Optimization of Insulin Dosing in Diabetic Patients with POMDPs

Chuma Kabaghe
Stanford University
chumak@stanford.edu

Seyed Shahabeddin Mousavi
Stanford University
ssmousav@stanford.edu

Abstract

In order to stay alive, people living with diabetes rely on insulin injections. The doses of these injections are decided by a healthcare team based on the patient's blood glucose level, how much their blood sugar fluctuates throughout the day and their lifestyle. In order to provide the team with the data needed to decide on the optimal doses, patients carry the burden of having to continually monitor their blood glucose levels, meal intake and potentially other changes in lifestyle and activity. To help ease this burden, we model the decision of determining the optimal insulin dose and injection timing as a partially observable Markov decision process (POMDP), where the ultimate goal is to maintain the patient's blood glucose level in the healthy range. We experimented with three online solvers: POMCP, AR-DESPOT, and POMCPOW. After hyperparameter tuning all three solvers yielded the ideal reward but took significantly different amounts of time to run with POMCPOW being the fastest.

1 Introduction

Diabetes is a global pandemic. An estimated 425 million people worldwide have diabetes, accounting for 12% of the world's health expenditures.[2] Diabetes mellitus refers collectively to a group of diseases resulting from dysfunction of the glucoregulatory system. Hyperglycemia, the hallmark of diabetes, is the primary consequence of this dysregulation and is associated with long-term complications and can be fatal. Artificial intelligence methods in combination with the latest technologies have the potential to enable the creation and delivery of better management services to deal with such chronic diseases.

Despite recent advances in glucose monitoring devices and insulin pumps, the burden of determining insulin doses still lies solely on the patient or caregiver. Throughout the day they have to make decisions about what amount of insulin to pump and how quickly and also how much food to eat and when. Warranting the obvious issues and complexities that come to mind including choosing an incorrect insulin dose or administering it at the wrong time which can be life threatening.

To ease the burden of making these various complex decisions, we modeled the decision of determining the optimal insulin dose and injection timing as a partially observable Markov decision process (POMDP), where the ultimate goal is to maintain the patient's blood glucose level in the healthy range of 70-180mg/dL.

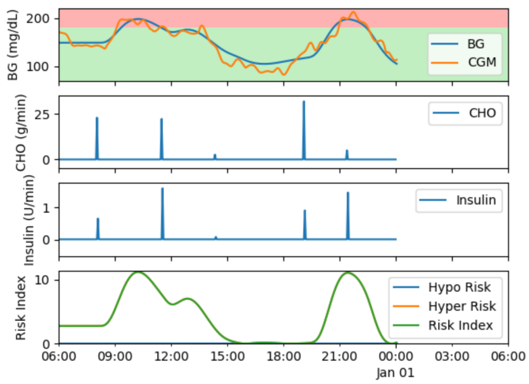

Sample visualization of changes in blood glucose levels (BG), carbohydrates (CHO), insulin and risk index

2 Related Work

There has been a great deal of prior work on the general applications of Artificial Intelligence to the monitoring, analysis, and control of the complications related to diabetes. A PubMed search yields at least 450 clinically-relevant and high-impact articles published in the last decade related to the field of applied AI in diabetes care. The AI applications aim to improve a broad spectrum of diabetes care, from diabetes screening and detection to monitoring and treatment, and included apps, devices, and systems that aid patients, clinicians, and health systems. A review of these articles suggests that AI applications are aiming to transform diabetes care in 4 main areas: automated retinal screening, clinical decision support, predictive population risk stratification, and patient self-management tools, as summarized in Table 1.

Category	Number of articles	Most common clinical AI applications
Automated Retinal Screening	96	Detection of diabetic retinopathy, maculopathy, exudates, and other abnormalities from normal findings
Clinical Decision Support	126	Detection and monitoring of diabetes and comorbidities such as neuropathy, nephropathy and wounds
Predictive Population Risk Stratification	135	Identification of diabetes subpopulations at higher risk for complications, hospitalization, and readmissions
Patient Self-Management Tools	94	AI-improved glucose sensors, artificial pancreas, activity and dietary tracking devices
Total	450	

Table 1: Categorization of Artificial Intelligence and Diabetes Care

A diverse and complex set of AI approaches and cognitive computing systems have been employed in such studies. A number of the more common AI approaches described in the research and their clinical applications in diabetes care are:

- Multilayer perceptron: Used for prediction models & patient self-management tools
- Random forest: Used for retinal screening, decision support, prediction models, and patient self-management tools
- Support vector machine (SVM): Used for retinal screening, decision support, prediction models, patient self-management tools
- K-nearest neighbors algorithm (KNN): Used for retinal screening, decision support, prediction models, patient self-management tools

Among the panoply of research papers on the applications of artificial intelligence to health care in general and diabetes in particular, it is evident that work specifically focused on the application of MDPs and POMDPs to these fields, while existent, are underrepresented. There has been notable prior work on the applications of Markov Decision Processes (MDPs) for screening and treatment of chronic diseases, including but not limited to, Type 2 Diabetes. [3] But we believe there is significant room for growth in this field.

