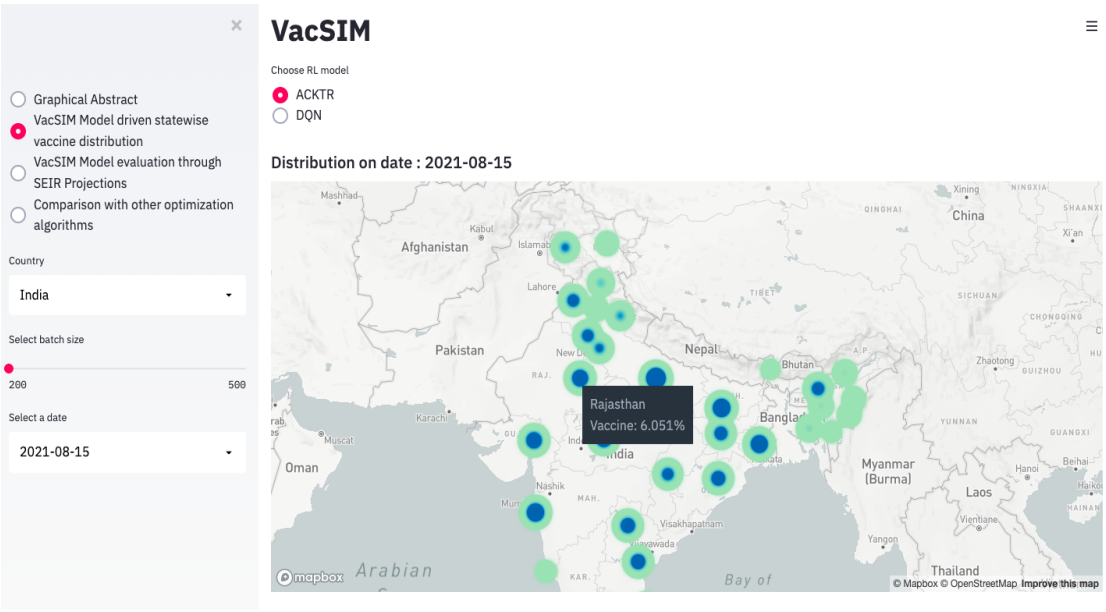


Figure 6.3: **VacSIM.**The VacSIM dashboard displays vaccine allocation by the state daily in a geomap and compares models using SEIR-based projection (Chapter 4). The VacSIM model output was also evaluated with a baseline optimization model. VacSIM has a separate page for presenting model comparisons.

must be condensed into evidence by utilizing natural language in ways that go beyond typical meta-analysis. Several research has concentrated on extracting evidence from text via the use of natural language processing. Using unsupervised word embedding evolution and machine learning, we demonstrate that new information can be recorded, monitored, and predicted. Our method for deciphering latent information flow in time-varying word-vectors networks. Our web application, EvidenceFlow **Figure 6.4**, is accessible to the public at [4]. It provides the results of online monitoring and forecasting of developing themes.

6.6 Conclusion

Our deployed dashboards helped with decision-making in the war against COVID-19. In the midst of the ambiguity surrounding COVID-19, *CovidNLP* and *EvidenceFlow* developed a paradigm and an interface for clinicians, academics, and policymakers to

[4]://evidenceflow.tavlab.iiitd.edu.in/

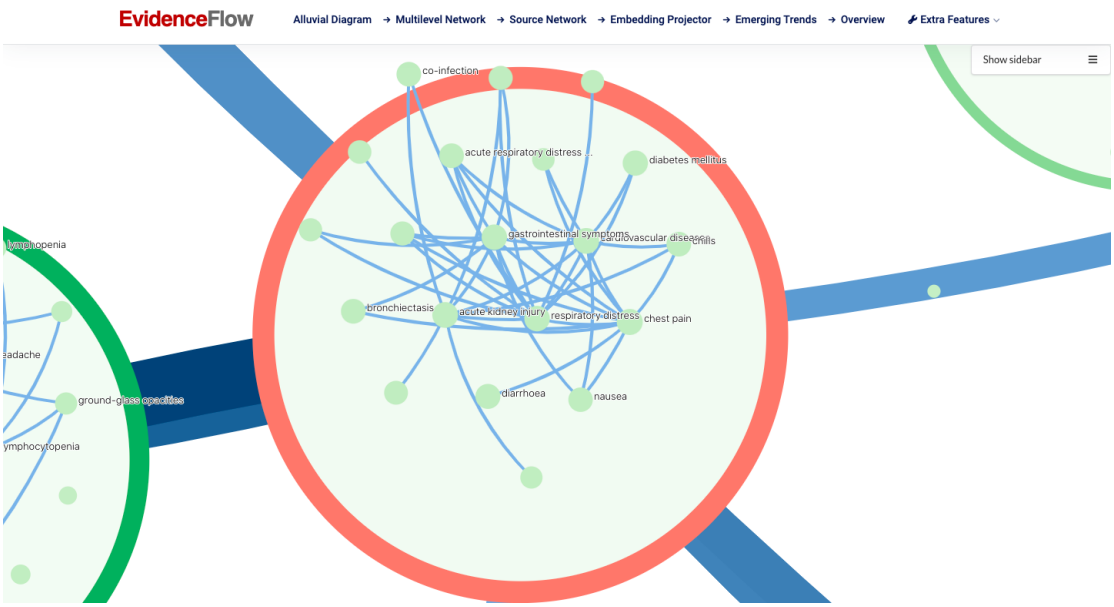


Figure 6.4: **EvidenceFlow**.The Evidence flow Dashboard interactively presents the latent themes in the form of an unsupervised work embedding network. It also includes a word embedding projector and alluvial diagram functionality.

extract pertinent information. *CoVision* offered visualisation of the COVID-19 infection trajectory, while *VacSIM* provided a solution and decision assistance for the distribution of COVID-19 vaccines.

CHAPTER 7

Conclusion and Future work

7.1 Conclusion

The social determinants of health (SDOH) influence health outcomes significantly. However, their impacts are difficult to measure. These are a part of some of the most challenging problems due to the complex nature of interdependencies and indirect effects, thus called "wicked problems.". In this thesis, we attempted to untangle the complexity of these problems using a Bayesian network framework for learning causal models, evaluating impact using counterfactual analysis, and finally, tuning a reinforcement learning agent to make sequential decisions. In our first objective, we proposed a probabilistic causal framework. A Bayesian network provides a probabilistic representation of the interconnection of variables, from which we may extract the most significant collection of features, known as the Markov blanket. Our learned Bayesian network was initialized with evidence-based causality. The learned network provided a probabilistic representation of the interconnection of variables, from which we extracted the most influential collection of features using a Markov blanket approach. Counterfactual analysis of the discovered influences provided potential quantification for "what if" scenarios. Finally, our framework included supervised machine learning to aid advanced prediction. The complete unified approach generated through this process was applied in use cases on antimicrobial resistance, mental health, and HIV transmission. Antimicrobial resistance is one of the most pressing global problems, often called the next pandemic. To enable policy for tackling Antimicrobial Resistance, a holistic approach has to be taken that includes SDOHs, global governance, and financial and healthcare factors into account. We analyzed 6,33,820 isolates from 70 middle and

