

CONFERENCE REPORT

SPEED

28 - 29 SEP 2018 /  @ 

**CORNELL
TECH**

HOME OF THE
**JACOBS
INSTITUTE**



SPEED

28-29 SEP 2018 /  @ CORNELL
TECH

CONFERENCE REPORT

With contributions by the Digital Life Initiative Fellows and Guests

Erica Du

Elizabeth O'Neill

Sabiha Sadeque

Marijn Sax

Nirvan Tyagi

Lauren van Haften-Schick

Hosted by the Digital Life Initiative

Helen Nissenbaum

James Grimmelman

Michael Byrne

Generously supported by



Contents

Introduction	4
Panel 1 Fast, Cheap, and Out of Control	6
Why Speed?	6
Theatre of the Future: Autonomous Vehicles as a Test Case for Designing for Speed.....	7
Artifice and Intelligence.....	8
Panel Discussion	8
Panel 2 Content Moderation	9
Halfway Around the Internet: Moderation in the Age of Memes	9
Facebook v. Sullivan	10
A History of the Instant in Media and Message Exchange	11
Public Pauses: Sociotechnical Dynamics of Temporal Whitespace in the Networked Press	12
Panel Discussion	13
Panel 3 Warfare and Policing	14
Decelerating Destruction: Speed, Stability, and Fragility in Military Arms Races	14
Speed Kills: Time, Space and Meaningful Human Control in Autonomous Weapons	16
The Influence of Speed on Human Decision Biases in Supervising and Designing Autonomy ..	19
Panel Discussion	21
Panel 4 With All Deliberate Speed (Security)	22
Fast, Slow, and In-Between: Speed & its Discontents	22
Flow Control: Internet Traffic Management & the Uneven Distribution of Delay.....	23
Desirable Inefficiency	23
Panel Discussion	24
Panel 5 Labor and Manufacturing	26
Fast Labor and Faster Machines in the Gig Economy	26
Interlude: I Love Lucy	27
The Problem with Digital Market Perfection	27
Panel Discussion	28
Panel 6 Finance	29
Man vs. Machine: The Past and Future of Human Collaboration with Machines in Financial Markets.....	29
High Speed Trading Algorithms & Human Manipulations: Agency & Accountability in Complex Financial Markets	30
Speed II: Have We Reached a Tipping Point?	31
Speed and Reliability: Cryptocurrency’s Irreconcilable Opposites?	32
Panel Discussion	33
Appendices	35
Speaker Biographies	35
Conference Flyer	40

Introduction

Speed Conference began with a provocation: How do we sustain meaningful human and societal oversight of AI given radically different processing speeds? Over the course of the two-day conference, twenty speakers from diverse fields presented suggestions for addressing these issues.

Along with scale and complexity, speed is one of the defining problems of algorithmic oversight. Intelligent systems and human actors operate at vastly different speeds, and these differences present challenges for managing and responding to algorithmic decisions. Despite speed's importance, discussions about the topic have been isolated in areas such as robotics, finance, warfare, and online communication.

How do we sustain meaningful human and societal oversight of AI given radically different processing speeds?

Speed Conference, held at Cornell Tech on September 28 and 29 2018, brought together nearly a hundred academics and practitioners to identify common themes and potential solutions in areas of speed, AI, and algorithmic oversight. Six panels of experts presented their research in areas of autonomous vehicles, warfare, information security, labor and manufacturing, content moderation, and finance. Panel discussions invited thought-provoking questions from the audience.

Speed Conference's presentations revealed several overarching themes. The first theme was that of excitement and anxiety over technological speed – and the fact that neither are new. Jason Farman's history of the New York pneumatic-tube mail system identified familiar tropes of "speed" around a technology that now seems charmingly antiquated. Numerous other participants gave examples of "old" technologies raising familiar questions: ARPANET, the original faux automaton Mechanical Turk, and the Dr. Strangelove-era nuclear weapons explored by Helen Nissenbaum in her opening presentation.

A second theme was the use of delay as a deliberate design element. This was most explicit in Fenwick McKelvey's presentation on flow control in Internet transmissions, where delay has been used for personalized service, to allocate scarce resources, and to create an opportunity sell "faster" (i.e. not deliberately delayed) plans. Larry Tabb's analysis of "speed bumps" in financial markets showed that these design choices connect both to systemic properties (how quickly markets incorporate new information) and fairness among participants (which types of trader are favored). Paul Ohm and Jonathan Frankle cited IEX's "magic shoebox", and linked it to Bitcoin's proof of work and iOS's password entry timeouts as examples of "desirable inefficiency," which they see as levers to add human values to technological systems. Finally, Mike Ananny described how the institutional press uses temporal gaps as a way of structuring the process of delivering news, introducing pauses and delays to make human sense out of events.

