Title: Precision medicine as a control problem: using simulation and deep reinforcement learning to discover adaptive, personalized multi-drug therapies.

PI(s) Grant: Gary An and Daniel Faissol

Institution(s): Lawrence Livermore National Laboratory, University of Vermont

Grant Number: UO1EB025825

Title of Grant(s): Agent-based Models to address the Crisis of Reproducibility and Precision Medicine

Abstract Authors: Daniel Faissol, Brenden Petersen, Gary An, Chase Cockrell

Abstract Text:

We explore the concept of using simulation and deep reinforcement learning to discover adaptive, personalized multi-drug therapies for complex diseases. We present sepsis as an initial test case for this approach. Sepsis is a life-threatening condition affecting one million people per year in the US in which dysregulation of the body's own immune system causes damage to its tissues, resulting in a 28 - 50% mortality rate. Clinical trials for sepsis treatment over the last 20 years have failed to produce a single currently FDA approved drug treatment.

Methods:
In this study, we attempt to discover an effective cytokine mediation treatment strategy for sepsis using a previously developed agent-based model that simulates the innate immune response to infection: the Innate Immune Response agent-based model (IIRABM). Previous attempts at reducing mortality with multi-cytokine mediation using the IIRABM have failed to reduce mortality across all patient parameterizations and motivated us to investigate whether adaptive, personalized multi-cytokine mediation can control the trajectory of sepsis and lower patient mortality. We used the IIRABM to compute a treatment policy in which systemic patient measurements are used in a feedback loop to inform future treatment.

Results:
Using deep reinforcement learning, we identified a policy that achieves 0% mortality on the patient parameterization on which it was trained. More importantly, this policy also achieves 0.8% mortality over 500 randomly selected patient parameterizations with baseline mortalities ranging from 1 - 99% (with an average of 49%) spanning the entire clinically plausible parameter space of the IIRABM. These results suggest that adaptive, personalized multi-cytokine mediation therapy could be a promising approach for treating sepsis. Moreover, we hope that this work motivates researchers to consider such an approach as part of future clinical trials.

Ongoing work:
This work has motivated the need for more robust deep reinforcement learning approaches. In particular, the results generated by reinforcement learning are sensitive to the parameterization of the agent-based model. Given the uncertainty in parameter values for agent-based models of diseases, we are working on a new approach consisting of: (1) allow increased

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC. LLNL-ABS-794180

heterogeneity/uncertainty in model parameters of the IIRABM, (2) constrain the parameter space by requiring the parameters to lead to plausible simulation output, and (3) search for cytokine mediation policies that are effective for all plausible parametrizations using more advanced DRL approaches.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC. LLNL-ABS-794180