high-income countries collected between 2004-2017. Before learning a Bayesian network for knowledge discovery, we integrated these AMR data with the Global Burden of Disease (GBD), Governance (WGI), and Finance data. Our study revealed that various risk factors, such as government effectiveness (50%), unsafe water and sanitation (19.87%), low fruit consumption (13.04%), and limited access to immunization and obstetric care (39%), impact the development of antibiotic resistance, albeit with different degrees of influence depending on the pathogens and antibiotics involved. Low-income countries often struggle with the problem of excessive antibiotic prescription, which is believed to contribute to the emergence of antimicrobial resistance. However, our analysis showed no significant correlation between WASH criteria (which positively correlate with a country's income) and the prevalence of antimicrobial resistance. This was unexpected, as poor sanitation and a lack of standard treatment guidelines likely led to widespread community spread of antimicrobial-resistant organisms and a positive correlation between WASH conditions and the prevalence of AMR. Additionally, our study found a positive correlation between the rotavirus vaccine coverage and the probability of resistance to meropenem. Interestingly, our analysis indicated that countries with lower access to immunization and obstetric care are more likely to have resistant bacteria when using most antibiotics. Specifically, there were notable associations between using meropenem and ceftazidime and access to healthcare. For instance, resistance to ceftazidime in E. coli and Klebsiella and to meropenem in Acinetobacter, Staphylococcus, Klebsiella, and Pseudomonas, were more prevalent in countries with higher access to healthcare. This could be because these antibiotics are commonly administered in intensive care unit settings, where they may be less administered and resisted in countries with limited healthcare access. It could also explain the rise of antimicrobial resistance in countries with low access to immunization and obstetric care. Fruit consumption was found to influence resistance to piperacillin and cefepime in Klebsiella, highlighting the potential impact of nutrition status on susceptibility to infections by drug-resistant organisms. Analysis of our study recommends developing and implementing strategies at

all levels of society to reduce the impact and spread of antimicrobial resistance. Global governance systems, financial regulators, industrial stakeholders, medical experts, and scientists must work together to tackle this problem head-on. The following use case focused on mental health during the COVID-19 pandemic. We analyzed the survey data of 17,764 adults in the United States from different age groups, genders, and socioeconomic statuses. We found that regular face-to-face interactions with friends and family, a fulfilling social life, and social stability could decrease the likelihood of mental stress by up to 5%. Furthermore, our study revealed that individuals with a prior diagnosis of any mental illness were the most vulnerable, with a chance of mental illness during the pandemic exceeding 7%. Therefore, we recommend implementing national-level policies to track their mental status and provide appropriate treatment regularly. Importantly, we observed that underscore the economic underpinnings of a collective mental health response. We found that income assistance through Social Security or Supplemental Social Security had a demonstrable effect on alleviating anxiety, improving the chances of good mental health by more than 10%. To the best of our knowledge, this provides the first scientific evidence proving the usefulness of such efforts. The effectiveness of these measures can be captured in similar modeling studies conducted in different parts of the world, with varying assistance structures during this time. Our third use case examined effective interventions to mitigate HIV transmission in female sex workers (FSWs), a marginalized group that often lacks access to quality care. Grassroots organizations face complex challenges in improving intervention efficacy. Our analysis revealed that having at least one financial investment, such as an RD, FD, or Chit Fund, was associated with a 6% increase in condom use with customers, highlighting the importance of financial stability for FSWs to negotiate safer sex practices. Additionally, depression symptoms significantly impacted an FSW's confidence to purchase condoms independently. Legal education and financial stability were identified as two significant determinants of HIV testing, with legal education increasing the likelihood of testing every six months by 16%. Our study demonstrates the poten-

tial of explainable models to help resource-limited FSWs prioritize anti-HIV interventions. Developing models for sensitive issues such as these requires transparency and professional review, and our approach offers a framework for learning and replicating interventions in any context. In our second objective on reinforcement learning-based policy, we proposed a novel pipeline for obtaining an optimal vaccine distribution that feeds forward connected Deep Reinforcement Learning models with a Contextual Bandits model, making the optimization process more robust, context-specific, and easily deployable in real-time. Numerous logistical issues (such as supply chain, legal agreements, quality control, and application, to mention a few) might need to be addressed to ensure the dissemination of healthcare resources. A robust decision support system is thus required to overcome these obstacles. Our approach used reinforcement learning to create a novel distribution policy model called VacSIM. For the model development, the output (action and rewards) of the ACKTR/DQN model is fed into the Contextual Bandit model to provide a reasonable context consisting of actions and rewards. We have pipelined an Actor-Critic using the Kronecker-Factored Trust Region (ACKTR) model and Deep Q-Networks (DQN) model in a feed-forward way. The policy is then optimized by contextual bandits using context-based COVID-19 features (susceptible population, recovery rate, death rate, and total infected cases) and demographic metrics such as the state's population, overall hospital beds, overall ICU beds, and ventilators. Finding the demographic segment that needs the vaccination the most is a difficult problem when dispensing it; in our case, we solved it using the previously mentioned context. Policymakers and researchers may use the new, open-source, and extendable solution we offer to the vaccine allocation challenge when making choices. As the reinforcement learning component of VacSIM can always be retrained to adapt to spatial and temporal variations by taking into account the most recent values of variables like recovery rate and age distribution, this feed-forward network can be modified and used for optimal allocation of various essential resources in varying scenarios and contexts. Additionally, model-driven evaluation is a major impediment. We solved this

problem in two ways. 1) Assuming the naive ground truth policy and compared to a model-driven policy, In our VacSIM study, for example, we assumed the nave policy in the state-by-state vaccine distribution scenario as follows: percent of vaccine delivered to a state = percent of infected people in that state. Then, we compared the Nave vaccine distribution strategy to the VacSIM policy. 2) Causal structural preservation: In this part, we built causal graphical models to examine the influence of vaccination distribution on indicators such as susceptibility, recovery, infection, and death in a directed network. We discovered that the vaccination distribution acted as a parent node of susceptibility, proving our hypothesis. Finally, we have used the policy-discovering and optimization models to roll out a deterrent against COVID-19. During COVID-19, researchers showed much interest in how AI might improve medical care. Several different AI models were offered as possible decision-making aids. We have also made several dashboards and tools that can be used on the web to help make decisions in real-time. Amid the ambiguity surrounding COVID-19, CovidNLP and *EvidenceFlow* developed a paradigm and an interface for clinicians, academics, and policymakers to extract pertinent information. *CoVision* offered visualization of the COVID-19 infection trajectory, while *VacSIM* provided a solution and decision assistance to distribute COVID-19 vaccines. Finally, The models presented in this thesis are not externally validated, but developed frameworks can extend beyond the specific applications we explored (Youssef *et al.* (2023)). By providing these innovative modeling frameworks, we aim to advance data-driven research and develop effective interventions across various fields.

