Penn State Journal of Law & International Affairs

Volume 7
Issue 3 *Symposium Issue*

April 2020

Autonomous Systems & International Norms

Charles Dunlap

Richard Jordan

Elsa Kania

Michael Klare

ISSN: 2168-7951

Penn State
Journal of Law & International Affairs

2020 SYMPOSIUM ISSUE

AUTONOMOUS SYSTEMS & INTERNATIONAL NORMS

Moderator: Tiyanjana Maluwa
Panelists: Charles Dunlap, Richard Jordan, Elsa Kania, and Michael Klare

Tiya Maluwa: Good morning. My name is Tiya. Tiya Maluwa, and I'm going to be moderating this panel. This panel is going to focus on Autonomous Systems and International Norms. We have four panelists, Elsa Kania is joining us from the West Coast. Elsa is on the West Coast, she couldn't be with us in person, but with us here on my left is a Dr. Michael Klare who is a secretary of the Arms Control Association Board and a former professor, emeritus professor from Hampshire College.

Next to him is Charlie Dunlap. Charlie is professor of the Practice of Law at Duke Law School, and he is also the executive director of the Center of Law, Ethics and National Security. Farthest left is Professor Jordan. Richard Jordan is professor of Political Science at Baylor in Waco, Texas. They're going to address us in this order, but I thought that we should start with Elsa Kania from the West Coast where we have her on the line just in case the technology fails. We hope not.

We have agreed that each presenter will do no more than 10 minutes, because we want to try and leave as much time as possible for Q&A at

the end. Yeah. I'm just advised that maybe the best way to do it is to have the panelists here go and then Elsa will come at the end, while they sort out the connection.

Michael Klare: As we look into the future of war, we find that the U.S. Air Force intends to send fighter planes into combat accompanied by groups of unmanned aircraft designed to attack enemy radars and air defense systems. The Army plans to send its troops into combat backed by unmanned caravans of supply systems and fighting vehicles. As current technology evolves, these unmanned systems–and, by the way, the military uses the term "unmanned," and I will continue to do that, but some analysts say we should call them "uninhabited" or "uncrewed," to suggest gender neutrality–will be capable of operating in swarms, identifying targets for attack and, in the absence of secure communication links from their human overseers, to make strike decisions on their own.

These developments pose significant challenges to international law and governance. Most obviously they raise questions regarding combat commanders' ability to adhere to the laws of war and international humanitarian law, (and this came up in the earlier panel). Generally speaking, these laws require that parties to war be capable of distinguishing between armed combatants and civilian noncombatants, and to spare the latter as much as possible from the consequences of fighting. Many analysts believe that fully autonomous weapon systems will never be capable of exercising such judgment and so should be banned.

Michael Klare: But concerns have also emerged over the implications for nuclear stability of deploying autonomous weapons. Many analysts worry that the introduction of these systems will increase the risk of accidental or unintended escalation from conventional to nuclear warfare, or will invite the adoption of first strike nuclear war plans, greatly increasing the potential for early nuclear weapons use in a crisis. It is these concerns that I want to address in my comments. For, as I see it, in a world of competing great powers armed with nuclear weapons, the most urgent task of international governance has to be to prevent the escalation of war across the nuclear firebreak, as the survival of human civilization cannot be guaranteed once that firebreak is crossed.

Brian this morning described himself as a futurist. I'd love to hear his estimate of the potential for nuclear war in the future. Five years ago, I would've said it's close to zero. Now I say it's hovering towards 50%, and I wish we could discuss that. But I think we have to begin with the observation that the development and deployment of the autonomous weapons is only part of the larger trends in world affairs involving increased competition among the great powers, especially Russia and China and the U.S., combined with the greater inclination to consider the employment of nuclear weapons in future great power engagements.

This represents a significant shift from recent years, in which counter-terrorism, rather than great power competition, was viewed as the main current international security affairs and the use of nuclear weapons was considered a

very distant possibility. Today, by contrast, great power competition and conflict has been designated by the U.S. Department of Defense (DoD) as the principle threat to U.S. security, and in this new threat environment, nuclear weapons have been accorded increased salience in U.S. military doctrine–and I believe the same thing is true of Russian and Chinese, or certainly Russian military doctrine. At the same time, the great powers, led by the United States under the Trump administration, have backed away from the arms control architecture that has constrained nuclear weapons developments in past decades.

In consonance with this altered international environment, the major powers have taken numerous steps to enhance their nuclear and conventional military capabilities, and to place themselves in a more combat-ready posture. This has entailed, among other things, increased spending on a wide variety of conventional and nuclear weapons and the deployment of additional combat forces in potential combat zones, such as the Baltic Sea region and the Black Sea region in Europe, and the South China Sea in the Asia Pacific. If we look at the nuclear arena, I don't have time to go at length in this, we see the modernization of the nuclear arsenals of the U.S., Russia, and China, and we see the acquisition of low-yield nuclear weapons by the U.S. and Russia. We also see renewed interest by the United States in the development of missile defenses, and we see the development of conventional or dual-use missiles with hypersonic velocities that could be used in strikes on enemy command, control, and communications facilities, and

other high-value systems deemed essential to the prosecution of modern war.

Michael Klare: These developments, troubling as they are, are being accompanied by a shifting attitude towards the *use* of nuclear weapons. During the Cold War period, of course, nuclear weapons were part of military strategy and weapons were deployed for that purpose in Europe and elsewhere. After the Cold War, these battlefield nuclear weapons were largely removed from deployment and nuclear weapons were thought of being used solely for the purpose of deterrence of a first strike by an enemy power. But more recently, that has shifted again toward a more Cold War-like environment in which nuclear weapons are being seen as potentially usable instruments of war by Russia and now by the United States, and under the Nuclear Posture Review (NPR) adopted by the Trump administration a year ago, the DoD envisions more uses for nuclear weapons. For example, and I highlight this, the NPR envisions a nuclear response to attacks on the nuclear command and control systems of the U.S. and its allies. The accelerated development of autonomous weapon systems by the U.S. and other countries has to be seen in this light.

The Pentagon's stated intent in developing autonomous weapons is, of course, to provide combat commanders with additional tools to identify, track and engage enemy forces, while minimizing the risk to U.S. military personnel. Autonomous weapons, it believes, provide various advantages in performing these functions as they're highly capable in tracking and identifying targets. They're also expendable. They don't carry humans who we

prize, and because they don't carry humans, they can be produced in large numbers and much more cheaply than conventional manned weapons. You could conceive of acquiring many of them, and because they can communicate with each other and coordinate with each other, they can be deployed in swarms, and spin and weave at machine speed.

