COMMERCIALIZING THE ROBOT ECOSYSTEM IN THE ANTHROPOCENE

Anders Sandberg, Future of Humanity Institute, Oxford University

Notes for a talk given at the Robotics Innovation Challenge, February 9, 2012.

THE BIG PICTURE

TECHNOLOGY

Let us begin with the big picture.

What is the purpose of technology?

Magic.

Yes, I mean this literally. We build machines in order to get what we want.

Here is a quote from a bona fide magician, the somewhat infamous British ceremonial magician Aleister Crowley: “Magic is the art and science of causing change according to will”.

Of course, he was more interested in supernatural methods of causing change, but he did acknowledge that tools and machines were magic too.

When we consider our technological dreams, we see that they aim at making the world into a kind of fairy tale world: wish for something in the right way, and it will become real.

Here is a picture of our modern fairy tale world. You might not recognize the enchanted forest, but it is there. Our stone age ancestors dreamed about a life where they did not have to do grueling hard work to survive, where food was easily accessible, where you could get as much fat, sugar and salt as you wished, where travel was easy, people did not have to fear predators or violence, where they were healthy and lived to a ripe old age.

Of course, many fairy tales do warn about wishes coming true. We do not always wish for good or wise things.

THE ANTHROPOCENE

Thanks to our technological magic we have become the most common large mammal on the planet. As a species we can be found in all ecosystems, sometimes by adapting to them but often simply be suborning the ones in existence or creating new ones that suit us.

This is the anthropocene, the geological era where the Earths geological and ecological cycles are dominated by human activities. This is at the same time great news and very bad news.

So, what are the properties of this new world we are creating, sometimes by mistake?

We have been changing things for a long time, but now is when the exponentially increasing scale of the changes have become so big that they actually change the face of the planet.

- We are the major form of erosion of rock
- We are using about a tenth of the energy processed by today's biosphere.
- 83% of the earth's land surface is influenced directly by human beings.
- We are fixing 190 megatons of nitrogen per year, about the same amount that the entire biosphere fixes.
- We appropriate 25% to 40% of the total net primary productivity of the planet for our use.
- Humanity have eight time the mass of all wild land vertebrates (about 40 million tons), and about the same biomass as all the fish and whales in the ocean. Domesticated animals have a biomass of roughly 100 megatons of carbon. The biomass of our animals is about 20 times the mass of all wild vertebrates on land, and 50% larger than the mass of all vertebrates in the ocean.
- Only 10% of the land area is more than 48 hours from a large city.

The world is essentially becoming domesticated: nature is as much a product of human decision and action as something natural.

Technology in the anthropocene is just as much about managing the human world as the natural, physical world.

THE USEFULNESS OF ROBOTICS

Back to technology as magic. To cause change according to will requires interpreting the intentions and effecting them.

This is where robotics comes in.

While non-autonomous tools are useful, autonomous devices are potentially even better at converting intentions into actions since we might not always know how to do that, or have the right physical abilities… beside our laziness. They can even act in our interest when we are not around, so that we do not need to formulate intentions all the time.

We might even desire robots for another reason: we are social creatures, and tend to view even the inanimate world as a social place. Building autonomy into our environment might make it fit our preconceptions better – we are literally adding the spirits and talking animals to the magical forest.

But, as I will argue a bit ahead, we do not necessarily want totally general robots. There are practical reasons to have different kinds of robots in different situations: benefits of specialization of software and hardware, economics, issues of trust and reliability. So the magical forest will not be filled with identical very general robots, but an ecosystem of different kinds of entities.

So, where do I think we have the big challenges and opportunities given this perspective?

THE BIG TRENDS

THE WORLD

- Demographics: a bigger and ageing population
- Urbanisation: most of us live in cities
- Wealthier: we are getting richer
- More globalized: we are getting more integrated
- Climate change: the geosystem is changing together with the technosystem
- Resource change: we need more and different resources

SOCIETY

- Security: people are desiring more safety, more risk averse
- Education: we need more education than ever (longer lifespans, accelerating change)
- Renegotiation: change means new rules have to be negotiated between classes, generations, subcultures, cultures etc.
- Democratization: people want individual control over their own lives and are increasingly willing to pay or fight for it.