7.2 Future Work

The World Health Organization defined social determinants of health a decade ago. However, it is very recently that the decision makers acknowledged SDOHs as the key factors of population health. Despite widespread acceptability, SDOH is seldom

included in local, global, demographic, or clinical decision-making. Although it is difficult to model the impacts of SDOH, increasing AI approaches and vast heterogeneous data sources can conceive and quantify SDOH. AI models can be utilized for real-time decision-making. We believe that the scope of the thesis is multidimensional and can be extended a future work as follows:

7.2.1 Application of Reinforcement Learning in Public Health

Making public health decisions is arduous due to the multifaceted complexity of data. The advancement of the reinforcement learning (RL) model provides a robust and explainable artificial intelligence framework for making decisions under uncertainty (Gottesman *et al.* (2019) Gandhi and Mishra (2021)). The usage of reinforcement learning in healthcare is defined as the fusion of medicine and computer science **Figure 7.1**. It is the culmination of public health decision-making and computational intelligence to best guide the interventions (Allioui *et al.* (2022)). The same implications are far-reaching, and such collaborations outweigh either effort by itself. To carry out such cross-disciplinary research, machine learning researchers and doctors actively involved in patient care and public health decision-making must come together. RL is a technically and mathematically sound solution for optimal decision-making in a variety of healthcare tasks that are complicated by noisy, multidimensional, incomplete data, nonlinear and complex dynamics, and in particular, sequential decision processes with delayed evaluation feedback. In these ongoing treatment regimens, decision-making is automated using RL. Utilizing the extensive data gathered in different settings), RL has already assisted in formulating dynamic treatment regimes (DTRs) for chronic diseases, including cancer and HIV. Despite its significant accomplishments, RL has received much less attention from academia than other research paradigms in the healthcare domains (Riachi *et al.* (2021)), such as standard machine learning, statistical learning, control-driven techniques, and deep learning, whether in computer science or medicine.

The use of reinforcement in healthcare still faces several other difficulties. The issue of adequately evaluating RL algorithms in healthcare contexts is the most important. The quantity of data required for training, the non-stationary nature of the data, and the fact that it is only partially viewable present additional difficulties.

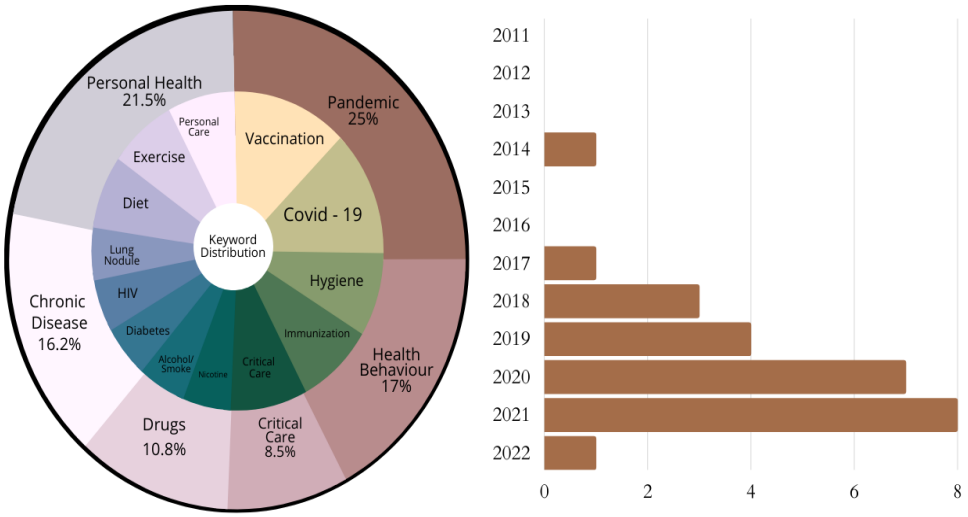


Figure 7.1: **An Overview of RL's Public Health Applications.** a)In the 27 articles reviewed, we observed the application of RL in Public Health in Intervention modeling, Mitigation Policies, Personalised Treatment, and in general complex decision making. Five of the studies that were reviewed dealt with COVID-19. b)Year-wise number of papers selected for the reviews. *Credits: Awasthi R, Devadiga A, Gupta M, Tewatia J, Sethi T Application of reinforcement learning in public health: a scoping review of literature JMIR Preprints. 24/07/2022:41381*

7.2.2 Plug and play Resource Allocation Platform for Healthcare

In chapter 5, To distribute the COVID-19 vaccine, we have introduced the VacSIM architecture. In addition to vaccines, there are several healthcare resources, including nurses, doctors, hospital beds, oxygen, and ventilators, that need to be allocated in the best possible methods(Cookson *et al.* (2008) Lane *et al.* (2017) Munthe *et al.* (2021) Kluge (2007) Yip (2021) (Ordu *et al.*, 2021)). The resources should be given to the most vulnerable people first. We propose that any general healthcare resource allocation can

benefit from a context-driven approach by using the contextual model we employed in VacSIM. We need the Plug n Play platform so that users may utilize the models more generally and understandably since stakeholders find it difficult to interpret the model parameters and hyper-parameters.

7.2.3 Integration of Social Media Data to Model Social Determinants of Health

Social media users have surpassed 1 billion in the previous decade (Knoke and Yang (2019)), making it one of the fastest-growing big data sources. In addition, the number of individuals in need of healthcare continues to rise in every culture. This sort of ever-growing, social media subscriber-based platform may be of a major utility in enhancing healthcare delivery to society via accessibility, communication, and engagement between health practitioners and patients. However, modeling health policy using social media data is inadequate or nonexistent. Consequently, social media with SDOH will serve as a crucial informational resource for intervention development in the future.

REFERENCES

1. **Abadie, A.**, **A. Diamond**, and **J. Hainmueller** (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, **105**(490), 493–505.

2. **Abou-Ismail, A.** (2020). Compartmental models of the covid-19 pandemic for physicians and physician-scientists. *SN comprehensive clinical medicine*, **2**(7), 852–858.

3. **Allioui, H.**, **M. A. Mohammed**, **N. Benameur**, **B. Al-Khateeb**, **K. H. Abdulkareem**, **B. Garcia-Zapirain**, **R. Damaševičius**, and **R. Maskeliūnas** (2022). A multi-agent deep reinforcement learning approach for enhancement of covid-19 ct image segmentation. *Journal of personalized medicine*, **12**(2), 309.

4. **Arango, M.** and **L. Pelov** (2020). Covid-19 pandemic cyclic lockdown optimization using reinforcement learning. *arXiv preprint arXiv:2009.04647*.

5. **Arias, D.**, **S. Saxena**, and **S. Verguet** (2022). Quantifying the global burden of mental disorders and their economic value. *EClinicalMedicine*, **54**, 101675.

6. **Athey, S.** and **G. W. Imbens** (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic perspectives*, **31**(2), 3–32.

7. **Avorn, J.** and **A. Kesselheim** (2020). Regulatory decision-making on covid-19 vaccines during a public health emergency. *Jama*, **324**(13), 1284–1285.

8. **Badham, J.**, **E. Chattoe-Brown**, **N. Gilbert**, **Z. Chalabi**, **F. Kee**, and **R. F. Hunter** (2018). Developing agent-based models of complex health behaviour. *Health & place*, **54**, 170–177.

9. **Bambra, C.**, **R. Riordan**, **J. Ford**, and **F. Matthews** (2020). The covid-19 pandemic and health inequalities. *J Epidemiol Community Health*, **74**(11), 964–968.

10. **Barten, F.**, **D. Mitlin**, **C. Mulholland**, **A. Hardoy**, and **R. Stern** (2007). Integrated approaches to address the social determinants of health for reducing health inequity. *Journal of Urban Health*, **84**(1), 164–173.