How many F-35 fighters can the Pentagon afford at hundreds of millions of dollars each? Just a few. How many destroyers can it afford at $1 billion each? Very few, and they're not planning to build many of those anymore. Instead, they will deploy hundreds or thousands, maybe tens of thousands of unmanned drone aircraft, ships, and submarines.

Michael Klare: Moreover, these weapons are not drones operating in a counter-terror environment where they have total control of the sky, and they're hunting individual targets as discussed in the previous panel. These are intended to operate in hotly-contested environments where enemy forces will have jamming capabilities, and they're intended to attack high-value targets such as air defense radars, early warning systems, airfields, and other very highly capable systems that are crucial to the defense of Russia and China. The question is, in this environment, when you have swarms of these systems in use, what will be the impact on the escalatory potential of future conflicts? That I think is the question that we have to ask. That's the ethical and international relations question that we have to ask.

Michael Klare: The answer, I believe (and I'm very much influenced by our keynote speaker, Paul Scharre, in this matter), is it is going to increase the potential for escalation of conflict to the nuclear realm. Why? Because if you lose contact with these systems–and as we've been told, it is entirely possible that in a hotly contested area it will be increasingly difficult to retain human communication with them–that they could go rogue or they'll continue fighting even though a political decision may be made to halt the fighting or to keep it at a limited level, that these machines will go on attacking high value systems that are essential to the defense of these countries and may be seen as a prelude to a nuclear attack, and therefore prompt early use of nuclear weapons by an adversary.

More worrisome still is the possibility that war planners will conclude that with the introduction of more capable ICBMs and SLBMs and bomber planes and missile defenses that are being conceived, that you can conceive of a nuclear first strike using swarms of undersea submarines, like the Orca system just being funded by the Navy, to attack enemy missile-carrying submarines, the most safe, secure, reliable, retaliatory second-strike system. You could also conceivably use thousands of drone aircraft to search out and strike enemy mobile missile systems, which China and Russia rely on for their retaliatory capability.

If you can achieve these kinds of attack, targeting second-strike capabilities, nuclear stability is out of the window. You can conceive of a first strike, and even if you are

not contemplating a first strike, if you just have that capability deployed in a crisis, an enemy has to assume you're thinking about it, and if it sees these swarms coming, you may very well decide you had better launch your weapons now before they're destroyed in the first strike. It's this kind of an environment, these kinds of scenarios that I think we have to worry about much more. The question of whether a drone could distinguish a civilian from a combatant is a moral question, I agree. But even more, I think, we have to worry about the implications for the survival of all of us from deploying these weapons in an environment that's already hotly contested, where the risk of escalation is becoming increasingly great.

Tiya Maluwa: Thank you, Michael. You have noticed that my one minute was stretched to three minutes, but that's because it's a very fascinating discussion and I didn't really want to cut you short.

Michael Klare: I apologize.

Tiya Maluwa: Some of the things that you skipped might come in at the end during Q&A. Without further ado, I pass the floor to Charlie.

Charles Dunlap: Well, thank you very much. I must say, I often don't hear a presentation with which I disagree more, but I've heard one. I hope to address a few of those things, but what I want to do is maybe take you back and talk about autonomous weapons in relationship of the development of international norms. International norms – as I would use the term – aren't necessarily treaties, they aren't necessarily customary international law, but rather they're more or less voluntary expected norms. And the way they develop or different

ways they can develop, but at some point you have to have countries agreeing to them and following them.

The U.S. has started doing this. If you look at the original DoD directive that talks about autonomous weapons, it discusses keeping a man in the loop, in some way, and what does that mean? Of course, that's the devil in the details because we need to understand the systems better. I'm not one of the people. I agree with some of our previous speakers who would say that we are not going to see a "Terminator" anytime soon by any country, let alone the United States. Because I think we have to understand it's not just the autonomy in the system, but it's also the mechanics that would go into it. Imagine trying to build a robot with autonomy, but also with the hydraulics and everything else: power supply, and fuel and everything else. We shouldn't look at these issues in isolation from other technologies, because a lot of things now are getting lumped under autonomy and we haven't even really described or settled on an internationally accepted definition. If I was looking to develop norms, I would at least try to agree on a definition.

Charles Dunlap: Is that going to happen anytime soon? I don't think so. Because as many of you know, at the UN, a Group of Governmental Experts has been meeting for years and they still haven't settled on a definition. Part of the problem – and where I think we ought to focus on – is that there's lots of *automated* systems. We've heard about them, but I don't think we should think about them so much as *autonomous* systems. What we ought to focus our energy on

is *learning* systems, because that is where the real difficulty comes. Because with automated systems you can take them out to a range, you can run scenarios through them until you get to the point which, what international law would require, is what a reasonable commander with the reasonable understanding of the system would expect, that is, that the use of the system would not cause, for example, civilian casualties which are excessive in relation to the anticipated military gain.

The problem with using fully autonomous systems is that you have *subjective* decisions to be made with respect to targeting. What is excessive in a particular circumstance? What is the anticipated military gain? That to me is still going to be in the mind of the human being for some time in the future. I often think of the Inchon landings during Korea. What autonomous system would have ever decided that was a good idea because all the analysis said it was a horrible idea. But MacArthur thought it was a good idea and it turned out to be a good idea. There are things that go on in human mind that will give an asymmetric advantage to those armed forces that don't go fully autonomous.

Charles Dunlap: We do face threats that we are going to have to have a lot of autonomy. Cyber is one. When things are happening so quickly, you're going to have to have a system that can respond almost automatically. But what do you build into an international norm which would help you with that? My personal view is, I don't like it. I think it's unwise when we pick out a particular system and give it its own unique legal regime, and I think that would be

impossible with autonomy. I think it's better to focus on the existing law of war and insist on adherence to it, rather than trying to throw up your hands and say, we need something entirely new.

I think my colleague here, his reference to the nuclear treaties is a little bit of an illustration of how difficult it is to have a unique treaty for just a particular kind of weapon system. When we talk about the INF, the INF was really a bilateral treaty between us and Russia. It didn't stop China from building those exact kind of missiles that were prohibited under the treaty. We have to look at that in the future, because these systems are coming, they are going to be available to many militaries.

As other people talk, some of the technology is off the shelf, but here's the good news: autonomy is going to be in every part of our life. We will have civilian systems, different kinds of systems that will . . . autonomous cars and so forth, that will help us develop the testing processes that we will need to have the security that these weapons are going to operate as we intended.

