A Framework for Machine Learning in Digital-Based Tuberculosis Treatment

Jackson A. Killian[1] , **Bryan Wilder**[1] , **Amit Sharma**[2] , **Vinod Choudhary**[3] , **Daksha Shah**[4] , **Bistra Dilkina**[1] , **Milind Tambe**[1]
[1]University of Southern California
[2]Microsoft Research India
[3]RNTCP, Mumbai
[4]Jt Exec. Health officer, MCGM Mumbai
{jakillia, bwilder, dilkina, tambe}@usc.edu, amshar@microsoft.com, drvinod06@gmail.com, dtomhbmc@rntcp.org

Abstract

Digital Adherence Technologies (DATs) are an increasingly popular method for verifying patient adherence to many medications. We analyze data from one city served by 99DOTS, a phone-call-based DAT deployed for Tuberculosis (TB) treatment in India where nearly 3 million people are afflicted with the disease each year. The data contains nearly 17,000 patients and 2.1M phone calls. We lay the groundwork for learning from this real-world data, including a method for avoiding the effects of unobserved interventions in training data used for machine learning. We then construct a deep learning model and show how it can be adapted and trained in different clinical scenarios to better target and improve patient care. In the real-time risk prediction setting our model could be used to proactively intervene with 21% more patients and before 76% more missed doses than current heuristic baselines. We also present a case study demonstrating how our model can be trained in an end-to-end decision focused learning setting to achieve 15% better solution quality in an example decision problem faced by health workers.

1 Introduction

The World Health Organization (WHO) reports that the lung disease tuberculosis (TB) is one of the top ten causes of death worldwide [WHO, 2018], yet in most cases it is a curable and preventable disease. The prevalence of TB is caused in part by non-adherence to medication, which results in greater risk of death, reinfection and contraction of multidrug-resistant TB [Thomas *et al.*, 2005]. To combat non-adherence, digital adherence technologies (DATs), which give patients flexible means to prove adherence, have gained popularity globally [Subbaraman *et al.*, 2018].

DATs allow patients to be "observed" consuming their medication electronically, e.g. via two-way text messaging, video capture, electronic pillboxes, or toll-free phone calls. Health workers can then view real-time patient adherence on a dashboard such as **Figure 1**. In addition to improving patient flexibility and privacy, the dashboard enables health workers to triage patients and focus their limited resources on the highest risk patients. Preliminary studies suggest that DATs can improve adherence in multiple disease settings [Haberer *et al.*, 2017; Corden *et al.*, 2016], prompting its use and evaluation for

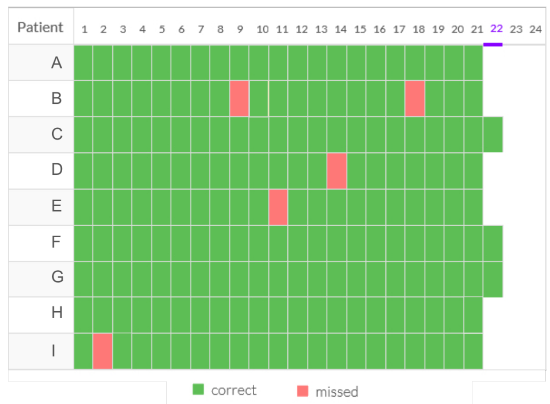

Figure 1: 99DOTS electronic adherence dashboard seen by health workers. Missed doses are marked in red while consumed doses are marked in green.

managing TB adherence [Garfein *et al.*, 2015; Liu *et al.*, 2015]. The WHO has even published a guide for the proper implementation of the technology in TB care [WHO, 2017].