11. **Beautrais, A. L.** (2000). Risk factors for suicide and attempted suicide among young people. *Australian & New Zealand Journal of Psychiatry*, **34**(3), 420–436.

12. **Biddle, S. J.**, **S. Ciaccioni**, **G. Thomas**, and **I. Vergeer** (2019). Physical activity and mental health in children and adolescents: An updated review of reviews and an analysis of causality. *Psychology of Sport and Exercise*, **42**, 146–155.

13. **Biswas, M.**, **M. S. Kaiser**, **M. Mahmud**, **S. Al Mamun**, **M. Hossain**, **M. A. Rahman**, *et al.*, An xai based autism detection: The context behind the detection. *In International Conference on Brain Informatics*. Springer, 2021.

14. **Brådvik, L.** (2018). Suicide risk and mental disorders. *International journal of environmental research and public health*, **15**(9), 2028.

15. **Brauer, F.**, Compartmental models in epidemiology. *In Mathematical epidemiology*. Springer, 2008, 19–79.

16. **Braveman, P.**, **S. Egerter**, **D. R. Williams**, *et al.* (2011). The social determinants of health: coming of age. *Annual review of public health*, **32**(1), 381–398.

17. **Braveman, P.** and **L. Gottlieb** (2014). The social determinants of health: it's time to consider the causes of the causes. *Public health reports*, **129**(1_suppl2), 19–31.

18. **Bukhari, Q.**, **Y. Jameel**, **J. M. Massaro**, **R. B. D'Agostino, Sr**, and **S. Khan** (2020). Periodic Oscillations in Daily Reported Infections and Deaths for Coronavirus Disease 2019. *JAMA Network Open*, **3**(8), e2017521. ISSN 2574-3805.

19. **Campbell, D. T.** and **J. C. Stanley**, *Experimental and quasi-experimental designs for research*. Ravenio books, 2015.

20. **Carey, G.** and **B. Crammond** (2015). Action on the social determinants of health: views from inside the policy process. *Social science & medicine*, **128**, 134–141.

21. **Chandler, C. I.** (2019). Current accounts of antimicrobial resistance: stabilisation, individualisation and antibiotics as infrastructure. *Palgrave communications*, **5**(1), 1–13.

22. **Chelouati, M.**, **M. S. Jha**, **M. Galeotta**, and **D. Theilliol**, Remaining useful life prediction for liquid propulsion rocket engine combustion chamber. *In 2021 5th International Conference on Control and Fault-Tolerant Systems (SysTol)*. IEEE, 2021.

23. **Chen, L.**, *Deep Learning and Practice with MindSpore*. Springer Nature, 2021.

24. **Chen, M.**, **V. Chernozhukov**, **I. Fernandez-Val**, and **B. Melly** (2017). *Counterfactual: Estimation and Inference Methods for Counterfactual Analysis*. R package version 1.1.

25. **Chen, Q.**, **A. Allot**, and **Z. Lu** (2021). Litcovid: an open database of covid-19 literature. *Nucleic acids research*, **49**(D1), D1534–D1540.

26. **Chen, Q.**, **M. Liang**, **Y. Li**, **J. Guo**, **D. Fei**, **L. Wang**, **L. He**, **C. Sheng**, **Y. Cai**, **X. Li**, *et al.* (2020). Mental health care for medical staff in china during the covid-19 outbreak. *The Lancet Psychiatry*, **7**(4), e15–e16.

27. **Cheng, S. H.**, **M. C. McKinnon**, **Y. J. Masuda**, **R. Garside**, **K. W. Jones**, **D. C. Miller**, **A. S. Pullin**, **W. J. Sutherland**, **C. Augustin**, **D. A. Gill**, *et al.* (2020). Strengthen causal models for better conservation outcomes for human well-being. *Plos one*, **15**(3), e0230495.

28. **Chua, A. Q.**, **M. Verma**, **L. Y. Hsu**, and **H. Legido-Quigley** (2021). An analysis of national action plans on antimicrobial resistance in southeast asia using a governance framework approach. *The Lancet Regional Health-Western Pacific*, **7**, 100084.

29. **Collier, M.** and **H. U. Llorens** (2018). Deep Contextual Multi-armed Bandits. *arXiv:1807.09809 [cs, stat]*. ArXiv: 1807.09809.

30. **Collignon, P.**, **J. J. Beggs**, **T. R. Walsh**, **S. Gandra**, and **R. Laxminarayan** (2018). Anthropological and socioeconomic factors contributing to global antimicrobial resistance: a univariate and multivariable analysis. *The Lancet Planetary Health*, **2**(9), e398–e405.

31. **Cookson, R.**, **C. McCabe**, and **A. Tsuchiya** (2008). Public healthcare resource allocation and the rule of rescue. *Journal of medical ethics*, **34**(7), 540–544.

32. **Cotta, R. M.**, **C. P. Naveira-Cotta**, and **P. Magal** (2020). Mathematical parameters of the covid-19 epidemic in brazil and evaluation of the impact of different public health measures. *Biology*, **9**(8), 220.

33. **Das, P.** and **R. Horton** (2015). Bringing sex workers to the centre of the hiv response. *Lancet (London, England)*, **385**(9962), 3–4.

34. **Davies, J.** and **D. Davies** (2010). Origins and evolution of antibiotic resistance. *Microbiology and molecular biology reviews*, **74**(3), 417–433.

35. **De Camargo, C.** (2022). 'it's tough shit, basically, that you're all gonna get it': Uk virus testing and police officer anxieties of contracting covid-19. *Policing and Society*, **32**(1), 35–51.

36. **Deaton, A.** and **N. Cartwright** (2018). Understanding and misunderstanding randomized controlled trials. *Social science & medicine*, **210**, 2–21.

37. **Deo, S.**, **S. Manurkar**, **S. Krishnan**, and **C. Franz** (2020). COVID-19 Vaccine: Development, Access and Distribution in the Indian Context. *ORF Issue Brief*, (378), 16.

38. **Dimick, J. B.** and **A. M. Ryan** (2014). Methods for evaluating changes in health care policy: the difference-in-differences approach. *Jama*, **312**(22), 2401–2402.

39. **Exworthy, M.** (2008). Policy to tackle the social determinants of health: using conceptual models to understand the policy process. *Health policy and planning*, **23**(5), 318–327.

40. **Fink, D. S.**, **K. M. Keyes**, and **M. Cerdá** (2016). Social determinants of population health: a systems sciences approach. *Current epidemiology reports*, **3**(1), 98–105.

41. **Foster, D.**, **C. Mcgregor**, and **S. El-masri**, A Survey of Agent-Based Intelligent Decision Support Systems to Support Clinical. *In Management and Research", 1st Intl. Workshop on Multi-Agent Systems for Medicine, Computational Biology, and Bioinformatics*.

42. **Fougère, D.** and **N. Jacquemet**, Policy evaluation using causal inference methods. *In Handbook of Research Methods and Applications in Empirical Microeconomics*. Edward Elgar Publishing, 2021, 294–324.

43. **Gail, M. H.**, **J. P. Costantino**, **D. Pee**, **M. Bondy**, **L. Newman**, **M. Selvan**, **G. L. Anderson**, **K. E. Malone**, **P. A. Marchbanks**, **W. McCaskill-Stevens**, *et al.* (2007). Projecting individualized absolute invasive breast cancer risk in african american women. *Journal of the National Cancer Institute*, **99**(23), 1782–1792.

