



Center for a
New American
Security

Ground Robotics

Preparing for Disruptive Change

Paul Scharre

Director, Future of Warfare Initiative

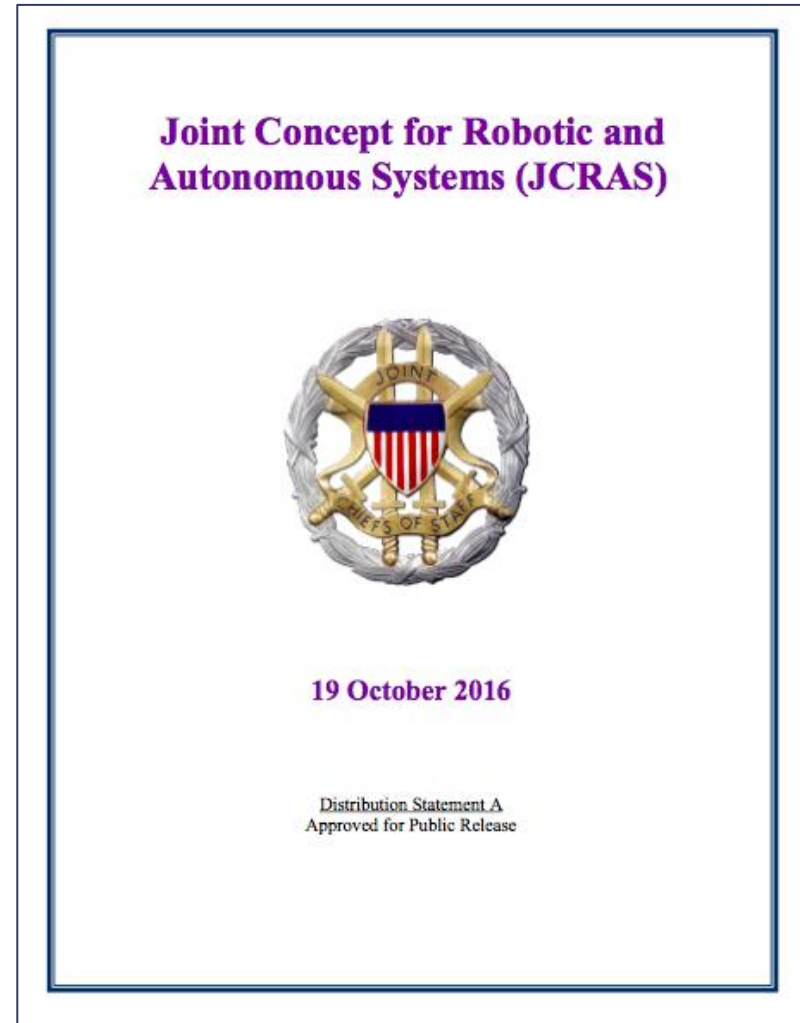

The Accidental Robotics Revolution

- Thousands of air and ground robots (“unmanned” or uninhabited systems) deployed to Iraq and Afghanistan
- Advantage of removing the human from the platform:
 - Take more risk
 - Performance advantages (speed, maneuverability, endurance, stealth, etc.) – this advantage diminishes as size of the vehicle increases
- Major downside of uninhabited platforms:
 - Lose the most advanced cognitive processing system on Earth: the human brain
 - Dependent on communications links to remote human controllers or onboard automation

Centaur Warfighting

- Solution: Human-machine teaming! (Third Offset)
- Physical teaming between inhabited and uninhabited vehicles
 - Leverage large numbers of expendable robotic vehicles
 - Human-inhabited vehicles forward in the fight for command-and-control to “quarterback” the fight
- Cognitive teaming between humans and automation
 - Leverage advantages of automation in speed and precision
 - Value of human judgment and adaptability in ambiguous, complex situations

Robotics and Autonomous Systems



Ambitious Near- and Mid-Term Goals

- “Soldiers operating dismounted for long periods will shift physical burdens to RAS platforms that provide a mobile power source and carry equipment, weapons, ammunition, water, food, and other supplies.”
- “As autonomous off-road technology fully matures, the Army will not wait for perfection in off-road navigation and tactical, inferential decision-making software before fielding autonomous prototypes for testing.”
- “The Army introduces exoskeleton technologies that lighten the Soldier load and allow for increased Soldier protection (armor) during close combat.”

Long Term Vision

- “Expendable RAS platforms will provide commanders the ability to take operational risks previously unimaginable with solely manned formations. Machines will take the place of humans maneuvering through the most dangerous avenues of approach and will make contact with likely threats without costing commanders valuable Soldiers. With less human exposure to hazards, the risks inherent with deception operations, penetrations behind enemy defenses, and exploitation and pursuit operations become less costly, giving commanders greater options and more reliable freedom of maneuver.”

Is the Strategy Executable?

- Is the strategy resourced?
 - Is the Army investing enough in ground robotics?
 - Is the Army going to field exoskeletons in the 2020s (4+ years from now) with no active program today?
- Is our acquisition system agile and flexible enough to stay ahead in a rapidly changing field?