In this work, we study how the wealth of longitudinal data produced by DATs can be used to help health workers better triage TB patients and deliver interventions to boost overall adherence of their patient cohort. The data we analyze comes from a partnership with the 99DOTS program [99DOTS, 2019] and the healthcare technology company Everwell [Everwell, 2019] who have implemented a DAT by which patients prove adherence through daily toll-free calls. 99DOTS operates in India where there were an estimated 2.7 million cases of TB in 2017 [WHO, 2018]; they shared data from one major city in Maharashtra (referred to as "The City.") Patients enrolled in 99DOTS in The City currently receive interventions according to the following general guidelines. If they have not taken their medication by the afternoon, they (and their health worker) receive a text message reminder. If the patient still does not take their medication by some time later, the worker will call the patient directly. Finally, if a patient simply does not respond to these previous interventions after some number of days, they may be personally visited by a health worker. Note that many of these patients live in low-resource communities where each health worker manages tens to hundreds of patients; far more than they can possibly visit in a day. Thus, models that can identify patients at risk of missing doses and prioritize interventions by health workers are of paramount importance.

However, this observational data was collected via an extensive rollout to real patients, so it contains health care worker intervention effects which we must consider when training our models. Thus, a key challenge was that health workers rarely record interventions on the 99DOTS system. While there is a well-developed literature on estimating heterogeneous treatment effects, standard techniques uniformly require knowledge of which patients received an intervention [Morgan and Winship, 2014; Athey and Imbens, 2016]. We note that such gaps will be common as countries eagerly adapt DAT systems in the hopes of benefiting low-income regions; to support the delivery of improved care, we must be able to draw lessons from this messy but plentiful data.

In this work, therefore, we introduce *a general approach for learning from adherence data with unobserved interventions*, based on domain knowledge of the intervention heuristics applied by health workers. Through our partnership with Everwell, we construct a proxy for interventions present in the historical 99DOTS data and develop a model that can help prioritize targets of interventions for health workers in two highlighted clinical scenarios: 1) Real-time non-adherence risk prediction where we enable health workers to accurately **identify 21% more high-risk patients** and **catch nearly 76% more missed doses** 2) Intervention path- planning problem where we build on decision focused learning to gain an **additional 15% improvement** over standard modeling.

2 Related Work

Outcomes and adherence research are well studied in the medical literature for a variety of diseases [Kardas *et al.*, 2013], and particularly tuberculosis [Shargie and Lindtjørn, 2007; Kliiman and Altraja, 2010]. Typically these studies gather demographic and medical statistics on a cohort, observe adherence and outcomes, then retrospectively apply survival [Kliiman and Altraja, 2010] or logistic regression [Roy *et al.*, 2015] analysis to determine covariates predictive of failure. Newer work has improved modeling accuracy via machine learning techniques [Hussain and Junejo, 2018; Sauer *et al.*, 2018]. While such studies have improved patient screening at the time of diagnosis, they offer little knowledge about how risk changes *during* treatment. Here, we show how a patient's real-time adherence data can be used to track and predict risk changes throughout the course of their treatment.

We leverage the fact that, in recent years, new technologies like an electronic pill bottle cap that records the date/time when the cap is removed have made measuring daily adherence feasible. While some previous work used this data to determine predictors of non-adherence [Pellowski *et al.*, 2016; Cook *et al.*, 2017], few have studied changes in adherence over time, except one study which *retrospectively* categorized patient adherence [Kim *et al.*, 2018]. Our focus is on *prospective* identification of patients at risk of missing doses *before* failures occur.

Methodologically, our work is related to the large body of research that deals with estimating the causal impact of interventions from observational data [Morgan and Winship, 2014; Athey and Imbens, 2016]. However, they crucially require exact knowledge of when interventions were carried out. This information is entirely absent in our setting, requiring us to develop new methods for handling *unobserved* interventions in the training data.

Table 1: Data Summary. *Doses per patient was calculated only on patients enrolled at least 6 months before Sept 2018.