44. **Galanis, G.** and **A. Hanieh** (2021). Incorporating social determinants of health into modelling of covid-19 and other infectious diseases: a baseline socio-economic compartmental model. *Social science & medicine*, **274**, 113794.

45. **Gandhi, N.** and **S. Mishra**, Applications of reinforcement learning for medical decision making. *In RTA-CSIT*. 2021.

46. **Gao, J.**, **P. Zheng**, **Y. Jia**, **H. Chen**, **Y. Mao**, **S. Chen**, **Y. Wang**, **H. Fu**, and **J. Dai** (2020). Mental health problems and social media exposure during covid-19 outbreak. *Plos one*, **15**(4), e0231924.

47. **Gentleman, R.** and **V. J. Carey**, Unsupervised machine learning. *In Bioconductor case studies*. Springer, 2008, 137–157.

48. **Goldsztejn, U.**, **D. Schwartzman**, and **A. Nehorai** (2020). Public policy and economic dynamics of covid-19 spread: a mathematical modeling study. *Plos one*, **15**(12), e0244174.

49. **Gottesman, O.**, **F. Johansson**, **M. Komorowski**, **A. Faisal**, **D. Sontag**, **F. Doshi-Velez**, and **L. A. Celi** (2019). Guidelines for reinforcement learning in healthcare. *Nature medicine*, **25**(1), 16–18.

50. **Greenwood, B.** (2014). The contribution of vaccination to global health: past, present and future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **369**(1645). ISSN 0962-8436.

51. **Hassouni, A. e.**, **M. Hoogendoorn**, **M. v. Otterlo**, and **E. Barbaro**, Personalization of health interventions using cluster-based reinforcement learning. *In International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 2018.

52. **Hennessy, M. B.**, **B. McCowan**, **J. Jiang**, and **J. P. Capitanio** (2014). Depressive-like behavioral response of adult male rhesus monkeys during routine animal husbandry procedure. *Frontiers in Behavioral Neuroscience*, **8**, 309.

53. **Hester, T.**, **M. Vecerik**, **O. Pietquin**, **M. Lanctot**, **T. Schaul**, **B. Piot**, **D. Horgan**, **J. Quan**, **A. Sendonaris**, **I. Osband**, **G. Dulac-Arnold**, **J. Agapiou**, **J. Leibo**, and **A. Gruslys** (2018). Deep Q-learning From Demonstrations. *Proceedings of the AAAI Conference on Artificial Intelligence*, **32**(1). ISSN 2374-3468. Number: 1.

54. **Ho, B.-J.**, **B. Balaji**, **M. Koseoglu**, and **M. Srivastava**, Nurture: notifying users at the right time using reinforcement learning. *In Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. 2018.

55. **Hogan, J. W.**, **N. Galai**, and **W. W. Davis** (2021). Modeling the impact of social determinants of health on hiv. *AIDS and Behavior*, **25**(2), 215–224.

56. **Holmes, E.** (2020). O connor rc, perry vh, tracey i, wessely s, arseneault l, et al. multidisciplinary research priorities for the covid-19 pandemic: a call for action for mental health science. *Lancet Psychiatry*, **7**(6), 547–60.

57. **Imbens, G. W.** (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics*, **86**(1), 4–29.

58. **Irvin, J. A.**, **A. A. Kondrich**, **M. Ko**, **P. Rajpurkar**, **B. Haghgoo**, **B. E. Landon**, **R. L. Phillips**, **S. Petterson**, **A. Y. Ng**, and **S. Basu** (2020). Incorporating machine learning and social determinants of health indicators into prospective risk adjustment for health plan payments. *BMC Public Health*, **20**(1), 1–10.

59. **Islam, M. R.**, **T. Oraby**, **A. McCombs**, **M. M. Chowdhury**, **M. Al-Mamun**, **M. G. Tyshenko**, and **C. Kadelka** (2021). Evaluation of the united states covid-19 vaccine allocation strategy. *Plos one*, **16**(11), e0259700.

60. **Iwata, K.** (2016). Extending the peak bandwidth of parameters for softmax selection in reinforcement learning. *IEEE Transactions on Neural Networks and Learning Systems*, **28**(8), 1865–1877.

61. **Janssen, M. A.** and **E. Ostrom** (2006). Empirically based, agent-based models. *Ecology and society*, **11**(2).

62. **Javad, M. O. M.**, **S. O. Agboola**, **K. Jethwani**, **A. Zeid**, **S. Kamarthi**, *et al.* (2019). A reinforcement learning–based method for management of type 1 diabetes: exploratory study. *JMIR diabetes*, **4**(3), e12905.

63. **Jewell, N. P.**, **J. A. Lewnard**, and **B. L. Jewell** (2020). Predictive mathematical models of the covid-19 pandemic: underlying principles and value of projections. *Jama*, **323**(19), 1893–1894.

64. **Kaelbling, L. P.**, **M. L. Littman**, and **A. W. Moore** (1996). Reinforcement learning: A survey. *Journal of artificial intelligence research*, **4**, 237–285.

65. **Kahn, B.**, **L. Brown**, **W. Foege**, **H. Gayle**, **E. National Academies of Sciences**, **Medicine**, *et al.*, *A framework for equitable allocation of COVID-19 vaccine*. National Academies Press (US), 2020.

66. **Kallakuri, S.**, **S. Devarapalli**, **A. P. Tripathi**, **A. Patel**, and **P. K. Maulik** (2018). Common mental disorders and risk factors in rural india: baseline data from the smart mental health project. *BJPsych open*, **4**(4), 192–198.

67. **Kerr, C. C.**, **R. M. Stuart**, **D. Mistry**, **R. G. Abeysuriya**, **K. Rosenfeld**, **G. R. Hart**, **R. C. Núñez**, **J. A. Cohen**, **P. Selvaraj**, **B. Hagedorn**, *et al.* (2021). Covasim: an agent-based model of covid-19 dynamics and interventions. *PLOS Computational Biology*, **17**(7), e1009149.

68. **Khamsi, R.** *et al.* (2020). If a coronavirus vaccine arrives, can the world make enough. *Nature*, **580**(7805), 578–580.

69. **Kim, H.** and **M. H. Kim** (2015). Package 'discretization'. *CRAN*.

70. **Kino, S.**, **Y.-T. Hsu**, **K. Shiba**, **Y.-S. Chien**, **C. Mita**, **I. Kawachi**, and **A. Daoud** (2021). A scoping review on the use of machine learning in research on social determinants of health: Trends and research prospects. *SSM-population health*, **15**, 100836.

71. **Kluge, E.-H. W.** (2007). Resource allocation in healthcare: implications of models of medicine as a profession. *Medscape General Medicine*, **9**(1), 57.

72. **Knoke, D.** and **S. Yang**, *Social network analysis*. SAGE publications, 2019.

73. **Kohl, J.**, **A. E. Autry**, and **C. Dulac** (2017). The neurobiology of parenting: A neural circuit perspective. *Bioessays*, **39**(1), 1–11.

74. **Koller, D.** and **N. Friedman**, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

75. **Komorowski, M.**, **L. A. Celi**, **O. Badawi**, **A. C. Gordon**, and **A. A. Faisal** (2018). The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, **24**(11), 1716–1720. ISSN 1546-170X.