TECHNOLOGY

- Several ”Moore’s law”: technology accelerates in many domains
- Converging technologies: micro, nano, info, cogno, bio, identity…
- Big data: we are acquiring enormous amounts of data, learning how to harness it for science and business.
- We are wirelessly connected 24/7
- Internet of things: more and more objects are smart and connected.
- Surveillance: more and more devices and data that allow observation and surveillance. Identity technology.
- Robot-ready: more and more devices are not just smart, but they have built in sensors and actuators that allow them to become robots.
- Low power, cheap electronics (because of wireless, smart objects and resource)
- Automation: people become more expensive, machines smarter and cheaper
- Customization: Due to individualization, smarter and cheaper technology

INTERSECTIONS

These trends intersect in interesting ways for robotics.

HEALTH

An older population will need support for work and especially health care – they will have chronic diseases for longer. The health care sector is growing enormously, and most western societies have trouble paying for it.

Telemedicine looks like one way out: allow remote checkups. At first merely services and sensors, but actuators will follow. Lab-on-a-chip and labtops might outsource testing to primary care or the home.

From an ethical perspective automating medicine is important. The cost of gadgets and pills tends to go down exponentially, while services (doctors, nurses) remain constantly costly. Software also allows encapsulation of skills and the wide copying and distribution of them, bringing down the cost of rare skills. While this is nice for health care budgets, it is also necessary for making it more widely available in the world. Given the moral importance of providing health and that we regard healthcare inequalities as among the worst failings of the current world system, it would seem we have an imperative to automate health care as much as it is possible.

Actually implementing automation and robotics in health is another matter. It is an environment dominated by large, rigid organizations with people (patients, professionals and politicians) with strong views on how it should be done. It might be that the imperatives of cost help force through robotics, but it could equally well be globalization: emerging markets inventing their own approaches to health care with less entrenched interests. So watch out for the Chinese nurse applications and Brasilian micropharmacies!

TELEPRESENCE

Concerns about climate change and resource use collide with increasing globalization. We want to interact globally, yet much travel is deeply wasteful. Concerns about pandemics and security might also make travel harder or restricted. Telepresence might be a solution. Skype and videoconferencing are still limited to merely communication, not action. Might remote controlled bodies be the future? Remote controlled presence might take off as augmented and virtual reality finally become easy thanks to computer game-driven fast graphics processing, broadband communications and consumer electronics lower equipment prices, and someone figures out a suitable user interface.

Telepresence robotics will depend on how people frame the use: is it the green thing to do, a high-status activity that sends social signals about caring and formality as well as a personal meeting? Or is it a cheap way of getting remote or outsourced labor? Is there a whiff of teledildonics or gadgetry, or cool space age engineering? These factors are more about business plans, marketing and design than engineering, but will determine the success far more than the actual capabilities.

If telepresence takes off it will help robotics, since it will spread hardware that can be put into “auto” mode. Conversely, a spread of consumer or office robotics might enable telepresence – just tale control over a suitable robot.

SMART CARS

Rapid technological progress on control and autonomy of cars (when Google lobbies state governments to change laws, it is a pretty big sign). Cars are already fairly “robotics ready” thanks to internal sensors, servos and controllers. But autonomous cares are a good demonstration that to succeed with an innovation it is not just enough to have the technology: it needs to be fitted in with the culture, economy and legal frameworks.

There are also interesting challenges when using such systems in a real world environment. Consider the problem of traffic lights – the car should stop for red lights, yet there are lots of red lights in the street environment (other cars, shop windows). Is the solution to build a better map of where they are? Mandating that they emit wifi or lifi signals? Much better smarts in the car? Or rather, all three (and more), looking for multiple mutually supporting data streams to make robust decisions.

A key challenge: robots that act in the real human world need behavior that is trustworthy, predictable, and can deal with messiness. Robustness is about handling uncertain and contradictory inputs, making use of whatever resources are available, and failing in graceful ways. Being trustworthy and predictable requires either some understanding of human expectations or ways of making the robot transparent.

QUADROTORS AND OTHER DRONEBOTS

Right now the net is abuzz with quadrotor robots. Leaving aside the hype, they represent a very interesting confluence of technologies and ideas.

Hobbyists and researchers are using 4 rotor choppers for a variety of purposes – monitoring, swarm control, building structures, etc. In many ways this is a reflection of the booming business of drones, in turn powered by the globalization and high human capital cost of warfare.

