

The Coming of 'Killer Robots'

Written by Christopher Coker

This PDF is auto-generated for reference only. As such, it may contain some conversion errors and/or missing information. For all formal use please refer to the official version on the website, as linked below.

The Coming of 'Killer Robots'

<https://www.e-ir.info/2013/05/16/the-coming-of-killer-robots/>

CHRISTOPHER COKER, MAY 16 2013

There will be one million robots toiling away in Foxconn's Chinese factories by 2015, a mighty mechanised army, tirelessly snapping together the world's iPhones, Nintendo games consoles and DVD players. It is a major threshold to cross in a country where cheap labour has been the key to dramatic economic growth and which has placed China in pole position to overtake the United States quite soon. In an era of systemic wage and benefits inflation and increasing instances of labour unrest Foxconn thinks that the ideal worker is no longer a poor migrant but a machine that never goes on strike and never demands a wage hike. Welcome to the future.

What is true of industry is also likely to be true of war; indeed the two have been intimately linked since the industrial revolution. That is why we are being warned that 'killer robots' are heading our way. A few years ago a high level study of the US Army, *Star 21: Strategic Technologies for the Army of the 21st Century* concluded that whilst the core 20th century weapons system had been the tank, in the 21st it was likely to be the 'unmanned system'. The study even predicted that robots would be running and walking on the battlefield by 2020. Another study appropriately entitled, *Unmanned Effects: Taking the Human out of the Loop* (2003) envisaged that five years after that autonomous robots would be fully networked and integrated on the battlefield. It hasn't happened and it won't for some time. Progress in robotics is painfully slow.

But if the post-human future may be further off than we think we need to ask questions about where we are heading. The future is not a destiny, it is a choice, and it is that knowledge which is the great curse of the modern consciousness, as well as a source of hope. And there is no doubt that robotics is the future of war. Gordon Johnson, the Unmanned Officer Team Leader for Project Alpha, envisaged some time ago that tactical autonomous combatants (TACs) will be operating with minimal human supervision. At the operational and strategic level war will remain 'manned', but on the battlefield, soldiers will find themselves increasingly co-existing with intelligent machines and it is the soldiers, not the machines which may be rendered redundant. "They don't get hungry. They're not afraid. They don't forget their orders. They don't care if the guy next to them has just been shot. Will they do a better job than humans?" For the technophiles, there is an inbuilt redundancy in the very question.

Robots were first used extensively in the non-conventional phase of the Second Gulf War (2003-7). Within five years of the initial invasion the US military was deploying 5000; the first armed robot was in operation in Baghdad in 2006. As technology develops, the robots we are discussing will take many shapes, very few of them are likely to be human. Only excessive anthropomorphism would lead one to conclude that a robot should look something like the Hollywood Terminator, or the human facsimiles represented in another film, *I, Robot*. With the development of nano-technology, some may 'swarm', others may look like tractors, tanks, even cockroaches or crickets. All sorts of shapes and locomotion styles are being tested. New research projects include robots that can fold themselves, fly and crawl, walk uphill and roll down. Some roboticists are even looking to the humble amoeba for inspiration – the result is the chembot made up of particles that are quite stiff when compressed but, given space, flow like liquids thus allowing it to enter any space no smaller than its fully compressed state, more or less regardless of the shape of the space in question.

Robots, in short, are positively protean in their possibilities and they won't always be called by that name. 'Robo' and 'bot' are two different affixes and of the two the latter seems to be easier on the ear. That the two should sound different isn't so surprising given the linguistic 'principle of contrast'. When we learn a language we expect that two

The Coming of 'Killer Robots'

Written by Christopher Coker

words should never be exact synonyms. So we speak of 'emotional baggage' but not of 'emotional luggage'. Robo is the name we seem to give more menacing systems – like Robocop. Bots are nicer, like nurse-bots and computer programs like searchbots and spam bots and chatbots which were developed to engage in more-or-less human sounding conversation. So the future may see warrior bots and military bots, and a host of others, less threatening in the imagination than the robots of today.

Robots are on the march, to use a military metaphor, and there is no going back. We have committed ourselves to reducing the human space of war still further. At the very heart of this desire is a wish as old as war itself: to take out of the equation such existential elements as courage, fear, cruelty and remorse, the traditional emotional features which have made war such an intensely human activity. War, we must remember, has only been rendered humane even at its most bloody, because of the human values, capacities and emotions which infuse it. But will we be able to render it more humane still, not only for ourselves but others?

Our own age, of course, has a particular predisposition for discovering historical turning points, or new ways of seeing and asserting the coherence of the world. Best-sellers market these axial moments on which everything turns; they fix our eye on the future. "Every age has its eye pasted to a key-hole," wrote Mary McCarthy, and she was right. Trying to second guess the future has become a mark of the post-modern world in which we live. The coming of robots is culturally a watershed, a Rubicon which we have decided to cross, as important in many ways as the arrival of gunpowder, and it will be far more radical in its implications and consequences however long it takes for robots to go from being automatic to autonomous (capable of making decisions themselves). The fact that we have already begun thinking through the implications shows that we are fully aware of this.