Are we ready for what is coming?

Challenge Our Assumptions

- Are ground robots for maneuver really in the far future (2030-2040)? Are we ready if technology matures faster?
- Will robotic systems really always be used only to augment or support human formations? Are we being visionary enough? Or is this the equivalent of “tanks support infantry”?
- Will we be able to overcome internal cultural resistance to certain applications of ground robots?
 - Armed ground robots: “the Army will introduce unmanned combat vehicles”
 - Casualty evacuation: “Future unmanned systems assist in enabling CASEVAC”

What Will Adversaries Do?

- “Another factor influencing the essence of modern means of armed conflict is the use of modern automated complexes of military equipment and research in the area of artificial intelligence. While today we have flying drones, tomorrow’s battlefields will be filled with walking, crawling, jumping, and flying robots. **In the near future it is possible a fully robotized unit will be created, capable of independently conducting military operations.**
- How shall we fight under such conditions? What forms and means should be used against a robotized enemy? What sort of robots do we need and how can they be developed? Already today our military minds must be thinking about these questions.”
 - General Valery Gerasimov, Russian Chief of the General Staff



НАШ ЭКСКЛЮЗИВ



© Rostec State Corporation



The Artificial Intelligence Revolution

The AI Revolution

- “The Fourth Industrial Revolution” (World Economic Forum)
- “Next Industrial Revolution” (Bank of America–Merrill Lynch)
- Kevin Kelly:
 - “AI will enliven inert objects, much as electricity did more than a century ago. Everything that we formerly electrified, we will now cognitize. ... the business plans of the next 10,000 startups are easy to forecast: *Take X and add AI*. This is a big deal, and now it’s here.” (*Wired*, 2014)

How big?

- Bank of America–Merrill Lynch predicts by 2020:
 - \$153 billion market for AI-enabled technology, including:
 - \$83 billion for robotics
 - \$70 billion for AI-based analytics
 - With an estimated \$14-33 trillion creative disruption impact annually
 - \$8-9 trillion in cost reductions in manufacturing and health care
 - \$9 trillion cuts in employment costs due to AI-enabled automation
 - Manufacturing labor costs cut 18-33%
 - \$1.9 trillion in efficiency gains due to autonomous drones & cars
 - Productivity boosted 30% in many industries
 - 47% of jobs have the potential to be automated



Who is betting big on AI?



Microsoft



What are we talking about?

Autonomy

Automation

Artificial intelligence

Machine intelligence

Machine learning

Robotics

What's everyone so excited about?

Deep Learning

Deep learning

Cheap parallel processors (GPUs)

+

Big data

+

Deep neural networks


+

A little bit of black magic

=

Machine intelligence

Deep neural network (DNNs)




Credit: Michael Nielsen

Just a few of the things AI can do

- Chess
- Jeopardy
- Go
- Atari
- Poker
- Object recognition
- Facial recognition
- Recognizing human emotions
- Stock trading
- Driving
- Accounting
- Medical diagnoses
- Programming subway repair schedules
- Categorizing song genres

Today's state of AI technology





Center for a
New American
Security

The Terminator Dilemma

LETHAL AUTONOMOUS WEAPONS

Paul Scharre
Director, Future of Warfare Initiative

- “[T]he notion of a completely robotic system that can make a decision about whether or not to inflict harm on an adversary is here. It’s not terribly refined. It’s not terribly good. But it’s here....”
 - GEN Paul Selva, Vice Chairman of the Joint Chiefs of Staff, August 2016
- “If our competitors go to Terminators and we are still operating where the machines are helping the humans and it turns out the Terminators are able to make decisions faster, even if they’re bad, how would we respond?”
 - Robert Work, Deputy Secretary of Defense, May 2016
- “We have to think about what autonomous kinetic options really look like ... We need to understand and know that it doesn’t necessarily need to happen, but we also have to put the options on the table.”
 - Melissa Flagg, DASD Research, March 2016
- “Others are going to do it. They are not going to be as constrained as we are, and we’re going to have a fundamental disadvantage if we don’t.”
 - Frank Kendall, USD AT&L, August 2016

What is the Terminator Dilemma?

- The basic technology to build autonomous weapons that could select and attack/engage targets on their own is here today. It isn't very sophisticated, but it is here.
- Do we build them?
 - On the one hand, it probably seems like a bad idea for a variety of reasons (law, ethics, safety, risk...)
 - On the other hand, our enemies are unlikely to be so concerned. Could autonomous weapons give a decisive advantage to the enemy? If so, can we afford to fall behind? The job of the military is to win wars.
- How should we approach this problem?

Wait, Terminators? What are we talking about?

- Autonomy
- Automation
- Levels of autonomy
- Autonomous weapon
- Semi-autonomous
- Supervised autonomous
- Fully autonomous
- Artificial intelligence

This is a confusing mess

What is “autonomy”?