Metric	Count
Total doses recorded	2,169,976
—By patient call	1,459,908
—Manual (entered by health worker)	710,068
Registered phones	38,000
Patients	16,975
Health centers	252
Doses recorded per patient*	
—Quartiles	57/149/188
—Min/Mean/Max	1/136/1409
Active patients per center per month	
—Quartiles	7/18/35
—Min/Mean/Max	1/25/226

3 Data Description

99DOTS provides each patient with a cover for every sleeve of pills that associates a hidden unpredictable phone number with each daily dose (note that one dose may consist of 2-5 pills.) As patients expose pills associated with each dose, they expose one phone number per day. Each patient is instructed to place a toll-free call to the indicated number each day. The dataset contains over 2.1 million dose records for about 17,000 patients, served by 252 health centers across The City from Feb 2017 to Sept 2018. Also included for each patient were demographic features such as weight-band, age-band, gender and treatment center ID, and treatment outcome (if assigned). **Table 1** summarizes the data, but we refer readers to the full paper for more detail, accepted to KDD 2019 and in preparation [Killian *et al.*, 2019].

4 Unobserved Interventions

The TB treatment system operates under tight resource limitations, e.g. one health worker may be responsible for more than 100 patients. Thus, it is critical that workers be able to accurately rank patient risk and prioritize interventions accordingly. Machine learning can be used to accomplish such risk ranking with promising accuracy, but it requires taking special care to understand how intervention resources were allocated in the existing data.

Therefore, a key challenge is that users of the 99DOTS platform generally do not record interventions: workers may make texts, calls, or personal visits to patients to try to improve adherence, but these interventions are not routinely logged in the data. While far from ideal, such gaps are inevitable as countries with differing standards of reporting adopt DATs for TB treatment. Given the abundance of data created by DATs and their potential to impact human lives, we emphasize the importance of learning lessons in this challenging setting where unobserved interventions occur. We next resolve this challenge by formulating a general screening procedure to reshape data around intervention effects to build valid models.

Intervention Proxy. The key is to identify a conservative estimate for where interventions occur to ensure that data with intervention signals are not included. Note that we only develop this procedure for the house visit intervention, which we consider a "resource-limited" intervention since workers cannot visit all of their patients in a timely manner. Generally, this is a last resort for health workers when patients will not respond to other "non-resource-limited" interventions like calls or texts.

To formulate our proxy, we first searched for health worker guidelines for carrying out house visits. The 2005 guide by India's Revised National Tuberculosis Control Program (RNTCP) [RNTCP, 2005] required that workers deliver a house visit after a single missed dose, but updated guides are far more vague on the subject. Both the most recent guide by the WHO [WHO, 2017] and by the RNTCP [RNTCP, 2016] leave house visits up to the discretion of the health worker. However, our partners at Everwell observed that health workers prioritize non-adherent patients for resource-limited interventions such as house visits. Thus, we formulated our proxy based on the adherence dashboard seen by health workers.

The 99DOTS dashboard gives a daily "Attention Required" value for each patient, as follows. If a patient misses 0 or 1 calls in the last 7 days, they are changed to "MEDIUM" attention, whereas if they miss 4 or more they are changed to "HIGH" attention. Patients with 2-3 missed doses retain their attention level from the previous day. As our conservative proxy, we assumed that only "HIGH" attention patients were candidates for resource-limited interventions since the attention level is a health worker's primary summary of recent patient adherence. This "attention required" system for screening resource-limited interventions is generalizable to any daily adherence setting; one need only to identify the threshold for a change to HIGH attention.

With this screening system, we can identify sequences of days during which a patient was a candidate for a resource-limited intervention, and subsequently avoid using signal from those days in our training task. We accomplish this with our formulation of the real-time risk prediction task which every day predicts the risk that a MEDIUM patient will become HIGH in the next week.