The robots the US military has built so far are simple devices, a blend of artificial intelligence which allows them to reason, and mechanical engineering which allows them to perform physical tasks informed and directed by reason. As currently defined, robots exhibit three key elements or functions:

- 1) *Programmability* – they have computational symbol-manipulative capabilities that a designer can combine as desired.
- 2) *Mechanical Capability* – they can act in their environment rather than merely function as an information processing or computational device.
- 3) *Flexibility* – they can operate using a wide range of programmes.

The first 'function' makes a robot a computer; the second a machine; the third a computer-enhanced machine that can respond to external stimuli. And these responses are much more complex than they would be if we were to use just mechanical or electro-mechanical components alone. Robotics adds a new element: complexity. The fact that in theory they may soon be able to learn and adapt to their environments suggests that behaviour will emerge over time. When they begin to learn independent of our programming they will be well on the way to achieving autonomy.

What is an autonomous system? It is one that is distinguished by several characteristics: self-repair, self-maintenance, self-improvement (it learns), and self-reproduction (the biggest challenge of all). Until recently, machines could not do all of these tasks, but some were capable of at least one. An aircraft, for example, can steer on auto-pilot without human interference, but not repair itself in flight when things go wrong (systems failure). A communications network can repair itself but not reproduce itself. Computer viruses can reproduce themselves in ways that even their programmers cannot anticipate, if programmed with evolutionary algorithms but they can't learn as they 'evolve'. One day, however, robots will be able repair, reprogram and maintain themselves without human involvement.

Although that day is some way off, machine life is already taking on a life of its own. John Smart, a development systems theorist, argues that human generated innovation is 'trending down' at the same time that technology-generated innovation is rapidly increasing. In other words, what we have been witnessing for some time is a reduction in human-initiated innovation which we have failed to register, not because machines are taking over, but

The Coming of 'Killer Robots'

Written by Christopher Coker

because we ourselves are becoming increasingly integrated with the machines we design. Smart insists that all the crude indicators on which we rely: Moore's Law (processing power); Gilder's Law (bandwidth increase); Cooper's Law (wireless band width); Kurtzweil's Law (price performance computation) continue to suggest that innovation is increasing exponentially. But they also suggest that human beings are catalysts, no longer controllers of the process. One example is the series of innovations required to make something we take for granted – a gasoline/electric hybrid auto like the Toyota Prius. To any outsider, including its owner, it looks pretty much the same as other models. Yet it is radically more complex, and many of the innovations are not the result of human thought but the computations done by the technological systems involved (CAD-CAM Programs, Infrastructures and Supply Chains).

But this is not the same as designing machines that can think for themselves and with whom eventually we may co-evolve together. When it comes to making robots autonomous we rely on two quite different engineering approaches. The first is top-down, taking its cue from mathematics. Mathematics start with an axiom and applies rules of logical inference to transform that axiom into the desired statement. Some in the AI community have adopted a similar approach called 'means-ends analysis'. A problem (the initial state) is set which requires resolution; we start from this starting point using information about the problem to be solved (data/premises); we set a goal (the end state) that we wish to attain; and we agree a set of operations that can turn the initial into the end state.

The bottom-up approach is very different and is inspired by what scientists call 'connectionism' which takes its cue from the human brain and the infinite number of connections which neurons make within it. A neural-network type of machine which mimics as closely as possible the human brain may develop intelligence in the same way that children do: by observing the world around them and using their observations, together with the instruction they receive from their parents or teachers. Scientists call it 'subscription architecture'. Take the attempt to get robots to walk, one of the most difficult challenges of all. Engineers do so by ensuring that each sensor (or leg) sees the world in a different way from the others. Instead of building a coherent system of the world (as we do in our own lives) the robot learns what it needs to walk across a room by interacting with its environment. In other words, its behaviour (in this case, walking) emerges over time. The point is that each leg (learning independently) is eventually able to co-ordinate its actions with others and navigate in a three dimensional space without a head.

In time, cognitive abilities may even emerge, too. Consciousness may even evolve. Consciousness does not need a sophisticated representation of the world, only a reliable interface with it. Cognition does not require logic to get it going. In human beings it is an operation of the nervous system (in the case of robots, sensors). The goal of Artificial Intelligence is to evolve a consciousness very similar to our own; the goal of *artificial life* (which I have been describing) is much less anthropomorphic. It aims to evolve intelligence within the machine through pathways found by the machines themselves. Whether silicon life will evolve is uncertain but if it does its evolution will be different again from that of carbon-based life forms such as human beings. As such we cannot *know* how it will turn out because although designed by us, it will evolve independent of our own programming.

Some scientists question whether computers will ever be able to think for themselves. Many more believe the day will come when robots will be able to share their thoughts with us. My own bet is that this will happen within the next 30 years, possibly earlier than that. But even that time frame is short by historical standards which is why we need to ask questions now. We need to be vigilant about what we construct. We must be very careful to understand what is involved and where we may be heading which is why Human Rights Watch produced its report on 'Killer Robots'. They are coming to a theatre of war near you, and they may arrive sooner than expected. It was Richard Smalley who wrote that when a scientist claims that something is possible he usually underestimates how long it will take; but if he claims it is impossible he is almost certain to be proved wrong.

—

Christopher Coker is Professor of International Relations at the London School of Economics and Political Science. He is the author of *Warrior Geeks: how C21st technology is changing the way we fight and think about war* (Hurst, 2013).

The Coming of 'Killer Robots'

Written by Christopher Coker

About the author:

Christopher Coker is Professor of International Relations at the London School of Economics and Political Science.