76. **Kompella, V.**, **R. Capobianco**, **S. Jong**, **J. Browne**, **S. Fox**, **L. Meyers**, **P. Wurman**, and **P. Stone** (2020). Reinforcement learning for optimization of covid-19 mitigation policies. *arXiv preprint arXiv:2010.10560*.

77. **Kordonouri, O.** and **M. C. Riddell** (2019). Use of apps for physical activity in type 1 diabetes: current status and requirements for future development. *Therapeutic advances in endocrinology and metabolism*, **10**, 2042018819839298.

78. **Kwak, G. H.**, **L. Ling**, and **P. Hui** (2021*a*). Deep reinforcement learning approaches for global public health strategies for COVID-19 pandemic. *Plos one*, **16**(5), e0251550. ISSN 1932-6203. Publisher: Public Library of Science.

79. **Kwak, G. H.**, **L. Ling**, and **P. Hui** (2021*b*). Deep reinforcement learning approaches for global public health strategies for covid-19 pandemic. *Plos one*, **16**(5), e0251550.

80. **Lane, H.**, **M. Sarkies**, **J. Martin**, and **T. Haines** (2017). Equity in healthcare resource allocation decision making: a systematic review. *Social science & medicine*, **175**, 11–27.

81. **Laxminarayan, R.**, **A. Duse**, **C. Wattal**, **A. K. Zaidi**, **H. F. Wertheim**, **N. Sumpradit**, **E. Vlieghe**, **G. L. Hara**, **I. M. Gould**, **H. Goossens**, *et al.* (2013). Antibiotic resistance—the need for global solutions. *The Lancet infectious diseases*, **13**(12), 1057–1098.

82. **Leng, Y.**, **Y. Zhai**, **S. Sun**, **Y. Wu**, **J. Selzer**, **S. Strover**, **H. Zhang**, **A. Chen**, and **Y. Ding** (2021). Misinformation during the covid-19 outbreak in china: cultural, social and political entanglements. *IEEE Transactions on Big Data*, **7**(1), 69–80.

83. **Lewis, E. F.** (2022). Machine learning and social determinants of health—an opportunity to move beyond race for inpatient risk prediction in patients with heart failure. *JAMA cardiology*.

84. **Li, M. Y.** and **J. S. Muldowney** (1995). Global stability for the SEIR model in epidemiology. *Mathematical Biosciences*, **125**(2), 155–164. ISSN 0025-5564.

85. **Liu, S.**, **K. C. See**, **K. Y. Ngiam**, **L. A. Celi**, **X. Sun**, and **M. Feng** (2020). Reinforcement Learning for Clinical Decision Support in Critical Care: Comprehensive Review. *Journal of Medical Internet Research*, **22**(7). ISSN 1439-4456.

86. **Liu, Y.**, **N. Qiao**, and **Y. Altinel** (2021). Reinforcement learning in neurocritical and neurosurgical care: principles and possible applications. *Computational and Mathematical Methods in Medicine*, **2021**.

87. **Ly, A.**, **M. Marsman**, **J. Verhagen**, **R. Grasman**, and **E.-J. Wagenmakers** (2017). A Tutorial on Fisher Information. *arXiv:1705.01064 [math, stat]*. ArXiv: 1705.01064.

88. **Maglio, P. P.** and **P. L. Mabry** (2011). Agent-based models and systems science approaches to public health. *American journal of preventive medicine*, **40**(3), 392–394.

89. **Mahapatra, B.**, **M. Walia**, **S. K. Patel**, **M. Battala**, **S. Mukherjee**, **P. Patel**, **B. Subramanium**, **Y. Atmavilas**, and **N. Saggurti** (2020). Sustaining consistent condom use among female sex workers by addressing their vulnerabilities and strengthening community-led organizations in india. *Plos one*, **15**(7), e0235094.

90. **Maheshwari, S.**, **K. Pahwa**, and **T. Sethi** (2021). Wiser: An end-to-end structure learning and deployment framework for causal graphical models. *arXiv preprint arXiv:2108.07046*.

91. **Maragakis, L. L.**, **E. N. Perencevich**, and **S. E. Cosgrove** (2008). Clinical and economic burden of antimicrobial resistance. *Expert review of anti-infective therapy*, **6**(5), 751–763.

92. **Marmot, M.** and **R. Wilkinson**, *Social determinants of health*. Oup Oxford, 2005.

93. **Martens, J.** (2020). New insights and perspectives on the natural gradient method. *arXiv:1412.1193 [cs, stat]*. ArXiv: 1412.1193.

94. **Martinho, D.**, **J. Carneiro**, **J. Neves**, **P. Novais**, **J. Corchado**, and **G. Marreiros**, A reinforcement learning approach to improve user achievement of health-related goals. *In EPIA Conference on Artificial Intelligence*. Springer, 2021.

95. **Martins, D. C.**, **O. Babajide**, **N. Maani**, **S. M. Abdalla**, **E. J. Gómez**, **M. J. Pongsiri**, **S. Tlou**, **G. M. Leung**, **G. C. Benjamin**, **E. Goosby**, *et al.* (2021). Integrating social determinants in decision-making processes for health: insights from conceptual frameworks—the 3-d commission. *Journal of Urban Health*, **98**(1), 51–59.

96. **Medina, M.-j.**, **H. Legido-Quigley**, and **L. Y. Hsu** (2020). Antimicrobial resistance in one health. *Global Health Security: Recognizing Vulnerabilities, Creating Opportunities*, 209–229.

97. **Mnih, V.**, **A. P. Badia**, **M. Mirza**, **A. Graves**, **T. P. Lillicrap**, **T. Harley**, **D. Silver**, and **K. Kavukcuoglu** (2016). Asynchronous Methods for Deep Reinforcement Learning. *arXiv:1602.01783 [cs]*. ArXiv: 1602.01783.

98. **Mooney, S. J.** and **V. Pejaver** (2018). Big data in public health: terminology, machine learning, and privacy. *Annual review of public health*, **39**, 95.

99. **Munthe, C.**, **D. Fumagalli**, and **E. Malmqvist** (2021). Sustainability principle for the ethics of healthcare resource allocation. *Journal of Medical Ethics*, **47**(2), 90–97.

100. **National Academies of Sciences, E.**, **Medicine**, *et al.*, *Framework for equitable allocation of COVID-19 vaccine*. National Academies Press, 2020.

101. **Nawaz, A.**, **X. Su**, **S. Iqbal**, **H. Zahoor**, **A. Asad**, **S. Asghar**, **F. Basit**, **M. Q. Barkat**, **A. Souhail**, and **S. A. Raheel Shah** (). Validating a phenomenological mathematical

model for public health and safety interventions influencing the evolutionary stages of recent outbreak for long-term and short-term domains in pakistan. *Complexity*, **2020**.

102. **Naylor, N. R.**, **R. Atun**, **N. Zhu**, **K. Kulasabanathan**, **S. Silva**, **A. Chatterjee**, **G. M. Knight**, and **J. V. Robotham** (2018). Estimating the burden of antimicrobial resistance: a systematic literature review. *Antimicrobial Resistance & Infection Control*, **7**(1), 1–17.

103. **Ordu, M.**, **E. Demir**, **C. Tofallis**, and **M. M. Gunal** (2021). A novel healthcare resource allocation decision support tool: A forecasting-simulation-optimization approach. *Journal of the operational research society*, **72**(3), 485–500.