Charles Dunlap: What is particularly tricky about learning system, unlike our existing systems where you normally, you test the weapon, you give it a review and then you deploy it, and you govern how it's used. But with learning machines, you're never going to be done testing because the machine you tested when you first deploy it will be a little bit different – or a lot different – from the machine a year down the road or even less, so we have to look at that. If I was to look to try to develop international norms, I

would try to develop something . . . What should the testing protocol be? Who should be involved with that? And so forth. But I do think it's going to be very difficult to develop norms. This is why the U.S., the UK, and Russia aren't even interested in trying to negotiate a particular treaty with respect to these kinds of particular weapons, in part, because they haven't agreed on the definition. I think in part they realized that they are going to be integrated.

Charles Dunlap: Just a couple more things. Another problem I have with treaties like we are talking about, if you focus on a particular system, there's a treaty that says it is illegal under international law to develop a weapon that is intended to blind combatants. You can develop a weapon that is intended to incinerate the adversary, but you just can't develop one that's intended to blind. Why? Because in the 1990s, if you went to the ICRC website, they say, well, being blind is like – these are not their words – but in essence it was like worse than being dead. But there they captured the technology at a moment in time, and today as horrific as blindness would be, it's better than being dead, and we're developing technology to help people see. That's the problem with these kinds of specific treaties.

I agree with my friend here that nuclear weapons are horrible, and their use would be horrific. But let's keep in mind we have never had a war between two nuclear armed countries. That is a complete turnaround since World War II. We have never had that, because there is an understanding. When I look to the future, I'm less concerned about nuclear

weapons because, for a variety of reasons, what we can talk about in the Q&A, and I want some students to ask questions during Q&A because I may call on them if they don't have questions. But this, I'm more concerned about biological weapons, I'm more concerned about DNA based weapons, I'm more concerned about cyber weapons, because cyber weapons can have really catastrophic effects. Now we're seeing some talk about EMP based weapons.

Charles Dunlap: I do think that there is a place for norm building, but we have to look at it rather modestly, and not try to put these in their own category, because at least now we have a track record. We have a history of how we interpret the international law of war with respect to weapons. We don't want to get to the mode that there's this whole new system, so everything goes out the window. In terms of nuclear command and control, all great. Technology, not just autonomous technology. Technology – not just autonomous technology – always presents challenges for controlling escalation. Absolutely. But I think that what we will see in the near term with autonomous systems is the development of decision-support systems, which will be helpful to and better than a lot of the analog systems that we depend upon now.

The problem that I see is the relationship between autonomous systems and the data that they use to build their decision-making process. Garbage in, garbage out. I think what we're going to really have to focus on the future is how do we determine what is the appropriate data, how much data is necessary to get the results that we want? In this respect,

ironically, China has an advantage to us, because they have the ability to coerce a billion human beings into providing data, which helps them build algorithms that may be more predictive than what we would be able to do. On the other hand, it's a different culture and maybe that will make them build algorithms that are mistaken when they use it to interpret what we are doing. In other words, there's a lot of opportunity here, the devil is in the detail. I don't want to take too much time because I really do want to get to the students' questions.

Tiya Maluwa: Thank you Charlie, and talking of China, we hope that we will be able to reconnect with Elsa because she intends to talk about China's approach to the development of autonomous systems.

Charles Dunlap: Believe me, Elsa's forgotten more than I've ever known. She is the person on this topic.

Tiya Maluwa: We hope to reconnect with her shortly. Yes, Richard.

Richard Jordan: It's a real privilege to be here amongst so many distinguished figures. I appreciate the diversity of the people you have here and being part of it. What I'm going to talk about is game theory and autonomous systems. I want to start with a few common normative concerns about autonomous systems, and these are familiar to everyone here in the audience. First, they're distant. Second, they often involve civilian losses. Third, they're automatic, sometimes even removing the decision for lethal force. And finally, they're abstract: the violence you're seeing at best through a screen, at worst as a few data points or even not at all. This all

leads them to be rather dehumanizing, and dehumanizing in a particularly insidious way.

Richard Jordan: I think in today's modern polarized political climate, we tend to see "dehumanizing" as calling the other guy an animal or something like that. But that's not the most dangerous form of dehumanization. Autonomous systems resemble what Hannah Arendt was talking about with Eichmann in Jerusalem. She said, "Eichmann was not Iago and not Macbeth. Except for an extraordinary diligence in looking out for his personal advancement, he had no motives at all. He merely, to put the matter colloquially, never realized what he was doing." The danger here is not that you think of the other person as subhuman, but that you don't think about the other person at all. That's the danger that we're seeing ethically and normatively with autonomous systems.

But what I also see when I look at this list is that we faced all of this before. In fact, we faced it even worse than we're facing it now—with nuclear weapons. (They're becoming a theme on the panel, right?) Nuclear weapons aren't just distant, they're intercontinental. It's as far as you could possibly be. They don't just risk civilian casualties: there were debates about whether they should even target civilians deliberately. For many people, being automatic was not a bug, but a feature: an automatic response was sought because it improves your bargaining position. They're not just abstract, but they're studied with Game Theory, which as a branch of mathematics is literally as abstracted as it is possible to be. The critique here is that nuclear weapons are also dehumanizing. In fact, I think they're probably

more dehumanizing than anything we've yet seen emerge with autonomous systems. That should be *encouraging* to us because we've made enormous moral progress in thinking about nuclear weapons and then cooperating around them. If we can solve that problem, I think we can solve the other.

Richard Jordan: In the rest of these remarks, I want to compare these two systems and what the development of international norms around nuclear weapons can tell us about the potential for norms around autonomous systems. To do that, I want to go back to a classic debate between two Game Theorists in the 1950s and 60s, Thomas Schelling and Anatol Rapoport. Schelling was an economist with RAND and instrumental in development of brinkmanship and other nuclear strategies. Rapoport was a game theorist in biology, and he was appalled by the idea that these methods that were being developed to study life in evolving systems, were being applied to millions of human beings and their potential extinction. He argued, "seduction lurks also in the mental habit of rational analysis. This analysis requires detachment,"—and that detachment word should be setting off alarm bells in all our heads.

He goes on to say, "one cannot play chess, if one becomes aware of the pieces as living souls." That's a quotation I think that has haunted me from the first time I heard it. But Rapoport lost that debate, and he lost it decisively, and he lost it for the simple reason that in fact we do play chess with human souls—and that is a grim fact but one that's necessary to face head on.