The quadcopter devices appear to have emerged thanks to the development of ever more capable toys, in turn empowered by better batteries (for portable gadgets), accelerometers (MEMS and cellphones) and cheap manufacturing. Some use smartphones as controls – already present sophisticated computing devices able to act as wireless sources. By adding a camera to the quadcopter it already becomes a drone – old regulations such as ITAR about drones mean very little against toys. Amateurs also use them in more or less subversive ways, such as the “occucopter” intended to help demonstrators by monitoring the police – and employing a distributed control system to make them less vulnerable to one operator being seized.

More radical ideas involve Matternet, using bigger quadcopters to transport small packages (1 kg) over a few kilometers, potentially integrated into a literal packet network for bridging rural areas.

This is a demonstration of the democratization of robotic technology. By using inexpensive parts in new ways, sharing the methods online, pooling expertise and ideas, people are developing robotic infrastructures based on their own ideas. The maker movement is trying to give people control over inexpensive means of physical production, open source aims at something similar for software. While the quality might be beatable by commercial or institutional projects, these grassroots projects act as a low end competition that will force the others to up their game – or they will be beaten by the free products.

They also provide something more important: ideas and experimentation. By lowering the threshold to entry they unleash a torrent of individual ideas, customized solutions and tests of wild possibilities. By sharing the results they speed up the accumulation of knowledge and skill.

This challenges traditional business models. Who is the designer, who is the buyer? When customers begin to share modifications of your robot, do you attempt to control the ecosystem or do you welcome the free development?

There is another important aspect of open designs: the feeling of control and trust. We all desire control over our lives, and chafe or fear when we do not experience it. Technology that is remote, handled by a priesthood and not encountered will be regarded as inessential and/or risky. Technology that is close and open for tinkering – even if that never happens – becomes part of life. In order to have buy in people must be able to play with the technology, to explore it, to make it theirs.

THE DILEMMA OF AUTONOMY

Safety is a function of intelligence:

Simple systems are fairly safe, and moral responsibility solely on the user.

Dumb but autonomous systems are more risky, and the responsibility becomes complex. The smarter the system gets the safer it might become, but validation and responsibility becomes very tricky. And once systems reach human level intelligence risks might go up significantly again in a nondeterministic manner.

This is the AI principal agent problem:

- Stupid systems misbehave because we cannot explain what we want (and they are bad at doing it)
- Smarter systems misbehave because we misunderstand each other
- Very smart systems misbehave because different goals from ours (and they are good at achieving them)

Programming in goals is a form of communication. Making the goals and states of robots transparent helps build trust.

CONCLUSIONS

Robotics help making humans even more powerful: the trick is also to get our intentions to be more accurate and *better*.

It is not just that we *can* develop certain forms of robotics: we *ought* to. Anything that helps solve the big problems or improve human wellbeing belongs here. Smart technology can have an enormous impact because it can be spread widely and relatively cheaply.

But our devices will have to function in an environment filled with humans and human thinking. Designing robots for the anthropocene means designing them so that they function in a human-shaped world. This is hard: humans have the most complex social behavior of any species on the planet, and we tend to expect anything that joins our social space to obey our own weird rules. Humans often have unclear intentions or stupid goals. And we are protective of our existing institutions and networks.

The world is increasingly "robot ready", not because our software is getting much smarter but because we are putting sensors, identifiers, connectors, and actuators everywhere. Robots might lose their shape: the robots themselves are distributed systems, with sensors, actuators and control distributed between multiple devices. This allows them to make use of the already existing systems: smart devices nearby can act as sensors: cloud robotics.

Together with the previous point, robustness suggests that data fusion may become even more important. Taking inputs from anything available to decide what to do or just what is going on.

Cheap and open: many forms of robotics are becoming open endeavors, with the borders between producers and customers, designers and users, blurred.

Robustness in interpreting human intentions, implementing them and behaving is a tough goal. It can be helped by making the device transparent or non-anthropomorphic in the right way (for example animal-shaped). But a key variable is to add the right level of autonomy and smarts: not everything has to be able to discuss the weather.

The anthropocene is a magical era. It is up to us to design it so that it is a good world to live in for everyone and everything.