- At its core, **autonomy is the ability of a machine to perform a task on its own.**
 - Obviously, this can take a number of different forms. A Terminator is not the same as a toaster. There are degrees of “autonomous.”
- What is confusing is that we use the same word – autonomy – to refer to *three completely different concepts*:
 - Human-machine command-and-control relationship
 - Sophistication of the machine
 - Task the machine is performing
- There is not a single spectrum of autonomy – there are 3 different dimensions to autonomy!

1. Human-Machine Relationship

Machine Freedom



Fully Autonomous - aka "Human out of the loop"

No ability for human to intervene in real time

Supervised Autonomous - aka "Human on the loop"

Human can intervene in real time

Semi-Autonomous - aka "Human in the loop"

Machine waits for human input before taking action

2. Sophistication of the Machine

"Automatic"
Simple,
Threshold-based



"Automated"
Complex,
Rule-based



Patriot Air
Defense
System

"Autonomous"
Self-learning /
Evolving



Nest learning
thermostat



Self-learning
robot

"Intelligent"
Human-level
cognition of a
problem




Spectrum of increasing "intelligence"

More capable ... but may be
harder to predict specific actions

3. Task being performed

- The most critical dimension of autonomy is the task being performed.
- Both a landmine and a toaster are simple automatic systems. Different task.
- A system is autonomous *with respect to what task?*
- No such thing as “fully autonomous” with respect to all tasks.

What is an “autonomous weapon”?



Do autonomous weapons exist?

- Human-supervised autonomous weapons
 - At least 30 countries have defensive human-supervised autonomous weapons, such as the Aegis or Patriot.
 - Limited use: used to defend human bases or ships, anti-vehicle, human supervised, humans co-located with system
- Fully autonomous weapons
 - Israeli Harpy drone (anti-radiation loitering munition). Sold to India, Turkey, South Korea, and China. China reported to have reverse-engineered their own variant.
 - Experimental U.S. systems (cancelled): LOCAAS, Tacit Rainbow

Why build autonomous weapons?

- Lots of advantages for incorporating autonomy into weapons, but there are advantages to keeping humans in the loop too.
 - For the foreseeable future, no AI will have the breadth, robustness, and flexibility of human cognition.
- So why take the human out of the loop?
 - Speed
 - Loss of communications (e.g., UxV in comms-denied environment)

Considerations

- Law
 - Is it legal?
- Ethics
 - Is it ethical?
- Risk
 - Is it safe?

Legal

- There is nothing in the laws of war that mandates human decision-making vice machine decision-making
- The laws of armed conflict cover effects on the battlefield (e.g., proportionality, discrimination, precautions in attack). If a machine can be used in a manner that meets these criteria, then it can be used lawfully.
- However ... one important asymmetry between people and machines under the laws of war is that machines are not legal agents. Humans are bound by the laws of war. Robots / autonomous systems are not combatants.

Ethical

- Sometimes what is legal and what is *right* are not the same. (The law may permit what many of us might perceive to be unethical actions.)
- Some legal decisions require value judgments that do not have a clear answer. For example, how much collateral damage is acceptable?
- Military professional ethics – What is the role of the military professional in decisions about the use of force? Can s/he off-load that decision to an engineer?

Risk

- How much do we trust an autonomous system? How much should we trust it? How much confidence should we place in our test & evaluation processes?
- What is the probability of unexpected behavior in a real-world combat environment with enemy adaptation?
- What are the consequences if/when the system fails?
 - Fratricide
 - Civilian casualties
 - Unintended escalation
- Other factors: interactions with adversary autonomous systems; hacking; speed; time; damage potential.
- Ex. 2003 Patriot fratricides, 2010 Wall Street flash crash

International Debate

- Over 60 non-governmental organizations as part of a Campaign to Stop Killer Robots have called for a legally-binding treaty banning autonomous weapons
- Over 3,000 AI and robotics experts signed a letter in 2015 calling for a ban on “offensive autonomous weapons beyond meaningful human control”
- For the past three years, nations have discussed autonomous weapons in the United Nations Convention on Certain Conventional Weapons
- Discussions may move to a more formal Group of Governmental Experts next year, but there is currently no momentum towards a treaty. Only a handful of states (and no major powers) have said they support a ban.

The decisions ahead

- It's not about whether or not we build "Terminators"
- Real questions:
 - If we have a loitering munition or a UxV that is operating in a comms-denied environment, how much freedom (autonomy) do we want to give it to attack emerging targets of opportunity?
 - How much prior information do we expect military commanders to have about specific targets for attack? How much specificity about target selection can they delegate to a machine?
 - If we have a UxV in a comms-denied environment and it is attacked, do we want it to use force to defend itself? What about preemptively?

Role of humans vs. automation

- What role do we want humans to play in use-of-force decisions? If we *could* automate everything, what decisions would we still want humans to retain, and why?
- VCJCS: “one of the places that we spend a great deal of time is determining whether or not the tools we are developing absolve humans of the decision to inflict violence on the enemy. And that is a fairly bright line that we’re not willing to cross. ... it is entirely possible that as we work our way through this process of bringing enabling technologies into the Department, that we could get dangerously close to that line. And we owe it to ourselves and to the people we serve to keep it a very bright line.”