Formally, for each patient who is MEDIUM at time t, use data from days $[t-6,t]$ to predict whether or not they change to HIGH at any time t_i where $t+1 \leq t_i \leq t+7$. We now demonstrate that, with our intervention proxy, resource-limited intervention effects cannot effect labels in this formulation. First, if a patient stays at MEDIUM for all t_i, then the label is 0. Since the patient was at MEDIUM for all t_i, our proxy states that no resource-limited intervention took place between our prediction time t and the time that produced the label, $t+7$. Second, if a patient changes from MEDIUM to HIGH on day t_i, then on day t_i we establish that the label is 1. By our proxy, any resource-limited intervention effect must happen in $[t_i+1, t+7]$, since attention is established at the end of a day t_i. So again, we have that no resource-limited intervention took place between our prediction time t and the time that produced the label, t_i. Thus, we ensure that intervention effects cannot influence our labels. Importantly, if we predict that a patient will have good adherence we can safely recommend no intervention since our combined screening and training method guarantees that their good adherence *is not contingent on* an intervention. This demonstrates that our classifier is suited to make predictions that prioritize resource-limited interventions.

5 Real-Time Risk Prediction

We now build a model for the prediction task formalized in **Section 4** which leverages our intervention screening proxy. Our goal was to develop a model corresponding to the health worker's daily task of using their patients' recent call history to evaluate adherence risk with the goal of scheduling different types of interventions. Better predictions allow workers to proactively intervene with more patients before they miss critical doses.

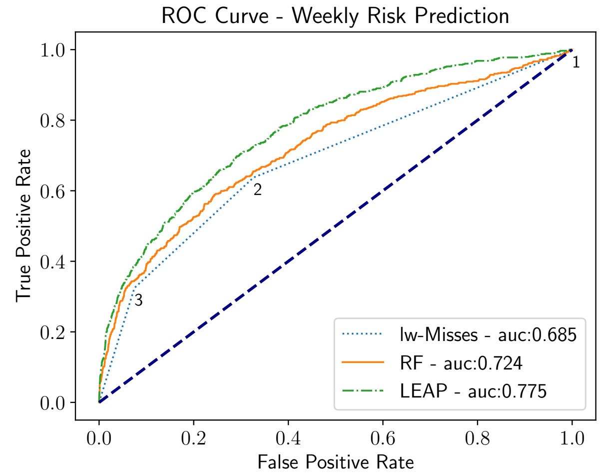

Figure 2: ROC Curve for the weekly risk prediction task comparing the missed call baseline (blue), Random Forest (yellow) and LEAP (green). Numbers under the blue curve give thresholds used to calculate the baseline's ROC curve.

Sample Generation. We started with the full population of 16,975 patients and generated training samples from each patient as follows. We considered all consecutive sequences of 14 days of call data where the first 7 days of each sequence were non-overlapping. We excluded any sequence of data with more than two manual doses or where no calls were missed. This generated 16,015 samples (2,437 positive). We implement a random forest as well as a combined LSTM and dense network, named LEAP. Both models receive demographic and descriptive features engineered from call data as input. LEAP also receives the raw call data as a binary sequence. We refer the reader to the full paper for more details about models and features [Killian *et al.*, 2019].

Model Evaluation. We randomized all data then separated 25% as the test set. The baseline we compared against was the method used by the existing 99DOTS platform to asses risk, namely calls made by the patient in the last week (lw-Misses). **Figure 2** shows that the random forest narrowly outperforms the baseline and LEAP clearly outperforms both. We next consider how each method might be used to plan house-visit interventions. Since this is a very limited resource, we set the strictest baseline threshold to consider patients for this intervention; that is 3 missed calls. Fixing the FPR of this baseline method, **Table 2** shows how many more patients in the test set would be reached each week by our method (as a result of its higher TPR) as well as the improvement in number of missed doses caught. To calculate missed calls caught, we count only missed doses that occur before the patient moves to HIGH risk. *Our model catches 21.6% more patients and 76.5% more missed calls, demonstrating substantially more precise targeting than the baseline.*

6 Decision Focused Learning

We now explore a case study of how our LEAP model can be specialized to provide decision support for a particular

