The Beginning of Real Machine Intelligence

I have raised the question "What is real intelligence?" and then went on to refine this question to "What criteria constitute intelligence on the same level as we humans?" This is not an easy problem to address, primarily because not everyone means the same thing when referring to the quality of intelligence and it can be used to mean different things in different contexts. For example, if a person makes a lot of mistakes we might declare that they are not very intelligent. If a person is very good at their job, especially when that job requires learning something complex and some degree of mental difficulty, such as solving problems or managing a financial portfolio, we refer to that person as smart, astute, or intelligent. Yet, if I were to create a machine that is capable of doing that same task, managing a financial portfolio for example, most of us would hesitate to say that this machine was intelligent. We might

say it a good system, perhaps even advanced. We might say things like “this machine is very capable.” But to confer intelligence on a machine is not a leap most of us are comfortable making. So exactly what is it we are looking for before we can confidently pronounce a machine to be intelligent? I will attempt to lay out a few basic criteria. Whether these criteria are all necessary for intelligence I leave as an open question along with whether these criteria are sufficient for intelligence. For the moment, the goal is only to come up with some minimal and rather basic notions for consideration. In other words, where does intelligence begin? At what point are we willing to say that a machine possesses intelligence? On the following pages I present what I believe are the minimum characteristics of any intelligent entity. This will be the basis for discussing how those criteria apply to machines along with caveats, concerns, and other factors to consider.

Criteria for Intelligence

Before delving further into how machines might become intelligent it is helpful to define, or at least describe, what is meant by intelligence when referring to machines. I.J. Good, who worked with Alan Turing during World War II and is credited as the originator of the oft cited term 'singularity', described what he called an “ultraintelligent” machine as "a machine that can far surpass all the intellectual activities of any man however clever". Good and many others tend to focus not just on what constitutes intelligence but specifically on how and when we will know that machines will surpass the capabilities of human. Regardless of whether we are defining intelligence or ultraintelligence, the type of the criteria should be the same. This gives us the basis for one possible criterion for intelligence, viz. the ability to perform a task or activity as well as an average person. In Superintelligence: Paths, Dangers, Strategies Nick Bostrom states : “…we use the term

"superintelligence" to refer to intellects that greatly outperform the best current human minds across many very general cognitive domains." He goes on to suggest that it is helpful to decompose this notion into three categories of superintelligence: speed superintelligence, collective superintelligence, and quality superintelligence.

Speed superintelligence is defined as "a system that can do all that a human intellect can do, but much faster."

Collective superintelligence is defined as "a system composed of a large number of smaller intellects such that the system's overall performance across many very general domains vastly outstrips that of any current cognitive system."

Quality superintelligence is defined as "a system that is at least as fast as a human mind and vastly qualitatively smarter."

Each of these three definitions covers a different form of intelligence and is in fact the product of a different type of system.

Speed Intelligence

Speed superintelligence is quite simply that, the ability to do the things a human typically does, but much, much faster. It is important to point out that we are not just talking about doing something twice as fast or even ten times as fast. When we use the term speed intelligence we mean very, very fast Intelligence!

As humans, we are able to get faster through training and practice. We do this by optimizing our neural circuitry. The result is that our reactions become faster, allowing our performance to increase. This is most evident when we pair a well-trained athlete with an athlete of a relatively lower skill level. To the beginner, everything seems to be progressing at a breakneck pace, barely able to keep up. To the well-

trained expert, time seems to be moving at a slow, almost relaxed pace. We often say that the pro makes it all seem effortless.

This optimization through training yields an improvement by a factor of two or three or perhaps even more. In a machine, we will realize improvements by a factor of hundreds or thousands. In specific problem areas, we already have machines which are faster than humans at a factor approaching one million times.

Imagine a scenario where you are competing with a machine for a job at a bank. It might be a position processing loans or mortgages, or perhaps managing a portfolio of investments. Both of these jobs already leverage computers for accessing and processing in formation. But today we still rely on human loan officers and portfolio managers to make decisions about which loans or mortgages to approve, which stocks to buy, and when to sell them. But what if we had enough confidence to let the computers make all of these decisions on their own, without supervision? How much more quickly could these computers process loans or make buying and selling decisions on investments? These automated financial agents could easily make decisions and process transactions at a rate that would allow them to take the place of tens or even hundreds of their human counterparts. In fact, they are used extensively in the financial sector today and continue to become more pervasive. Automated trading has been used to accelerate trades for years and has now become available even to the individual investor.

From the perspective of speed and efficiency, machines clearly are more intelligent than humans already. They have become so good at what they do that they make their task seem effortless.

Collective Intelligence

Nick Bostrom has described collective intelligence as “joint problem-

solving capacity". There are may examples of problems that have been solved by more than a single individual and which probably could not have been solved by a single person. When a new drug is discovered and developed it is not done by a single person but by a team of people working toward a collective goal, sometimes in collaboration with other teams performing similar, related research. Projects such as the Human Genome Project was a massive research project that was only able to achieve its goal by the collective research of twenty teams across six countries over thirteen years. In our daily lives we often tackle jobs at work that require the cooperation of several people or large teams to solve problems and implement solutions. We are surrounded by evidence that our problem-solving potential is increased by combining the skills of multiple individuals. This increase in potential is realized for two reasons.

