Part II Supporting Material

1. I have gravitated towards the nerd community for most of my life. As a child, I preferred hanging out at the Museum of Natural History over playing organized sports. I actually used to play hooky during elementary school to hang out at the Museum. This quite strong love of bleeding edge scientific and technological research has not wavered in the least since those early days of delinquency.

In high school, I joined my school's F.R.C. robotics team before my freshman year had even begun. During my first three years on the team, I completely revolutionized the robot's autonomous control system. Up until that point, autonomous mode programming was done manually, tediously, by tuning certain "gain" values which calibrated the system's PID controller. I proposed a much more intuitive and conceptually interesting system which drove the robot as a human would. The new system was inspired by pondering how I, a human being, can navigate a crowded New York City sidewalk without even thinking about it. In this new system, the autonomous driving code focuses on a distant geographical point which it then tries to navigate towards. By creating many points along a path (using bézier curves) we were able to get the robot to follow any given vector path. Drivers could now use a program like Adobe Photoshop to plot new autonomous courses instead of relying on the programming team to manually write each route.

During my high school summers, I interned with numerous companies. No other internship was as rewarding or beneficial as my experience with Infoexchange, a financial information trading web company. While at the company, I was able to work with top executives and suggest ideas of improvement at corporate board meetings. My work mainly focused on the company's technical backend. Many of my proposals to the structure of their codebase wound up being implemented and put into production. In the words of my former boss: "Having worked in the Information Technology field for the better part of 15 years, I can say without a doubt that Jordan rivals the maturity level and work performance of people twice his age. He would often think of new approaches to designs and enthusiastically come back with sketches and ideas he made at home. Jordan has a strong ability to communicate complex ideas, and I enjoyed running through designs with him on a white-board. His exuberance and desire to help bring forward new ideas set him apart as a truly unique individual."

I believe that these examples show my complete representation of the characteristics and traits put forth in the question. I have shown that I have the drive, the skills, and the desire to become a strong researcher in the future academic landscape.

2. Finding new heuristic measures (like Integrated Information) might help future simulators "cheat" evolutionary time and "guide" simulated populations down a more intelligent evolutionary path. Improving raw computational power also has the ability to advance the demonstrated abilities of artificial evolution. In the words of Virgil Griffith, a leading researcher in this field: "Increasing CPU cycle speed has proven an easy feat in modern times, creating artificial intelligence has not. This is a way to convert the raw advancement in CPU speed to advancement into raw advancement in artificial intelligence." I suppose one major long-term question for this field might be: To what degree will using heuristic guidance actually improve simulation results? Will these guided tests actually show significant speed improvement over a more traditional evolution simulation? Is there any sort of point of diminishing return after which speed improvement is unnecessary for intelligent system development?

My work will provide future researchers with confidence that Integrated Information will indeed yield significant improvement to simulation time. My research is one of many crucial pieces of work that must be completed before these heuristics are put to the test in developing more advanced systems.

To take a wildly fantastic voyage into the future: The scientific community hopes to soon obtain a full functional map of the nervous system of *C. Elegans*, a small yet independently living nematode worm. This creature is a promising specimen for this kind of full mapping because it has one of the simplest and smallest nervous systems of any creature on Earth. The anatomy of *C. Elegans'* nervous system has already been fully mapped, yet a full functional map of the creature's brain at work has yet be obtained. Once the scientific community has access to the full functional network of this specimen, researchers will undoubtedly test different types of heuristic measures, such as Integrated Information, with it's naturally occurring biological network. How will Integrated Information behave when tested on real biological networks? What might that mean for artificial analogs of *C. Elegans* that hold similar heuristic values? Should they be considered as living systems?

I cannot begin to answer those questions, but with any luck, I will be around to lead the scientific inquiry that may one day uncover their answers.