Table 2: LEAP vs. Baseline - Missed Doses Caught

Method	True Positives	Doses Caught
Baseline	204	204
LEAP	248	360
Improvement	**21.6%**	**76.5%**

LEAP vs. baseline for catching missed doses with a fixed false positive rate. Our method learns behaviors indicative of non-adherence far earlier than the baseline, allowing for more missed doses to be prevented.

intervention. We exploit end-to-end differentiability of the model to replace our earlier loss function (binary cross-entropy) with a performance metric tailored to the objective and constraints of specific decision problem. To accomplish this end-to-end training, we leverage recent advances in *decision-focused learning*, which embeds an optimization model in the loop of machine learning training [Wilder *et al.*, 2018; Donti *et al.*, 2017].

We focus on a specific optimization problem that models the allocation of health workers to intervene with patients who are at risk in the near future. This prospective intervention is enabled by our real-time risk predictions and serves as an example of how our system can enable proactive, targeted action by providers. However, we emphasize that our system can be easily modified to capture other intervention problems. Such flexibility is one benefit to our technical approach, which allows the ML model to *automatically* adapt to the problem specified by a domain expert.

Our optimization problem models a health worker who plans a series of interventions over the course of a week. The health worker is responsible for a population of patients across different locations, and may visit one location each day, allowing them to intervene with all of the patients at that location. The optimization problem is to select a set of locations to visit which maximizes the number of patients who receive an intervention *on or before the first day they would have missed a dose*. We refer to this quantity as the number of *successful interventions*.

We now show how this optimization problem can be formalized as a linear program. We have a set of locations $i = 1...L$ and patients $j = 1...N$ where patient j has location ℓ_j. Over days of the week $t = 1...7$, the objective coefficient c_{jt} is 1 if an intervention on day t with patient j is successful and 0 otherwise. Our decision variable is x_{it}, and takes the value 1 if the health worker visit location i on day t and 0 otherwise. With this notation, the final LP is as follows:

$$\max_x \sum_{t=1}^{7} \sum_{i=1}^{L} x_{it} \left(\sum_{j:\ell_j = i} c_{jt} \right)$$

$$\text{s.t.} \sum_{i=1}^{L} x_{it} \le 1, t = 1...7$$

$$\sum_{t=1}^{7} x_{it} \le 1, i = 1...L$$

$$0 \le x_{it} \le 1 \quad \forall i, t$$

where the second constraint prevents the objective from double-counting multiple visit to a location. The machine learning task is to predict the values of the c_{jt}, which are unknown at the start of the week. We compare three models: 1) an extension of the lw-Misses baseline 2) LEAP trained using standard cross-entropy loss and 3) LEAP trained to predict c_{jt} using performance on the above optimization problem as the loss function via the differentiable surrogate given by [Wilder *et al.*, 2018] (LEAP-Decision). We created instances of the decision problem by randomly partitioning patients into groups of 100, modeling a health worker under severe resource constraints (as they would benefit most from such a system).

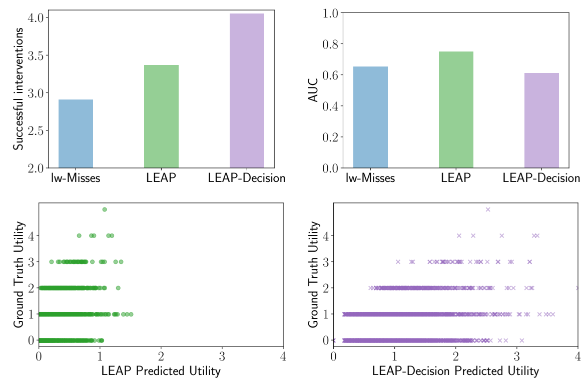

Figure 3: Results for decision focused learning problem. Top row: successful interventions and AUC for each method. Bottom row: visualizations of model predictions.