Richard Jordan: What can Game Theory tell us? Harrison Wagner says that there are two fundamental insights from Game Theory. First, the idea of strategic interdependence, that what I do to pursue my ends will depend fundamentally on what you do to pursue your ends, and vice versa. Second, the idea of multiple equilibria, that any strategic interaction can end in multiple ways (or just about any strategic interaction). To solve the first, you study it through abstract reason. It's essentially an optimization problem. It's studied through math.

You solve the second through imagination, if we may call it that. Through setting yourself in a culture, through trying to understand the normative and moral paths that might lead you from one equilibrium to another. Here I think you can really see what Rapoport's mistake was. He was trying to study the first problem using the tools of the second. He was trying to import into what was essentially a cold calculation, something that didn't belong there.

Richard Jordan: Where does moral imagination belong? The answer seems to be in discriminating between possible worlds. Asking, which one do we want to inhabit? How do we get there? The answer that Schelling came up with, and that was so persuasive, is this: you look for focal points. These are points which are—or can be made to be—psychologically, aesthetically, or morally appealing. They're something you notice and you are drawn to, and they have to be clear and they have to stand out. They have to be clear, because you need to be able to detect cheating. You need to know if

somebody else is not converging on the same point as you.

Perhaps even more importantly though, focal points need to stand out. They need to be axes about which all of your other normative expectations can revolve. If we think of focal points, the solution here is *not* to change the weapon, not to try to humanize a dehumanizing weapon. Rather, we need to humanize the moral arguments, the moral narratives we're constructing around it, and with nuclear weapons, this looked like the nuclear taboo.

Schelling dedicated his Nobel acceptance speech exclusively to the nuclear taboo and its success in preventing nuclear war over the past half century. The taboo is the idea that there's a qualitative difference between zero and one. If we're looking for an equilibrium in which were going to converge, nonuse is the most attractive.

Richard Jordan: It's become such a part of our culture, such a part of our moral upbringing, that when I teach about the nuclear taboo in my introductory class to the undergrads—most of them have never even thought about nuclear weapons before—even though they'd never heard this phrase, all of them know what it is, and it's not surprising to any of them. It's not surprising because somehow passively they just absorbed it in the culture around them, that this is the expectation, that this is the moral anchor for how we think about nuclear weapons.

Other famous focal points would include like the moratorium on whaling. Economically, it's suboptimal. We can sustainably fish whales at

higher rates than zero, but we converged on zero because it's simple and clear: save the whales! It's very clear to everyone involved and it's easy to detect cheating: if one person has commercially fished whales other than the few people who aren't on board with it, you know somebody broke the rules. Something perhaps a bit more familiar to a legal audience would be norms surrounding outer space. In the fifties and sixties there were debates, should we use the pattern of airspace or the pattern of the law of the sea to think about outer space? We really owe a lot to those early normative entrepreneurs and legal entrepreneurs who said, we'd much prefer to live in a world in which outer space is a global commons, like the sea, rather than parts of it being the private property of Russia or America.

Richard Jordan: What does all that mean for autonomous systems? What are the focal points for autonomous systems? I think there is a challenge here because we don't have anything like a nuclear taboo. It's not going to be zero. I think we're going to have to create the focal points ourselves. That's a challenge I think for everyone in this room, especially people who are thinking creatively about these weapons. (The way I think the nuclear taboo became really institutionalized in our culture through fiction came up during professor Jordan's talk earlier today; we need to do something similar with autonomous systems, and we haven't done that yet.)

I do have some suggestions, with just my last few minutes, of what these need to look like. The first, is a reminder that a focal point has to be an equilibrium. It has to be strategically

rational for the main players. As an aside, I've always found it irksome when people critique President Obama as the "drone warrior," because it's always seemed to me that particular policy flows very much from domestic incentives rather than from personal character. I think we have to recognize that as long as certain domestic incentives are pushing actors in the same direction, then any leader in that situation is going to take those actions. A total ban is just not going to be acceptable to democratic publics. Another point is, if it's a winning strategy to make a decision in a millisecond, it's rather a moot point about whether we want to make that decision or not. It will be made.

Richard Jordan: Let me turn to my second suggestion. I couldn't resist putting a clickbait headline up here, since we are talking about drones. This jumped out at me when I saw the Kalashnikov assault rifle: "now we have Kalashnikov drones." (What do you expect from your Facebook newsfeed, right?) I think there's a tendency to focus on trying to solve problems like the one in this headline because it's what draws our attention; it's what the public is talking about. But this is really the hardest possible problem we could set ourselves to solve. There are literally thousands of actors involved. State and non-state actors, all with competing incentives, all from different moral traditions. If you're looking for a moral or aesthetic focal point around which you're going to organize, good luck, because I don't see it.

But that suggests to me that we should *start at the top*. This is a chart taken from a colleague's

working paper down at Texas Tech. It plots the different distributions of drone technologies for consumers and non-state actors versus commercial actors versus state actors. Now drone technology is a continuum, but what I think this captures is there's something approaching a qualitative difference between the vast bulk of what states are using and everything else. By being drawn to the problem that, "oh, we can build these things in our garages," we're missing the easiest places to begin norm building, which is at the top, where there's only a dozen actors need to get on board and most of our incentives are aligned. Instead of trying to solve the problem from the bottom, let's go to the top of the distribution, and work our way down.

Richard Jordan: Third, and this is really driven from the idea of focal points, is *keep it simple*. I think especially because the technology is so fascinating, especially to people in this room, that it is easy and it is fun to talk about the ins and outs of technology; but it's also really, really confusing to talk about, and to start trying to draw distinctions in it, especially to outsiders. I think we need to set ourselves this aim: that any president, any congressman, any senator, any secretary needs to come into office the same way as my undergraduates come into my class. They need to know the guiding principles of how to use autonomous systems without knowing anything about the technology, because they just inherited them from the culture around them. For people to do that, the principles have to be something really simple that even, forgive me, that even as senator could understand. I think that's what we should set as our goal.

I'll conclude. I think the central question here is, can we draw clear moral distinctions that align with our incentives, focused our cooperation, and by capturing our imaginations, rehumanize this fundamentally dehumanizing technology? If that seems like a tall order, I would just say—we've already done it once. Why can't we do it again? Thank you.

Tiya Maluwa: Thank you, Richard. We'll now try and reconnect with Elsa.

Elsa Kania: Okay. I'm glad to be here, at least in spirit and virtually and I thank you so much for the invitation to participate in this panel. It's been a fascinating conversation so far and I'm looking forward to your questions and to continuing the discussion. I'll provide some quick perspectives on how I see the Chinese military exploring the development of the range of applications of artificial intelligence, including for autonomous weapon systems. Then I'll comment quickly on some of the emergent legal and normative debates in which China is taking part, and where I see that going forward. Please somebody yell if I cut out again audio wise because I don't want to be talking to myself.

