Human Immunodeficiency Virus (HIV) is a potential precursor to Acquired Immune Deficiency Syndrome (AIDS), which commonly treated with a mixture of reverse transcriptase inhibitors (RTI) and protease inhibitors (PI). Adams, et al (2004) used a simulator to dynamically model the patients health as a state space of patient health indicators that can be modified by actions, providing different types of treatment. Ernst, et al, (2006) found, by applying reinforcement learning to this model of states and actions, they could predict potentially effective HIV treatment policies. However, Ernst 2006 assumed that patients reliably present themselves for treatment every five days and that all patients respond to prescribed treatments similarly. These assumptions are unrealistic.

We investigate whether we can produce similarly effective treatment policies even when faced with sporadically sampled data, and we account for the physiological variation of individual patients. By imputing the patients missing personal health indicators, we find that we can often still compute effective treatment policies. We also show that parameter variation results in different optimal policies, a first step to motivate the need for transfer learning in this domain.