Figure 3 shows results for this task. In the top row, we see that LEAP and LEAP-Decision both outperform lw-Misses, as expected. LEAP-Decision improves the number of successful interventions by approximately 15% compared to LEAP, demonstrating the value of tailoring the learned model to a given planning problem. LEAP-Decision actually has worse AUC than either LEAP or lw-Misses, indicating that typical measures of machine learning accuracy are not a perfect proxy for utility in decision making. To investigate what specifically distinguishes the predictions made by LEAP-Decision, the bottom row of **Figure 3** shows scatter plots of the predicted utility at each location according to LEAP and LEAP-Decision versus the true values. Visually, LEAP-Decision appears better able to distinguish the high-utility outliers which are most important to making good decisions. Quantitatively, LEAP-Decision's predictions have worse correlation with the ground truth overall (0.463, versus 0.519 for LEAP), but better correlation on locations where the true utility is strictly more than 1 (0.504 versus 0.409). Hence, decision-focused training incentivizes the model to focus on making accurate predictions specifically for locations that are likely to be good candidates for an intervention. This demonstrates the benefit of our flexible machine learning modeling approach, which can use custom-defined loss functions to automatically adapt to particular decision problems.

Acknowledgments

We thank the Everwell team, especially Brandon Liu, Priyanka Ivatury, Amy Chen, and Bill Thies for sharing the data and thoughtfully advising us each step of the way. This research was supported by MURI Grant W911NF-18-1-0208.

References

[99DOTS, 2019] 99DOTS. 99dots, 2019.

[Athey and Imbens, 2016] Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.

[Cook *et al.*, 2017] Paul F Cook, Sarah J Schmiege, Whitney Starr, Jane M Carrington, and Lucy Bradley-Springer. Prospective state and trait predictors of daily medication adherence behavior in hiv. *Nursing research*, 66(4):275–285, 2017.

[Corden *et al.*, 2016] Marya E Corden, Ellen M Koucky, Christopher Brenner, Hannah L Palac, Adisa Soren, Mark Begale, Bernice Ruo, Susan M Kaiser, Jenna Duffecy, and David C Mohr. Medlink: A mobile intervention to improve medication adherence and processes of care for treatment of depression in general medicine. *Digital Health*, 2:2055207616663069, 2016.

[Donti *et al.*, 2017] Priya Donti, Brandon Amos, and J Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, pages 5484–5494, 2017.

[Everwell, 2019] Everwell. Everwell, 2019.

[Garfein *et al.*, 2015] Richard S Garfein, Kelly Collins, Fátima Muñoz, Kathleen Moser, Paris Cerecer-Callu, Fredrick Raab, Phillip Rios, Allison Flick, María Luisa Zúñiga, Jazmine Cuevas-Mota, et al. Feasibility of tuberculosis treatment monitoring by video directly observed therapy: a binational pilot study. *The International Journal of Tuberculosis and Lung Disease*, 19(9):1057–1064, 2015.

[Haberer *et al.*, 2017] Jessica E Haberer, Nicholas Musinguzi, Alexander C Tsai, BM Bwana, C Muzoora, PW Hunt, JN Martin, DR Bangsberg, et al. Real-time electronic adherence monitoring plus follow-up improves adherence compared with standard electronic adherence monitoring. *AIDS (London, England)*, 31(1):169–171, 2017.

[Hussain and Junejo, 2018] Owais A Hussain and Khurum N Junejo. Predicting treatment outcome of drug-susceptible tuberculosis patients using machine-learning models. *Informatics for Health and Social Care*, pages 1–17, 2018.

[Kardas *et al.*, 2013] Przemyslaw Kardas, Pawel Lewek, and Michal Matyjaszczyk. Determinants of patient adherence: a review of systematic reviews. *Frontiers in pharmacology*, 4:91, 2013.