Okay, so here we go. I think that it has become quite clear in the past couple of years that artificial intelligence is a new frontier for strategic competition among great powers. China's approach to date has been informed by its close study of initiatives the U.S. third offset, which provoked a lot of concerns among Chinese leaders because that was aimed at reestablishing U.S. operational advantage relative to China, and that focus on concepts

like human-machine teaming and learning autonomous systems, and the notion that today's emerging technologies could change conflict in ways that enable potentially disruption of the existing balance of power. I would argue that a lot of those ideas have been very deeply ingrained into how the Chinese military thinks about warfare today.

Elsa Kania: Increasingly, there appears to be a consensus among Chinese military strategists that we are in the midst of a continuing revolution in military affairs. I will note that the notion of an RMA has some history. It has fallen out of favor in the U.S., but there has been some consistency in Chinese military thinking on the fact that information technology along with today's emerging technologies are starting to transform conflict in ways that may have far reaching implications. This may be a little bit jargon-y, but the Chinese military talks about today's work there as informatized and the military strategic guideline focuses on fighting and winning and informatized wars within the region. Increasingly, the notion is that warfare is becoming intelligentized or smartified you could say, but that sounds a little bit less serious, and that AI is really going to be pervasive across all aspects of military power for a range of capabilities.

I think when we talk about norms, when we talk about some of the legal considerations, I think it's important to keep in mind this backdrop of fairly intense military competition. There is a lot of talk these days that we are in a new era of great power rivalry. I would argue, in some respects, there is more continuity than change, because the Chinese military has been

looking at the U.S. as both a teacher and also as a target of their development since at least the 1990s. I think across the range of applications today, some do involve autonomous systems including swarming drones and those varying degrees of autonomy.

Elsa Kania: I happened to be walking through China's military museum in Beijing a few summers ago and saw a small exhibit on future warfare that had a depiction of a swarm combat system going up against an aircraft carrier. Although they did not specify whose aircraft carrier might be the intended target of such swarms, I think then we can luckily come to some conclusions there. I think that the Chinese military appears to be thinking fairly asymmetrically about how they can leverage these capabilities in fighting a more powerful adversary, which is how they tend to see the United States. I think it is far too soon to say which military will ultimately be advantaged by these developments. I would also add that AI is not a singular technology, but a range of techniques that have quite multifaceted applications, some of which are more concerning than others. For instance, I do believe that the introduction of greater autonomy into cyber defense and automated offensive operations in cyberspace could pose a major concern in terms of escalation or a momentum driven conflict emerging in that domain. When we think about some of the advances in autonomy and cruise missiles, even hypersonic glide vehicles that China is pursuing, some of that may be incremental in nature.

Elsa Kania: But in other respects, when you start to combine the speed and increased capability and precision of these weapon systems across the board, it does start to become more significant as perhaps enabling changes in the military balance, or conduct of operations going forward. There appears to be a consistent theme in Chinese military writings—although I will note that to my knowledge, the Chinese military does not have an official policy on autonomy yet to date—the notion that given these advances and artificial intelligence and autonomy, the tempo of operations will increase progressively. Eventually there will be a point where humans can no longer remain fully in the loop, and instead we'll have to be on the loop or even out of the loop at a certain point in time. There is intense attention to the notion of AI and decision support, particularly to improve commander's decision-making capabilities.

Although herein otherwise, but in some respects, the PLA's approach to AI may be informed by self-diagnosis of their weaknesses. For instance, Chinese military leaders have complained that they believe their commanders are not capable of making good decisions, evaluating the situation and executing orders well. Just as AI could be one response to that, organizationally in a sense, or the notion that introducing AI for decision support could augment a commander's capabilities and will also increase the speed and efficiency of commanding control including through data fusion and integration of information on the battlefield. Though of course, I think the tendency of bureaucracies and organizations towards stove piping and the

challenges of managing military data, which quite often is data that is relevant, is often quite distinct from the data that is more prevalent in the commercial ecosystem. I think a lot of these factors could be challenges and impediments in the near term.

Elsa Kania: As I mentioned, the Chinese military does not yet have a formal policy on how they will approach questions of autonomy, or at least if they have, they have not released it or been transparent about it. I hope that at some point they might, but at least so far there has not. There does seem to be some awareness of the challenges of testing, but I have seen less discussion than I would hope to see about some of the reasons for concern about the vulnerability and unpredictability of autonomous weapon systems. For instance, as we saw with my own technical difficulties at the start, technology can be unreliable at times, whether for reasons of connectivity or a human error in some cases, or the human machine interface. One question I would raise that I am not sure we have a full visibility on is whether the PLA, the Chinese military, may be too enthusiastic about some of these capabilities and not realistic enough in their awareness of some of the shortcomings in operating under real world conditions, including because they lack recent experience in combat.

Elsa Kania: The Chinese military has not fought a war since they fought Vietnam in 1979. They are trying to improve training and realism, including through war gaming. We're getting that this involves AI in some cases. But that is not a substitute for the harsh lessons of actual

operational experience, I would argue. There could be perhaps an overconfidence there that could prove dangerous or concerning. With regard to China's engagement in diplomatic forums so far at the United Nation Group of government experts on lethal autonomous weapons systems, China's position has evolved and has been a little bit peculiar at times. About a year ago, China's delegation came out in favor of a ban on the use, and not development, of fully autonomous weapon systems. However, the definition that they described in their working paper at the time appeared to exclude everything that the Chinese military and defense industry were interested in or already developing. They defined a fully autonomous weapon as one that had no human control at any stage in the process, no option of termination ones that had been launched, and was entirely in discriminant.

No professional military would want a weapon that is indiscriminate. That defeats the purpose. If you are fighting a major adversary, you want to be sure your weapon will function as intended, and essentially that support for advancing have been somewhat symbolic and some would characterize it as a legal warfare or law fare of sorts. Tactic to take common cause with activists pushing for a ban which could prove unrealistic, well framing it in a way that would not constrain their own developments. From what I've heard at the UNGG this year, China was rather quiet and reaffirmed some of these positions because I'd really moving the conversation forward. But it will be interesting to see if the Chinese government does start to become more engaged in these forums going forward.

Elsa Kania: It is worth noting that China's new generation AI development plan launched in December of 2017 does talk about China playing a greater role in leading in the global governance of AI, including developing a legal, ethical and normative frameworks. I think that's encouraging, but I would note that the ethics and norms that are prioritized in China by the Chinese government may be different from those that we'd see in debates in the U.S. and even if Chinese companies or engineers may care on a personal basis about privacy and personal freedoms and the Chinese government clearly has a strong interest in access to data and surveillance that may override any resistance though there has been some pushback at times.