First, we have an additive effect. Think of a simple example where our goal is to identify all red objects in a warehouse. A single individual may be able to pick out and identify each object in 1 second. Therefore they can sort through 60 objects per minute, 3600 objects in an hour. But if we can enlist ten people to take on this task, assuming they are all able to sort through objects at the same rate, we can sort through 36,000 in an hour. In other words, we can accomplish the task ten times as fast! This is of course a rather simple uninteresting example, but this is the way many tasks are accomplished. The pyramids of Egypt were built this way, and many intellectual tasks are as well.

The second reason we can accomplish difficult tasks more easily with a group is due to the breadth of skills required to complete the task. We live in an increasingly complex world and the problems we are faced with are constantly increasing in complexity as well. As the tasks become more complex, solving them requires an ever-increasing array of problem-solving skills. While this phenomenon has been evolving for a long time, the first good example with a significant impact is found in Henry Ford's approach to building the automobile. Rather than relying

on a single individual or even a few individuals with the requisite skills to build a car, he broke the overall tasks into smaller tasks requiring a degree of skill that could be developed in the workers in a short amount of time. Today we see an even greater divergence of skills required in our world. Think of the number of people that contribute to the treatment of a patient during their stay at a hospital. Even going to a store to buy something sometimes involves several people. We ask where we can find what we are looking for and are directed to the correct department. Once there we can't find what we are looking for and have to ask someone who works in that department. We then ask them about some feature of the product, they don't know the answer and go find someone who knows more about that product. Finally, we go to the register where we pay for what we have selected. While this may not seem like something that exemplifies the height of human intelligence, it demonstrates how much of what we undertake relies on the knowledge of multiple individuals. This is a simple form of collective intelligence.

There is one more aspect of collective intelligence that is important to recognize. In the past we see examples of a brilliant individual such as Thomas Edison or Alexander Bell who were great innovators. But even these great names relied on the discoveries and knowledge gained from others, some their contemporaries and some from the past. More than any other species, humans have the ability to learn from their predecessors. This ability is referred to by Michael Tomasello and Steven Mithen as cultural learning. In its simplest form it can be seen as a child learning not to cross the street without looking for oncoming cars. It can be seen as the fundamentals of reading, writing, and arithmetic that we learn in our first few years of school. It is seen in our education where we learn the necessary skills for our careers. Although we don't often think of it, when we attend a year of school to learn a trade or skill we are learning skills and knowledge that took hundreds of years to accumulate. The long term effect of this type of knowledge transfer is incredible. It is what has allowed us to develop rocket ships

that fly to the moon and understand the complex system we know as the human body. It is what has led to the development of the computers that are so advanced, so fast, so intelligent – that they may soon surpass us in intelligence in its every form.

Quality Intelligence

The notion of quality intelligence, unlike speed intelligence, is difficult to define in measurable terms. It can best be described by example. Consider the case where the ability of one person to perform a certain task is clearly superior to another. For example, one person may be good at math while another just doesn't "get it". One person may have a natural ability to learn to play a musical instrument with ease while another finds learning the same instrument an exercise in frustration. One person may be good at recognizing patterns, such as constellations in the sky, whereas another has difficulty seeing them. In these examples the issue isn't how fast one performs the task, it isn't a matter of speed intelligence. It also isn't a matter of how much information we can process or how much we know, it isn't a matter of collective intelligence. This intelligence requires a specific kind or aspect of intelligence, it requires the ability to think or to process information in a way which is specific to a particular task. The need for developing very specific cognitive skills is often what differentiates one individual from another, particularly in highly skilled domains. Mastery of advanced skills require a particular and often unique approach. We often hear people talk about the ability to "see" a problem in a specific way. It is one of the key characteristics which differentiates top performers in science, business, or the arts.

Up to this point we have only considered abilities and domains that are well known to us. But what about abilities that are unknown to us because they are out of our reach? Bostrom suggests "… the idea of possible but non-realized cognitive talents, talents that no actual human possesses…". What these abilities might be we can only guess. Might a person or other entity be able to foresee the future? Not by literally

being able to see the future, but by observing the obvious consequences of the present, not evident to the comparatively less intelligent population around them. While predicting the future may sound like any one of many films that come out of Hollywood, we have readily come to accept that we can now predict the weather with a fair amount of accuracy, at least for a few days. As our ability to understand the systems that manifest themselves as wind and rain, hot and cold, we have learned to "see" tomorrow's weather before it happens. In medicine, we are constantly striving to understand the system of human physiology and the malfunctions that we know generally as "disease". As physicians, chemists, and other scientists increase their collective intelligence, it in many cases increases their quality intelligence. While the knowledge that arises out of collective intelligence may be necessary to discover the cause and ultimately the cure for a particular malady, it is not necessarily sufficient. For that leap in advancement, we need quality intelligence. We need to be able to "see" a complex system of biology, the human body, and its interaction with another complex system, the environment, in a particular way. It is this "seeing" that leads to an abstract concept we call "understanding". It is this understanding that allows us to see that a person is becoming ill before they show symptoms.

As humans and machines become more knowledgeable, more intelligent, and more powerful, the ability to "see" events before they happen will become more and more prevalent. In a competitive environment, the individual who is better at seeing what will happen next has a distinct advantage over everyone else. In the future, whoever can see more clearly into the future will be ahead of the rest of the pack. That individual will be one of the top people….or one of the top machines.