[Killian *et al.*, 2019] Jackson A Killian, Bryan Wilder, Amit Sharma, Vinod Choudhary, Bistra Dilkina, and Milind Tambe. Learning to prescribe interventions for tuberculosis patients using digital adherence data. *arXiv preprint arXiv:1902.01506*, 2019.

[Kim *et al.*, 2018] Kyuhyung Kim, Bumhwi Kim, Albert Jin Chung, Keekoo Kwon, Eunchang Choi, and Jae-wook Nah. Algorithm and system for improving the medication adherence of tuberculosis patients. In *2018 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 914–916. IEEE, 2018.

[Kliiman and Altraja, 2010] K Kliiman and A Altraja. Predictors and mortality associated with treatment default in pulmonary tuberculosis. *The international journal of tuberculosis and lung disease*, 14(4):454–463, 2010.

[Liu *et al.*, 2015] Xiaoqiu Liu, James J Lewis, Hui Zhang, Wei Lu, Shun Zhang, Guilan Zheng, Liqiong Bai, Jun Li, Xue Li, Hongguang Chen, et al. Effectiveness of electronic reminders to improve medication adherence in tuberculosis patients: a cluster-randomised trial. *PLoS medicine*, 12(9):e1001876, 2015.

[Morgan and Winship, 2014] Stephen L Morgan and Christopher Winship. *Counterfactuals and causal inference*. Cambridge University Press, 2014.

[Pellowski *et al.*, 2016] Jennifer A Pellowski, Seth C Kalichman, Moira O Kalichman, and Chauncey Cherry. Alcohol-antiretroviral therapy interactive toxicity beliefs and daily medication adherence and alcohol use among people living with hiv. *AIDS care*, 28(8):963–970, 2016.

[RNTCP, 2005] RNTCP. Technical and operational guidelines for tuberculosis control, 2005.

[RNTCP, 2016] RNTCP. Revised national tb control programme: Technical and operational guidelines for tuberculosis control in india 2016, 2016.

[Roy *et al.*, 2015] Nirmalya Roy, Mausumi Basu, Sibasis Das, Amitava Mandal, Debashis Dutt, and Samir Dasgupta. Risk factors associated with default among tuberculosis patients in darjeeling district of west bengal, india. *Journal of family medicine and primary care*, 4(3):388, 2015.

[Sauer *et al.*, 2018] Christopher Martin Sauer, David Sasson, Kenneth E Paik, Ned McCague, Leo Anthony Celi, Ivan Sanchez Fernandez, and Ben MW Illigens. Feature selection and prediction of treatment failure in tuberculosis. *PloS one*, 13(11):e0207491, 2018.

[Shargie and Lindtjørn, 2007] Estifanos Biru Shargie and Bernt Lindtjørn. Determinants of treatment adherence among smear-positive pulmonary tuberculosis patients in southern ethiopia. *PLoS medicine*, 4(2):e37, 2007.

[Subbaraman *et al.*, 2018] Ramnath Subbaraman, Laura de Mondesert, Angella Musiimenta, Madhukar Pai, Kenneth H Mayer, Beena E Thomas, and Jessica Haberer. Digital adherence technologies for the management of tuberculosis therapy: mapping the landscape and research priorities. *BMJ global health*, 3(5):e001018, 2018.

[Thomas *et al.*, 2005] Aleyamma Thomas, PG Gopi, T Santha, V Chandrasekaran, R Subramani, N Selvakumar, SI Eusuff, K Sadacharam, and PR Narayanan. Predictors of relapse among pulmonary tuberculosis patients treated in a dots programme in south india. *The International Journal of Tuberculosis and Lung Disease*, 9(5):556–561, 2005.

[WHO, 2017] WHO. Handbook for the use of digital technologies to support tuberculosis medication adherence. 2017.

[WHO, 2018] WHO. *Global tuberculosis report 2018*. World Health Organization, 2018. Licence: CC BY-NC-SA 3.0 IGO.

[Wilder *et al.*, 2018] Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *AAAI Conference on Artificial Intelligence*, 2018.