Elsa Kania: I think China, as of this spring, the Chinese government has set up a new commission to think through issues of ethics, law and regulation for autonomous systems. It'd be interesting to see what comes out of there. There has been some talk by Chinese military strategists and scholars about concerns over an AI arms race and options for arms control. Although I think mechanisms that are feasible may prove challenging or perhaps unrealistic.

To conclude, since I think I am getting near to the end of my time, I'll tell you that I think going forward as we think about norms, we have to keep in mind the reality of this competition and the intense security dilemma that is exacerbated by uncertainties over the status of developments. Arguably the U.S. and China tend to overestimate or exaggerate in some cases, each other's capabilities. There's some amount of propaganda and perhaps

information in Chinese reporting on what they've accomplished. For instance, claiming to have a world record in swarms of drones and to have beaten us records. Size does not matter most in this context necessarily, and a large swarm does not necessarily mean it's the most sophisticated in terms of algorithms. I think definitely some propaganda or one-upsmanship there. But I think given this competition, I think it will be important to think very pragmatically about where the U.S., China and Russia have alignment of interests, and I would argue that in risk mitigation, and that could include the options for testing, verification, fail safes, sharing of best practices.

Hopefully there will be some opportunities for dialogues among great powers on these issues going forward because I think the risks to strategic stability, including because the nexus of AI nuclear weapons even simply in early warning does seem concerning. I think it will be really valuable to try to engage among competitors to see if there are pragmatic measures that all could agree upon would be a favorable path forward. I will stop there. I hope you could hear me fairly well for most of the time. Hopefully I've not just been talking to myself and I will look forward to your questions and conversation.

Tiya Maluwa: Thank you, Elsa. We heard you very well, and we have just under fifteen minutes for questions. I would like to encourage anybody who has a question or comment to step up to the mics on the right or left, and bear in mind the invitation that was issued particularly to students. So, please.

Audience: Hi. I welcome you all here. I've noticed just through my own research and the conversation, there's a huge emphasis on humanitarian law with autonomous weapons systems. I was wondering, especially in light of increased use of autonomous systems in domestic law enforcement if we need to, or the international community needs to address human rights norms or human rights implications with autonomous weapon systems?

Charles Dunlap: Oh, I'll take a quick stab at that. I think you're onto something. I'm not a big fan of international human rights law in the context of armed conflict, but it does have a place. I'm particularly concerned about how you build the algorithm and do you build in biases when you build the algorithm? Also, where do you take the data to build the algorithm? Or to fuel the algorithm so to speak. I think that that has a lot of implications for international human rights law. I think that we need to think about what kind of data are we about, individuals, for example, are we going to permit into the system? That's only scratching the surface. I don't know if anybody . . . But you're onto something.

Audience: Okay.

Charles Dunlap: Keep working on it, because there is a lot of discussion . . . the whole predictive sentencing phenomena in criminal trials, predictive law enforcement. It all sounds good. But in the end, and especially if you have a learning system where you start out with bad, with say, biased data, it's only going to get more . . . it has some potential to get more biased and

more unfair, and that's particularly something that, that you ought to be concerned about in that context.

Audience: Thank you.

Tiya Maluwa: Thank you. Chuck.

Chuck Diebel: Thank you panelists for being here. I'm Chuck Diebel, a third-year joint degree student here. A lot of talk about fitting new technologies into new legal regimes, new arms control agreements and treaties. My question is about countries pushing back against current international norms in new domains, cyber and space. What the United States can do in ensuring that the liberal order that's been created since World War II is sustained, and not necessarily focusing on a specific technology, but how to fit the current technologies in the current legal norms and ensuring that those norms are complied with.

Charles Dunlap: Actions speak louder than words. That's why I think that what the administration did to Internet Research Agency, right before the 2018 election, it at least established that you can't interfere with U.S. elections and suffer nothing for it. The other thing that I think our government needs to do more is we need more *opinio juris*. In other words, if you look in particular, I'm thinking the cyber domain is it's a little bit amorphous, we have the Harold Koh speech, we have a chapter in the DoD Law of War Manual, but we get sometimes inconsistent messages. I think that's a big part of it.

Charles Dunlap: But you're exactly right. In the cyber area, China has said, "we don't think the law of

armed conflict applies." We have to push back on that if we want to have norms, but there has to be some consequences. That's one reason why people say, "Oh, the indictments didn't do much with Internet Research Agency, it's just a couple guys." It really starts to lay down a marker, but we have to do it in other areas and including with respect to autonomous weapons.

I think that is what we ought to do especially as we look for law development. I think there are a lot of good ideas. But also in the testing and approval of a weapon prior to its deployment. There are things in the existing Geneva Conventions, and Protocol I that speak to that. But I think there are only like twelve or fourteen nations among the hundred and some signatories who actually do weapons reviews and testing before they buy them or field them.

Michael Klare: I think you raised a good question. I do represent something called the Arms Control Association, which is a non-nonpartisan organization. But our goal is to, as best as possible, to advocate for and promote disarmament and nuclear stability through international agreements. That's where we come from, and as an organization, we heavily supported the Iranian nuclear deal, the Joint Comprehensive Plan of Action which we thought was a very important step towards nuclear nonproliferation. We are very distressed by the administration's decision to withdraw from that. We hope that there will be a way to preserve that treaty or to restore that agreement down the road. We're very distressed by claims that Russia has violated the

INF Treaty. But we don't think that withdrawing from it was the right way to approach that problem. Rather, there should be more consultations, more effort to resolve the question of Russian violations, and see if there's a way to resolve them without breaking out of the treaty.

Michael Klare: I agreed with my colleague that China is not a member of the treaty and has pursued weapons which are dangerous. In addressing this problem, we should be working with Russia through a strategic dialogue to try to persuade China to enter into discussions to try to impose some kinds of limits on these weapons. Instead, the U.S. is now pursuing INF-noncompliant weapons which will put Russia's and China's command and control facilities at risk. Russia has said it will do the same, so we will be in a much more dangerous environment. My colleague here from Baylor said there is a nuclear taboo, but that doesn't last forever. These agreements and commitments to controls on nuclear weapons were intended to buttress the taboo by making the onset of nuclear war harder.