3 Approaches

The POMDP for this problem is defined as follows:

State Space $\mathcal{S}$: {s $\in$ [0.0, ∞]} Blood Glucose (BG) levels

Action Space $\mathcal{A}$: {a $\in$ [0.0, 30.0]} How much basal (background insulin) to administer

Rewards $\mathcal{R}$: {r $\in$ [-∞, 0.0]} A function of blood glucose measurements in the last hour. The reward at each step is risk[t-1] - risk[t]. risk[t] is the risk index at time t defined in the paper Statistical Tools to Analyze Continuous Glucose Monitor (CGM) Data [1]. Zero is the ideal reward.

Observation Space $\mathcal{O}$: {o $\in$ [0.0, ∞]}: blood glucose level measured by CGM sensor.

Explicitly defining a transition and observation model for this POMDP is intractable. Thus, we used a generative model provided by a modified Simglucose v0.2.1 [7]. In order to interface with POMDP.jl which we used to solve this problem, we had to modify Simglucose to return a next state, observation and reward given an action which is different from the default implementation that takes in a time range and returns observations and rewards obtained from running the simulation in the provided time range. Simglucose v0.2.1 is a python implementation of the FDA-approved UVa/Padova. All simulations were run on simglucose-adolescent2-v0 with the default controller and the same random seed. A discount factor of 0.9 was used in solving this POMDP.

3.1 Partially Observable Monte Carlo Planning (POMCP)

POMCP is a Monte-Carlo algorithm for online planning in large POMDPs. The algorithm combines a Monte-Carlo update of the agent's belief state with a Monte-Carlo tree search from the current belief state. POMCP has two important properties. First, Monte-Carlo sampling is used to break the curse of dimensionality both during belief state updates and during planning. Second, only a black box simulator of the POMDP is required, rather than explicit probability distributions. [6] These properties enable POMCP to plan effectively in significantly larger POMDPs like the one we are trying to solve.
We experimented with different values of the UCB exploration constant.

UCB exploration constant	1.0	10.0	50.0	100.0
Reward	0.0	0.0	0.0	0.0

Table 2: Hyperparameter tuning results for UCB exploration constant

As can be seen in Table 2, varying the UCB exploration constant always yielded reward of 0, which is the ideal result. Thus, the default value of 1.0 was used.

3.2 AR-DESPOT (Anytime Regularized DEterminized Sparse Partially Observable Tree)

AR-DESPOT is an online POMDP solver whose best policy has been shown to be near optimal. Determinized Sparse Partially Observable Tree (DESPOT) is a sparse approximation of the standard belief tree which focuses online planning on a set of randomly sampled scenarios and compactly captures the "execution" of all policies under these sampled scenarios.[5]

In addition to the default hyperparameters for AR-DESPOT solver, we used epsilon_0 = 0.5, D=1, max_trials = 1, K=2, bounds=(DefaultPolicyLB(RandomSolver()), 350.0).

3.3 POMCPOW

POMCPOW is an online solver based on Monte Carlo tree search for POMDPs with continuous state, action, and observation spaces. It solves problems specified using the POMDPs.jl interface and its requirements are the same as for an importance-sampling particle filter - a generative model for the dynamics and an explicit observation model. In this algorithm, the belief updates are weighted, but they also expand gradually as more simulations are added. Furthermore, since the richness of the belief representation is related to the number of times the node is visited, beliefs that are more likely to be reached by the optimal policy have more particles. At each step, the simulated state is inserted into the weighted particle collection that represents the belief, and a new state is sampled from that belief.[4] Experimenting with different values for tree_queries and max_time resulted in optimal values of tree_queries = 10 and max_time =10.

4 Results & Analysis

To evaluate the performance of our planners, we compared the rewards obtained by their policies to those of a random policy. The difference between the two values is our evaluation metric, which we will refer to as score. A larger score indicates a better policy.

For all experiments, our POMDP was initialized at state BG =169.66951984163407. For evaluation we used the same seed for all policies and simulated 100 steps which corresponds to 100 blood glucose measurements taken every 5 minutes. So 8.33 hours in the life of the adolescent represented by the simulation.

Given our large and continuous state space, using online POMDP solvers made solving the POMDP more tractable. After hyperparameter tuning, all three solvers returned policies that yielded an ideal reward of 0.0. This is not surprising given that our reward is a function of blood glucose measurements in the last hour. The reward at each step is risk[t-1] - risk[t]. Since our evaluation is over 8.33 hour period and the risk value does not change much if the observations are within the healthy region. We hypothesize that there are multiple policies that would result in risk[t-1] - risk[t]=0 after 8.33 hours. With the right hyperparameters, each of these solvers was able to produce one of these policies.