104. **Organization, W. H.** *et al.* (2008). Social determinants of health. Technical report, WHO Regional Office for South-East Asia.

105. **Paremoer, L.**, **S. Nandi**, **H. Serag**, and **F. Baum** (2021). Covid-19 pandemic and the social determinants of health. *bmj*, **372**.

106. **Pearl, J.** (1998). Graphical models for probabilistic and causal reasoning. *Quantified representation of uncertainty and imprecision*, 367–389.

107. **Pearl, J.**, *Causality*. Cambridge university press, 2009.

108. **Pearl, J.**, *Probabilistic reasoning in intelligent systems*, volume 88. Elsevier, 2014.

109. **Pearl, J.** and **D. Mackenzie**, *The book of why: the new science of cause and effect*. Basic books, 2018.

110. **Pedregosa, F.**, **G. Varoquaux**, **A. Gramfort**, **V. Michel**, **B. Thirion**, **O. Grisel**, **M. Blondel**, **P. Prettenhofer**, **R. Weiss**, **V. Dubourg**, *et al.* (2011). Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, **12**, 2825–2830.

111. **Pereira-Sanchez, V.**, **F. Adiukwu**, **S. El Hayek**, **D. G. Bytyçi**, **J. M. Gonzalez-Diaz**, **G. K. Kundadak**, **A. Larnaout**, **M. Nofal**, **L. Orsolini**, **R. Ramalho**, *et al.* (2020).

Covid-19 effect on mental health: patients and workforce. *The Lancet Psychiatry*, **7**(6), e29–e30.

112. **Petersen, M. L.** and **M. J. van der Laan** (2014). Causal models and learning from data: integrating causal modeling and statistical estimation. *Epidemiology (Cambridge, Mass.)*, **25**(3), 418.

113. **Pfefferbaum, B.** and **C. S. North** (2020). Mental health and the covid-19 pandemic. *New England Journal of Medicine*, **383**(6), 510–512.

114. **Polack, F. P.**, **S. J. Thomas**, **N. Kitchin**, **J. Absalon**, **A. Gurtman**, **S. Lockhart**, **J. L. Perez**, **G. Perez Marc**, **E. D. Moreira**, **C. Zerbini**, **R. Bailey**, **K. A. Swanson**, **S. Roychoudhury**, **K. Koury**, **P. Li**, **W. V. Kalina**, **D. Cooper**, **R. W. Frenck**, **L. L. Hammitt**, **O. Tureci**, **H. Nell**, **A. Schaefer**, **S. Unal**, **D. B. Tresnan**, **S. Mather**, **P. R. Dormitzer**, **U. Sahin**, **K. U. Jansen**, and **W. C. Gruber** (2020). Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine. *New England Journal of Medicine*, **383**(27), 2603–2615.

115. **Quante, A. S.**, **A. S. Whittemore**, **T. Shriver**, **K. Strauch**, and **M. B. Terry** (2012). Breast cancer risk assessment across the risk continuum: genetic and nongenetic risk factors contributing to differential model performance. *Breast Cancer Research*, **14**(6), 1–12.

116. **Rabel, M.**, **M. Laxy**, **B. Thorand**, **A. Peters**, **L. Schwettmann**, and **F. Mess** (2019). Clustering of health-related behavior patterns and demographics. results from the population-based kora s4/f4 cohort study. *Frontiers in public health*, **6**, 387.

117. **Raghu, A.**, **M. Komorowski**, **L. A. Celi**, **P. Szolovits**, and **M. Ghassemi** (). Continuous State-Space Models for Optimal Sepsis Treatment: a Deep Reinforcement Learning Approach, 17.

118. **Rahmandad, H.**, **R. Xu**, and **N. Ghaffarzadegan** (2022). Enhancing long-term forecasting: Learning from covid-19 models. *PLoS computational biology*, **18**(5), e1010100.

119. **Ramachandran, M.** (1997). A counterfactual analysis of causation. *Mind*, **106**(422), 263–277.

120. **Ranapurwala, S. I.**, **J. E. Cavanaugh**, **T. Young**, **H. Wu**, **C. Peek-Asa**, and **M. R. Ramirez** (2019). Public health application of predictive modeling: an example from farm vehicle crashes. *Injury epidemiology*, **6**(1), 1–11.

121. **Rella, S. A.**, **Y. A. Kulikova**, **E. T. Dermitzakis**, and **F. A. Kondrashov** (2021). Rates of sars-cov-2 transmission and vaccination impact the fate of vaccine-resistant strains. *Scientific Reports*, **11**(1), 1–10.

122. **Riachi, E.**, **M. Mamdani**, **M. Fralick**, and **F. Rudzicz** (2021). Challenges for reinforcement learning in healthcare. *arXiv preprint arXiv:2103.05612*.

123. **Riquelme, C.**, **G. Tucker**, and **J. Snoek** (2018). Deep Bayesian Bandits Showdown: An Empirical Comparison of Bayesian Deep Networks for Thompson Sampling. *arXiv:1802.09127 [cs, stat]*. ArXiv: 1802.09127.

124. **Roda, W. C.**, **M. B. Varughese**, **D. Han**, and **M. Y. Li** (2020). Why is it difficult to accurately predict the covid-19 epidemic? *Infectious disease modelling*, **5**, 271–281.

125. **Sachs, J.**, **C. Kroll**, **G. Lafortune**, **G. Fuller**, and **F. Woelm**, *Sustainable development report 2022*. Cambridge University Press, 2022.

126. **Saria, S.** (2018). Individualized sepsis treatment using reinforcement learning. *Nature Medicine*, **24**(11), 1641–1642. ISSN 1546-170X. Number: 11 Publisher: Nature Publishing Group.

127. **Schroeder, S. A.** (2007). We can do better—improving the health of the american people. *New England journal of medicine*, **357**(12), 1221–1228.

128. **Schulman, J.**, **P. Moritz**, **S. Levine**, **M. Jordan**, and **P. Abbeel** (2018). High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv:1506.02438 [cs]*. ArXiv: 1506.02438.

129. **Seabrook, E. M.**, **M. L. Kern**, and **N. S. Rickard** (2016). Social networking sites, depression, and anxiety: a systematic review. *JMIR mental health*, **3**(4), e5842.

130. **Segar, M. W.**, **J. L. Hall**, **P. S. Jhund**, **T. M. Powell-Wiley**, **A. A. Morris**, **D. Kao**, **G. C. Fonarow**, **R. Hernandez**, **N. E. Ibrahim**, **C. Rutan**, *et al.* (2022). Machine learning–based models incorporating social determinants of health vs traditional models for predicting in-hospital mortality in patients with heart failure. *JAMA cardiology*, **7**(8), 844–854.

131. **Sethi, T.**, **A. Mittal**, **S. Maheshwari**, and **S. Chugh**, Learning to address health inequality in the united states with a bayesian decision network. *In Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33. 2019.

132. **Singh, A.**, **N. Thakur**, and **A. Sharma**, A review of supervised machine learning algorithms. *In 2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom)*. Ieee, 2016.

133. **Smith, P. C.**, **A. Anell**, **R. Busse**, **L. Crivelli**, **J. Healy**, **A. K. Lindahl**, **G. Westert**, and **T. Kene** (2012). Leadership and governance in seven developed health systems. *Health policy*, **106**(1), 37–49.