Michael Klare: Yet, everything that I was talking about earlier was to show how those controls are being undermined and how all remaining arms control agreements are now at risk, including the New START Treaty, which could also be terminated. All of those were intended to protect the taboo by making it physically hard to start a nuclear war. Well, all of those are being eroded, and so whether or not the taboo remains intact is in doubt when the mechanisms in place to preserve it are at risk. Where autonomous weapons fit in, that is what

I was trying to get at, and my concern is that they are increasing the risk. I do think that there are mechanisms, some of which have been discussed, that could diminish the risk that autonomous weapons, if fielded, would play that destabilizing role, and we can talk about that.

Charles Dunlap: Well, number one, we've always tried to put enemies' command and control systems in at risk, so that's nothing new. Secondly, the Iran agreement was going to run out. I mean, it wasn't a permanent fix. Then thirdly, I question the Nuclear Weapons Ban Treaty as it raises moral questions of exactly the opposite. In other words, I think adherence to it raises moral questions and I just invite people to take a look at an article I wrote on War on the Rocks. It's called "Is a Nuclear Weapon Ban Treaty and Moral," and just make your own judgment.

Tiya Maluwa: Right. We do have four people lined up who'll ask questions, but before we do that, can I just pass the flow to Elsa? I'm informed she wants to come in, and then I'll come to you.

Elsa Kania: Oh, sure. I was just going to respond to an earlier point about some of the norms and considerations for human rights in play. Because I wanted to highlight that what scares me most about what China is doing in AI is not the weapon systems but the ways in which these technologies can be leveraged to increase the course of capability of the state, including for surveillance and censorship. I think the situation in Xingjiang today, which is absolutely and just beyond the pollen with the intense repression and detention of by some

estimates up to a million Weekers. If anyone saw the New York Times earlier this week a feature on how cities are essentially being turned into Open Air prisons, with pervasive surveillance and very granular monitoring of individuals for anything that is seen as an indication of what they're calling terrorism. But essentially this has often been characterized as a cultural genocide against minority population in China.

Elsa Kania: I think some of these tools are being experimented with in Xingjiang, and companies are making a lot of money in the process, including some like iFLYTEK and SenseTime and at some point some American companies have been suppliers so that does put some ethical questions in that context as well. But I think that these capabilities can have the potential for abuse and democracies or an authoritarian regimes. But I think the damage can be graver when there are protections and the freedom to have a full ethical debate. I think we should expect to see these capabilities diffusing or proliferating, including again, because of the profit mode. But a lot of companies can make a lot of money selling tools for surveillance and particularly leveraging facial recognition. Chinese companies are already increasingly providing this to other governments, including those of either shakily democratic or outright authoritarian, have recently an agreement in Zimbabwe, new partnerships in Malaysia and Singapore. Have programs in smart cities and facial recognition for urban environments that . . . and I think the aggregate impact of this could be to really challenge the future of

democratic governance in quite far reaching and troubling ways.

Elsa Kania: I do think that it's really an urgent question and there's been some talk of sanctioning Chinese companies that are involved in some of . . . involved or complicit in some of these activities including those that have a presence or activities in the U.S. I think from the . . . I'm not a lawyer I will say, but I think a lot of legal and normative questions too, and I don't have any good answers myself, but I do think it's something for all of us to consider of how do we think about the threat that these technologies may pose to democracy? And how we can start early and trying to really explore new and hopefully creative solutions to make sure that these technologies will be accountable. I think that does start in some of our own debates about how we're using AI and policing and Homeland Security, but also does extend to these international considerations as well.

Tiya Maluwa: Thank you. We have one minute, but I'm going to do the unthinkable, which is to take up five minutes of your lunch time because I don't want to lose the questions here. I would only ask that you keep the questions short, the answers short, and if by any chance the preceding question covers what you wanted to ask, feel free to pass up your chance. But Richard wanted to react very quickly, please.

Richard Jordan: Very quickly. I'll try to keep just three points to the gentleman's question earlier. First, I'd say that the benefit of norms as opposed to law is that they are flexible and evolve. I think we should not only understand, but expect that

these things will change over time. Second, a lot of these norms were articulated during either the Cold War or during the 90s or 2000s, during the unipolar moment. The world is very different. It's going to be very different, so we should expect that these norms will change. On the plus side, I don't think autonomous systems will be driving that change, except insofar as they disproportionately benefit revisionist powers, as they disproportionately shift the balance of power. But the third and perhaps the most important point is that as we adjust the secondary, the buttressing institutions (like the INF), you've got to leave the underlying norms in place. Similarly, as you're adjusting nuclear arms control agreements, you've got to leave the nuclear taboo in place, and not just leave it in place but reaffirm it. Because those sorts of norms take 50 years to establish, and they're lost very easily, and they're very hard to replace. So while we should accept evolutionary change, we need to hold the fundamentals in place.

Tiya Maluwa: Thank you. Please.

Audience: Yeah. I'll try to be quick. In terms of thinking about the distinctions between norms and treaties and laws, is there anything that we might be able to learn for this domain from the very successful establishment of a norm against human cloning without having to have any treaties? That's a case where there was a very rapid establishment of a norm. To the best of our knowledge, it actually hasn't yet been violated. But without any treaties. Can we learn anything from that domain and case study that might be useful over in this one or are they just too far distinct from one another?

Charles Dunlap: I would just, I'm guessing here is, I think they are two distinct, because there's so many uses of autonomy in other areas of society, by cloning is a narrower topic. Perhaps it's easier to wrap your arms around it, whereas here, and we at least we can define what it is. We're still struggling with exactly what we mean by autonomy, fully autonomous machine learning and so forth. I think we need to look at that. That's a great idea, by the way. But I don't know if it is going to be the template, but might be a piece of the template.

Tiya Maluwa: Thank you.

Audience: This is for Dr. Jordan and Dr. Kania. Given that asymmetric trying to get an asymmetric technology advantage, what does that deterrence look like in the South China Sea to get freedom of navigation, trade, self-determination?

Michael Klare: Could you repeat that please?

Audience: Okay. Basically how do you deter, how do these technologies, given the Chinese have a very aggressive stance of the South China Sea and these technologies seem to want to give you a little bit of an edge. How do you model a deterrent strategy for that?

Richard Jordan: The same question.

Tiya Maluwa: Thank you. Elsa I hope you heard the question because you might want to come in on that as well. True?