Though all the solvers yielded a perfect reward, each took a significantly different amount of time to run. All solvers were run on a 64 bit, Intel Core i7 CPU @2.40GHz, 16.0GB RAM.

Solver	POMCP	POMCPOW	AR-DESPOT
Runtime	4 hours	10 minutes	3 hours

Table 3: Run times for different solvers

Solver	Random	POMCP	AR-DESPOT	POMCPOW
Score	0	13.6646	13.6646	13.6646

Table 4: Results

Sample Policy Output from POMCPOW

Timestep	s	a	r	o
1	169.6695	1.1	-9.6993	56.5565
2	50.8030	21.2	0.0	56.5565
3	50.8030	14.7	-0.8070	55.6021

Table 5: Sample policy from POMCPOW

5 Conclusion

This article provides three different online POMDP solutions to the problem of insulin administeration. We have utilized Simglucose, a Type-1 diabetes simulator implemented in Python, for a generative model to feed data to three separate online solvers, POMCP, AR-DESPOT, and POMCPOW; from the JuliaPOMDP package. We concluded that using online solvers would best meet our needs given our model has a large continuous space, and grasping with the challenges of trying out the offline Fast Informed Bound (FIB) solver with preproduced datasets from the Simglucose simulator only reinforced our decision. After hyperparameter tuning & having run the different online models for extended periods of time, they performed well and all provided solutions with the optimal reward of zero. The main distinction between the different solvers was that one of them, POMCPOW, took significantly less time than the other two to complete. We would undoubtedly like to minimize the runtime of our algorithm if maintaining the optimal reward while doing so is possible, especially given that these are online solvers, so we proceed to elect the POMCPOW solver as our recommendation for future research in the field.

Further research directions would be to test with simulations of other individuals that Simglucose provides, as we only tested with simulations of one adolescent. Also testing on other reputable type 1 diabetes simulators (Simglucose was the only implementation of an FDA-approved simulator, UVa/Padova Simulator (2008 version), we found) that might exist. Given the opportunity, testing the results against real-world data in medical institutions to gauge the viability of the simulators and provided POMDP solutions would also be ideal.

Our overall aim is to ultimately succeed in developing a reliable, highly accurate system that would relieve diabetes patients and their family of the burden of guesswork when it comes to the health of the patient.

6 Group Member Contributions

Chuma - Researched various libraries for implementing the solvers, modified the Simglucose simulator, implemented POMCP, AR-DESPOT, and POMCPOW solvers and contributed to the writing of the proposal, update, and final report. Participated in several in-person project group meetings to brainstorm and collaborate on our project at various stages throughout the quarter.

Shahab - Researched various articles and papers on the state of the diabetic health industry and the applications of AI and technology to the industry, modified the Simglucose simulator, attempted the offline solver (FIB) with preproduced datasets from the simulator, and contributed to the writing of the proposal, update, and final report. Participated in several in-person project group meetings to brainstorm and collaborate on our project at various stages throughout the quarter.

A link to our code: https://github.com/chuma9/cs238-project

References

[CK09] William Clarke and Boris Kovatchev. "Statistical Tools to Analyze Continuous Glucose Monitor Data". In: *Diabetes Technology Therapeutics* 11.S1 (Feb. 2009). DOI: 10.1089/dia.2008.0138.

[Dan] et al. Dankwa-Mullan I. Rivo M. *Transforming Diabetes Care Through Artificial Intelligence: The Future Is Here (2019)*. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6555175/. (accessed: 4.12.2019).

[SD] Lauren Steimle and Brian Denton. *Markov Decision Processes for Screening and Treatment of Chronic Diseases (2017)*. URL: http://btdenton.engin.umich.edu/wp-content/uploads/sites/138/2015/10/Steimle-2015.pdf. (accessed: 23.11.2019).

[SK17] Zachary Sunberg and Mykel Kochenderfer. *Online algorithms for POMDPs with continuous state, action, and observation spaces*. 2017. arXiv: 1709.06196 [cs.AI].

[Som+13] Adhiraj Somani et al. "DESPOT: Online POMDP Planning with Regularization". In: *Advances in Neural Information Processing Systems 26*. Ed. by C. J. C. Burges et al. Curran Associates, Inc., 2013, pp. 1772–1780. URL: http://papers.nips.cc/paper/5189-despot-online-pomdp-planning-with-regularization.pdf.

[SV10] David Silver and Joel Veness. "Monte-Carlo Planning in Large POMDPs". In: *Advances in Neural Information Processing Systems 23*. Ed. by J. D. Lafferty et al. Curran Associates, Inc., 2010, pp. 2164–2172. URL: http://papers.nips.cc/paper/4031-monte-carlo-planning-in-large-pomdps.pdf.

[Xie] Jinyu Xie. *Simglucose v0.2.1 (2018) [Online]*. URL: https://github.com/jxx123/simglucose. (accessed: 14.11.2019).