134. **Stekhoven, D. J.** and **P. Bühlmann** (2012). Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, **28**(1), 112–118.

135. **Stone, R. L.** (2020). How an ohio-based physician organization overcame internal hurdles and launched a telehealth service as covid-19 shutdowns loomed. *NEJM Catalyst Innovations in Care Delivery*, **1**(5).

136. **Stuart, E. A.** (2010). Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, **25**(1), 1.

137. **Sutton, R. S.** and **A. G. Barto**, *Reinforcement learning: An introduction*. MIT press, 2018.

138. **Tagliabue, F.**, **L. Galassi**, and **P. Mariani** (2020). The "pandemic" of disinformation in covid-19. *SN comprehensive clinical medicine*, **2**(9), 1287–1289.

139. **Tracy, M.**, **M. Cerdá**, and **K. M. Keyes** (2018). Agent-based modeling in public health: current applications and future directions. *Annual review of public health*, **39**, 77.

140. **Truszkowska, A.**, **M. Thakore**, **L. Zino**, **S. Butail**, **E. Caroppo**, **Z.-P. Jiang**, **A. Rizzo**, and **M. Porfiri** (2021). Designing the safe reopening of us towns through high-resolution agent-based modeling. *Advanced theory and simulations*, **4**(9), 2100157.

141. **van Hasselt, H.**, **A. Guez**, and **D. Silver** (2015). Deep Reinforcement Learning with Double Q-learning. *arXiv:1509.06461 [cs]*. ArXiv: 1509.06461.

142. **Venkataramani, A. S.**, **J. Bor**, and **A. B. Jena** (2016). Regression discontinuity designs in healthcare research. *bmj*, **352**.

143. **Watson, G. L.**, **D. Xiong**, **L. Zhang**, **J. A. Zoller**, **J. Shamshoian**, **P. Sundin**, **T. Bufford**, **A. W. Rimoin**, **M. A. Suchard**, and **C. M. Ramirez** (2021). Pandemic velocity: Forecasting covid-19 in the us with a machine learning & bayesian time series compartmental model. *PLoS computational biology*, **17**(3), e1008837.

144. **White, D. B.** and **D. C. Angus** (2020). A proposed lottery system to allocate scarce covid-19 medications: promoting fairness and generating knowledge. *Jama*, **324**(4), 329–330.

145. **Wilder, B.**, **L. Onasch-Vera**, **G. Diguiseppi**, **R. Petering**, **C. Hill**, **A. Yadav**, **E. Rice**, and **M. Tambe**, Clinical trial of an ai-augmented intervention for hiv prevention in youth experiencing homelessness. *In Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35. 2021.

146. **Williams, R. J.** (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, **8**(3), 229–256. ISSN 1573-0565.

147. **Wilson, P. W.**, **R. B. D'Agostino**, **D. Levy**, **A. M. Belanger**, **H. Silbershatz**, and **W. B. Kannel** (1998). Prediction of coronary heart disease using risk factor categories. *Circulation*, **97**(18), 1837–1847.

148. **Wouters, O. J.**, **K. C. Shadlen**, **M. Salcher-Konrad**, **A. J. Pollard**, **H. J. Larson**, **Y. Teerawattananon**, and **M. Jit** (2021). Challenges in ensuring global access to covid-19 vaccines: production, affordability, allocation, and deployment. *The Lancet*, **397**(10278), 1023–1034.

149. **Wu, Y.**, **E. Mansimov**, **S. Liao**, **R. Grosse**, and **J. Ba** (2017). Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation. *arXiv:1708.05144 [cs]*. ArXiv: 1708.05144.

150. **Yip, J. Y.-C.** (2021). Healthcare resource allocation in the covid-19 pandemic: Ethical considerations from the perspective of distributive justice within public health. *Public Health in Practice*, **2**, 100111.

151. **Youssef, A.**, **M. Pencina**, **A. Thakur**, **T. Zhu**, **D. Clifton**, and **N. H. Shah** (2023). All models are local: time to replace external validation with recurrent local validation. *arXiv preprint arXiv:2305.03219*.

152. **Zhang, L.**, **S. F. Rashid**, and **G. Leung** (2022). Social determinants, data science, and decision making: The 3-d approach to achieving health equity in asia. *Health Data Science*, **2022**.

APPENDIX A

Variables Description

1. **Access to Immunisation and Obstetric care** A measure of health system access estimated using a principal component analysis of antenatal clinics, DTP3 immunization, measles immunization, hospital beds, in-facility delivery, and skilled birth attendance.

2. **Mean Temp Long Term** Population-weighted mean temperature.

3. **Fruits G ADJ** Fruits g/p/d avail adjusted.

4. **Dengue Outbreak** A binary indicator for each country and year of dengue outbreaks; 0 = no outbreak, 1= outbreak.

5. **Sev Scaler Diarrhea** All Risk Factors SEV(Summary exposure value) scalar for cause Diarrhea.

6. **Rule of Law** Rule of law captures perceptions of the extent to which agents have confidence in and abide by the rules of society, and in particular the quality of contract enforcement, property rights, the police, and the courts, as well as the likelihood of crime and violence. This table lists the individual variables from each data source used to construct this measure in the Worldwide Governance Indicators.

7. **Government Effectiveness** Government effectiveness captures perceptions of the quality of public services, the quality of the civil service and the degree of its independence from political pressures, the quality of policy formulation and implementation, and the credibility of the government's commitment to such policies. This table lists the individual variables from each data source used to construct this measure in the Worldwide Governance Indicators.

8. **Voice and Accountability** Voice and Accountability captures perceptions of the extent to which a country's citizens are able to participate in selecting their government, as well as freedom of expression, freedom of association, and a free media. This table lists the individual variables from each data source used to construct this measure in the Worldwide Governance Indicators.

9. **GFDD EI 04** Bank's Overhead Costs to Total Assets.

10. **Financial Literacy** Financial literacy is the ability to understand and effectively use various financial skills, including personal financial management, budgeting, and investing.

11. **Combined TI services** Targeted intervention (TI) services are provided to female sex workers (FSWs) to promote safe sex practices, reduce the risk of HIV transmission, and provide access to healthcare services. TI teams provide education on HIV prevention and safe sex practices, distribute free condoms, offer

voluntary counseling and testing for HIV, conduct regular screenings for sexually transmitted infections (STIs), provide treatment for those who test positive, and help FSWs access healthcare services for a range of issues. Additionally, TI teams train peer educators who can provide support and education to other FSWs in their community.

12. **soc5a** In the past seven days, how often have you Felt nervous, anxious, or on edge
13. **soc5b** In the past seven days, how often have you Felt depressed
14. **soc2a** In the past month, how often did you talk with any of your neighbors?
15. **soc2b** During a typical month before March 1, 2020, when COVID-19 began spreading in the United States, how often did you talk with any of your neighbors?
16. **phys1a** Have you experienced Fever related symptoms in the past seven days?
17. **phys1q** Have you experienced Loss of appetite-related symptoms in the past seven days?
18. **phys3a** Has a doctor or other health care provider ever told you that you have Diabetes?
19. **phys3m** Has a doctor or other health care provider ever told you that you have Overweight or obese?
20. **phys3k** Has a doctor or other health care provider ever told you that you have Cancer?