3. Writing intelligent code is *difficult* and *time consuming.* Over the past few decades, researchers have developed techniques by which systems (such as programs and neural networks) can be "evolved" in simulation without human interaction. Specifically, this project analyzes how Integrated Information (a physical property of any neural network) changes as the network is evolved. Scientists hope to use Integrated Information as an automatic intelligence "grade" that can be used by the simulator to "push" developing systems down more intelligent evolutionary paths. This paper presents new experimental data that supports the use of Integrated Information as such.

4. Evolution has proven to be a wildly successful autonomous process for creating intelligent systems in the natural world and in simulation. Since the early 1960s, researchers have used artificial evolution to find ingenious and novel solutions to complex problems such as series prediction and flight control. Recently, artificial evolution has been applied to neural networks with the aim of evolving more robust artificial intelligence. Several metrics have been proposed to chart the emergence of intelligence in these evolved networks. This work analyzes the behavior of a new metric, integrated information (ϕ). Observed data is analyzed, interpreted, and compared to more conventional properties of the artificial neural network. The data analysis shows that ϕ increases over evolutionary time and is therefore promising as a heuristic measure.

5a. This work was done solely by me, Jordan Perr-Sauer, with absolutely no full-time oversight or supervision of any kind. My personal laptop was sufficient for conducting most of the experiments mentioned in the paper. For certain high intensity tests I used a few of the machines in my school's computer lab. This was done under the supervision of one of the school's computer science teachers, **(Mike Zamansky: MZamans@schools.nyc.gov)**, though he did not interfere with my work in any way.

I was inspired to do research in this field by a Google Tech Talk given by **(Virgil Griffith: virgil@caltech.edu)** on Polyworld and artificial evolution. I reached out to him and was able to get in email correspondence. Once I decided upon my topic of study, Virgil made himself available for technical assistance and the odd higher mathematical question. While he did do a great deal of conceptual lecturing over the phone, Virgil did not explicitly participate in running any of my actual experiments. He was a great resource and a great technical help.

Before I came in contact with Virgil Griffith I reached out to the author and maintainer of the Polyworld simulation environment, **(Larry Yaeger: laryy@indiana.edu)**, through Polyworld's online development forum. Larry was incredibly helpful in providing support for Polyworld which I used extensively in my research.

5b. The development of my paper's idea began on a whim. I knew from the very start that I wanted to do something in artificial intelligence. As I researched this broad topic I found many online references to something called "artificial evolution." I had always admired the idea of evolving artificial creatures using software (though I was unaware that the practice had an official name) so I bought Lawrence J. Fogel's pioneering text (from the late 60s!) on the subject. Inspired by his work, I spent my entire April school vacation mimicking his experiments and exploring my results. Originally, I wanted to perform more philosophical research into the feasibility of creating conscious life through artificial evolution. I toyed with the idea for a few months before discounting it as too broad and probably not feasible. Later, while aimlessly poking around the Internet, I found a Google Tech Talk given by Virgil Griffith about a modern version of Fogel's work, Polyworld. I reached out to Virgil Griffith and during an amazing two hour phone conversation he briefed me on the current state of affairs in artificial evolution research. I left that conversation with a few excellent starting points for research and subsequently spent a good week reading through related research papers. We had another phone conversation in which we discussed the pros and cons of choosing Integrated Information analysis as my major topic.

5c. I had been intrigued by evolutionary programming and the simulation of life inside an artificial machine long before I started this project. Research for this work began in January of 2009 and continued until Intel's deadline in November of 2009.

5d. This work was individual.

5e. I learned almost everything by myself. As I stated above: Virgil Griffith helped verbally explain some of the finer concepts of Integrated Information, solidifying my understanding of the metric, and he helped construct an algorithm for calculating granger causality in artificial neural networks. All experimentation, all writing and all data analysis were carried out by me alone.

5f. I have been programming for a very long time (since I was 12) and cannot easily explain how I learned so much about computer science. The formal computer science classes offered at my school helped put standard names to practices and data structures I knew intuitively, but the content was not new to me. In doing this project, I tried to get out of my element and do something that I didn't have any training for. By completing this research paper I wound up learning more than I could ever have expected about the scientific process, scientific rigor, attention to detail, and technical writing.