

The Future of Humanity

Roger Bishop Jones

Contents

1	Introduction	2
2	Discussion Plan	2
3	Some Big Things Happening Now	2
3.1	Personal assistants, Audio interfaces	2
3.2	Self-driving vehicles	2
4	Some Bigger Things Happening Soon?	3
4.1	The Information Economy and Machine Intelligence	3
4.2	Designer Babies and The Evolution of Homo Sapiens	3
5	Reasoning about the Future	3
5.1	Designing the Future	4

The Future of Humanity

1 Introduction

Over recent decades it has become more respectable for philosophers to think about the future. One symptom of this shift is the establishment of various institutes dedicating to research into the future, often multi-disciplinary, in which philosophers play a role. Two such institutes in the UK are the Oxford University *Future of Humanity Institute* (www.fhi.ox.ac.uk) and the Cambridge University *Centre for the study of existential risk* (cser.org).

The pace of change continues to accelerate, making it ever more difficult to predict the future, even on quite modest timescales. But without forethought, the risk of serious problems arising are exacerbated. How can it be done?

2 Discussion Plan

The plan is for a philosophical discussion which illuminates the issues, so the principal elements in the plan are topics for discussion. and the principal purpose of the notes is to provide background for the discussions.

I propose that the discussion be in three main parts:

1. Some developments now taking place.
2. Some more radical changes which might not be far ahead.
3. General considerations about how to think ahead.

3 Some Big Things Happening Now

3.1 Personal assistants, Audio interfaces

Computers have been accepting human dication for many years now, but the capabilities have been very limited.

The technology has now become sufficiently good that we can expect this mode of interaction with computers (mobile phones, tablets,...) to become very widespread, pretty soon. Smartphones already have this technology, in the form of Apple's Siri, Google Now, Microsoft Cortana, and the Amazon Echo.

Huge resources can be expected to go into making these interfaces as capable and intelligent as possible, since they will in effect be sales assistants. They will advise us on how to solve any kind of problem and on where to buy the things we need to effect the solution. People talking to computers at home and in public places will soon become as ubiquitous as the use of texting and other social media have now become.

3.2 Self-driving vehicles

Google has been working on self-driving cars since 2009, and has now clocked a million miles driven by driverless (but supervised) cars in trials in California.

Google talk of making such vehicles available to the public in the timeframe 2017-20.

The CEO Uber the internet car ride company has said that if Tesla can produce self-driving cars by 2020 he would buy all 500,000 of the expected production volume.

It has been projected that once self-driving cars become available, few people will want to but their own cars, since hiring on demand will be convenient and much cheaper than running a car. This would decimate

the automobile industry since only a fraction of present car numbers will be needed. Furthermore, the technology required for such cars is so different from existing car technology that its not certain that any of the existing global manufacturers would be among the winners, they might lose out to google, apple, tesla, or a completely new corporation.

4 Some Bigger Things Happening Soon?

4.1 The Information Economy and Machine Intelligence

Even without machine intelligence, the continuous (exponential) growth in computing power and the already pervasive reach of networks bringing that computing power into every crevice, paves the way for progressive automation of very large swathes of what was previously done by people.

Just as the industrial revolution moved most people out of agriculture into manufacturing, industrial automation is now shrinking the numbers employed in manufacturing, and increasingly large parts of the economy involve information only. Information industries, such as social media, grow very rapidly and use very little human labour. (*WhatsApp* had only 55 employees after 5 years when it was acquired by *Facebook* for \$20B. Facebook's CEO was a billionaire by the age of 23.)

All the talk is now of superintelligence, machines, becoming more intelligent than humans. This would of course turn the world upside down, with very difficult to predict consequences. Is it possible to reason about what will happen and to take precautions against the worst outcomes. Bostrom and the Future of Humanity Institute think so, but are their arguments and conclusions credible?

4.2 Designer Babies and The Evolution of Homo Sapiens

At a slower pace perhaps than advances in information technology synthetic biology continues to advance, and we can anticipate that in this century the way in which the human species (and the entire ecosystem) evolves will be completely transformed.

The process of biological evolution is usually described as involving random genetic variation couples with environmental selection, resulting in the proliferation of just those few variants which prove advantageous.

Homo Sapiens has interfered with this mechanism by developing medical science which ensures the survival and even the propagation of all but the most severely disadvantageous genetic variants. We will soon be in the position that variation arises not by some primarily random process, but by design. It will be technically possible for parents to select the genome of their children, and we may then move forward to "evolution by design", which will transform the species at lightening speed.

Where will this take us?

See: Designer babies at Future of Human Evolution (futurehumanevolution.com/designer-babies).

5 Reasoning about the Future

Oxford now has an interdisciplinary "Future of Humanity Institute" directed by Nick Bostrom, also a Professor of Philosophy at the University of Oxford. His latest book on "superintelligence" is creating a stir.

I have read some of his work, but find much of what he writes unconvincing. What does it take to reason soundly and convincingly about the future in the face of the rapid and radical changes which we are going through?

5.1 Designing the Future

The following quote has been attributed to Abraham Lincoln, and similar observations to many other since then:

“The best way to predict your future is to create it”

This is obvious in certain specific contexts, for example, when the construction of a building starts, if you want to know how the building will turn out, there is not much detail to be gleaned from inspection of the foundations, or observation of the work in progress. The best way to discover how the building will turn out is to ask the architect.

Where people are trying to mould the future, even if they are not expected to be entirely successful, knowing what they are trying to achieve is certainly an important element in forecasting what will actually happen.

By contrast, if we are considering dystopian scenarios, particularly dystopian outcomes which are technically possible but very difficult to achieve, you have to ask, “who will be trying to make this happen”, because it is not going to happen by accident.

When we consider the possibility of rogue artificial intelligence, it's a reasonable presumption that people developing intelligent machines will be designing the machines to do as they are told, and special effort would be required for autonomous *motivation* which it seems doubtful anyone would have a reason to undertake.