

Title: What It Is To Be Conscious: Exploring the Plausibility of Consciousness in Deep Learning Computers

Author: (Peter) Zach Davis

Faculty Advisors: Kristina Striegnitz and David Barnett

As artificial intelligence and robotics progress further and faster every day, designing and building a conscious computer appears to be on the horizon. Recent technological advances have allowed engineers and computer scientists to create robots and computer programs that were previously impossible. The development of these highly sophisticated robots and AI programs has thus prompted the age-old question: can a computer be conscious? The answer relies on addressing two key sub-problems. The first is the nature of consciousness: what constitutes a system as conscious, or what properties does consciousness have? Secondly, does the physical make-up of the robot or computer matter? Is there a particular composition of the robot or computer that is necessary for consciousness, or is consciousness unaffected by differences in physical properties? My aim is to explore these issues with respect to deep-learning computer programs. These programs use artificial neural networks and learning algorithms to create highly sophisticated, seemingly intelligent computers that are comparable to, yet fundamentally different from, a human brain. Additionally, I will discuss the required actions we must take in order to come to a consensus on the consciousness of deep learning computers.