Elsa Kania: Well, firstly to clarify I am not a doctor. I have a long path before I can be Dr. Kania. But for the time being, I'll say, I think that's a great

question and a difficult one, because there aren't easy answers. I worry we missed our window of opportunity in the South China Sea and that a failure to push back forcefully earlier before China had made the militarization of these features, stocks and the water was really a last chance to try to nip this in the butt so to speak. I think now, unfortunately there is a new status quo and I think we are likely to continue to see militarization of these features. I know there's some debate as to whether these will alter the military balance in the South China Sea or prove less impactful given potentially some of their vulnerabilities. But at the very least this has enabled more placement of radars and weapons systems and does start to reinforce Chinese sea control in the South China Sea in conjunction with the autonomous and unmanned undersea systems that could be used to detect American submarines.

That's the problem side of things. In terms of as solutions, I think for . . . this won't be quite an answer to your question, but I think for deterrence, vis a vis China, we first have to evaluate what our own priorities are. For instance, if maintaining U.S. access to the South China Sea is a core priority, then we might have to be willing to accept a higher level of risk or, but again, I do think that the status quo was fairly locked in there. I think word deterrence will be a major challenge and concern going forward, that relates to the South China Sea will be Taiwan. I think the question of would the U.S. support and defend Taiwan? And how would some of the weapons systems we're talking about, including old fighter jets, converted into kamikaze drones that could be used against Taiwan's air

defenses, how could those factor in? I think part of this will depend on politics and perceptions, because many Americans do question whether the U.S. can and should defend Taiwan. China certainly has an interest in that uncertainty or that uncertainty could also engender miscalculation.

Elsa Kania: Again, I know we're short on time and I don't have a good answer here, but I do think that we will see autonomous systems become a major feature of some of these disputes and flashpoints in the region, including increased use of drones by China to establish persistent presence or improve their domain awareness in the Eastern South China seas. I think U.S. solutions have to first depend upon a political question of, where our red lines are for sure and what level of risk we're willing to accept if we decide we want to change the status quo? Which will be a matter of compelling at this point because the window for deterrence is passed, at least in the South China Sea.

Tiya Maluwa: Thank you very much. Richard.

Richard Jordan: Very quickly, I think I'm going to build on to comments that have come from earlier on the panel. First, I think that deterrence against China is going to differ fundamentally from how you come to terms against a non-state actor or even a rogue state, and that's because it's always in the shadow of nuclear power. That is, as strange and unusual as autonomous systems are, they're not going to change the fact that what we're most worried about is escalation. I think that the opportunities for peace are actually promising because we're bargaining in the shadow of that unspeakable,

horrific event, and that's always going to be what's in the back of people's minds. It's going to be a game of brinkmanship, and gradually ratcheting up the risk of disaster on both sides. Not actually, I think, fighting with most of these weapons. If it comes to that, we've crossed a line from which, well there's almost no going back.

Unfortunately, I'm going to use the general's comment to undermine his critique a little bit of Dr. Klare, which is, I think that we have to accept that the risk the nuclear war is going to increase. Since we are playing a game of brinkmanship as the only viable deterrent strategy in the South China Sea, we have to accept that the risk of nuclear war, it's going to ratchet up from 0.1% 0.5% or something like that.

Tiya Maluwa: Thank you. We will take the last two questions one after the other and then the panelists can decide which one they want to take on. So, you go first.

Audience: I've heard a lot of the panelists talk about accountability today and I think that is probably a biggest concern for the regular American about economy and autonomous citizens. I was wondering your thoughts as humans probably isolate, you'd be moved out of the loop on those decision making, how we account for accountability in that, and how well the international system currently can address those concerns?

Tiya Maluwa: Thank you. Sir.

Audience: I was wondering about the depersonalization effects of warfare because of autonomous

weaponry. You mentioned that that's a common fact between, well that and nuclear warfare, but the core difference is, we all seem afraid of nuclear warfare, whereas autonomous weapons seem to make it easier for us to fight without guilt or much thought of what's happening on the battlefield. I was wondering how those threats and the risk of escalation compares because of that?

Richard Jordan: To this quick question, I think the distinction between them is not as sharp as it might appear. One of the common critiques of someone like Herman Kahn in the fifties and sixties was that these people were contemplating the deaths of millions and millions and were almost gleeful in the fact, because it was so depersonalized for them. I think it's become very personal to us because we've been surrounded by fictional narratives that have made it more personal for us. If I asked you to visualize a nuclear war, you probably don't visualize a picture of the bombing of Hiroshima. You probably visualize Dr. Strangelove or if you're a Millennial of some dystopian landscape from a bad teen novel. That's what made it real to us. (Sorry, that's kind of a flippant comment). It's made it real to us in a way that no one has done that yet for drones. No one hasn't done that yet for autonomous systems, and I think in 20 years we'll have the same kind of perspective as we have now on nuclear weapons.

Tiya Maluwa: And question two?

Charles Dunlap: Well, I was going to say one thing about both of them.

Tiya Maluwa: Yeah.

Charles Dunlap: I think autonomy is going to allow not only the depersonalization of war, but the hyper-personalization of war because it'll enable adversaries go after very specific individuals and to collapse organizations by pulling point out specific capabilities. Regarding your question, I think accountability in the law of war sense in the international law sense, number one, command accountability has not been that good. We just had the Bemba case and that was a pretty clear one. It just was reversed by the ICC. I think that it's going to come down to the commander. The commander is going to be responsible. Ergo, the individual commander, he or she is going to have to have a reasonable understanding of the system. This is why that discussion earlier about explainable AI is critical to the ability to field these weapons and to hold people accountable.

I think that is where the focus is going to be. People talk about whether the computer manufacturer or the software writer would be liable. No. I mean, if he's a rogue and puts something deliberately malicious in there, but not just in the way that the weapon comes out. I've written a piece on that, I'll send it to you. Good question.

Michael Klare: Can I just say one thing? This is not exactly on this topic, but I began my presentation by emphasizing the shift in the international political environment that we're in. Sometimes spending time in Washington helps with this. That for many people in Washington and Moscow and Beijing, it is evident that we've moved into an era of great power competition and conflict. The military leadership and the

political leadership sees this more intense competitive environment, in which war is a very realistic possibility among the great powers, and that technology is making it possible to conceive of attacks on, as you say, not just on individuals, but on the leadership.

President Putin in his speech a few weeks ago specifically said that, if the U.S. deploys weapons in Europe of the type that are banned by the INF Treaty, Russia will respond by deploying weapons off the coast of the United States intended to attack the command capacity, meaning the president of the United States and other leaders of the U.S. There is a more intense competitive environment in which the national leadership sees war as a very real possibility, and it's in that environment that all of these developments have to be viewed.

Tiya Maluwa: Yeah. Thank you very much. We are running out of time, hopelessly. We should continue the conversation over lunch, I hope. But could you just join me in thanking our panelists for their very enriching presentations.