Ananny's observations point to a third theme: the way in which technological speed must be constructed and assimilated by human participants. Wendy Ju conceptualizes it as a kind of theater, in which autonomous vehicles perform for audiences of passengers, pedestrians, and other drivers. These performances must create the right social cues to negotiate interactions. Steve Jackson, on the other hand, observed that the apparent "speed" of some technologies, like ride-sharing apps and smartphones, is made possible by hiding the much slower timescales on which support and maintenance work takes place. Finally, Mary "Missy" Cummings observed that the increased speed of unmanned weapons platforms can have the counterintuitive effect of enabling *more* time for deliberation by more human participants. Taking one human (the pilot) out of the loop removes constraints on the time available for reflection before acting on a targeting decision.

Speed unleashes competitive pressures that warp human social systems, as black holes warp space and time.

The fourth theme, raised by several panelists, was that speed unleashes competitive pressures that warp human social systems around them, as a black hole warps space and time around it. James Grimmelman identified content moderation as one such domain: the pressure to be fast and first with viral content creates the conditions under which fake news flourishes. Argyri Panezi points to similar processes in gig economy platforms, which can force workers into precarious working conditions (illustrated by the famous *I Love Lucy* candy factory clip). Frank Pasquale's analysis of

AI and autonomous weapon systems as a (doubly literal) arms race was the most chilling example of this dynamic. Finally, Simona Abis's presentation on investment funds showed that strategic use of speed in trading does not necessarily translate into performance advantages.

Finally, panelists worried about the systemic consequences of acceleration. Rory van Loo pointed to speed-abetted market "perfection" as a source of instability. Andrea Matwyshyn presented a series of vignettes about how AI systems interact with people and with each other to show that speed's feedback loops can be chaotic and unpredictable. Katie Brennan invited participants to think about "assemblages" of high-frequency trading systems as sources of emergent behavior not easily traceable to any single actor. Michael O'Connor's presentation on blockchain technologies examined governance tradeoffs between different systemic goals: encouraging broad participation, processing transactions quickly, and minimizing the environmental footprint.

The overriding concern throughout the conference was how best to achieve what Peter Asaro called "*meaningful* human control": the exercise of properly informed and properly responsible moral agency. From weapons systems and vehicles to financial markets and communication platforms, careful consideration of agency and accountability is vital when deliberating the future of speed.

Panel 1 | Fast, Cheap, and Out of Control

Reported by Lauren van Haaften-Schick, DLI Doctoral Fellow

Why Speed?

Helen Nissenbaum, Cornell Tech

Why *Speed*? Helen Nissenbaum, Director of the Digital Life Initiative at Cornell Tech, started the day describing the impetus behind the conference’s provocative title. As Nissenbaum explained, a focus on *Speed* draws our attention to the temporalities of emergent technologies and the ways in which our interactions with those technologies are impacted by the ever-increasing *speed* of their development and of computing itself.

We are, she observed, in an intense moment of growth in Artificial Intelligence, with countries and corporations alike competing and collaborating in investing in AI development and applications. AI-driven automation is increasingly controlling material systems and decision making, ranging from advertising to healthcare. Thus far many of the public conversations around AI have been directed at function and performance, dominated by questions like “Will these systems work well enough? Will AI work as well as humans and human experts? Or as well as smart systems are already working?”

Beyond these questions, Nissenbaum urged that we also must consider the impact of AI – and the impact of its literal *speed* – upon *societal values, cultural and political values, and quality of life*. This second tier of concerns forces the stakes of the conference’s overall inquiry to surface: “Who are going to be the winners and losers in the future of AI?” “What about accountability?” How do we move popular attention past the spectacle of sci-fi scare scenarios around “machines run amok” to address these actual social concerns? And how do we move past the assumption that a “human in the loop,” or human oversight or intervention, would be sufficient protection against the over-determination of AI, when human oversight must also be meaningful and not just a token gesture?

Nissenbaum re-grounded this web of provocations back to the concept of *speed*, and the radical difference in speed between AI and humans. Here, Nissenbaum’s hint of a fictional nightmare scenario re-entered via reference to Stanley Kubrick’s *Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb* (1964). The film’s plot hinges on tensions between human ethics versus bureaucratized control in decision making – here regarding the decision of whether or not to deploy an atom bomb, who should possess the power to make that decision, and the *speed* at which a human actor might make that decision impulsively.

It would seem, given the severity of the choice of whether to deploy such a weapon, that a rational and ethical human actor or actors would be the desired authority for such a decision, perhaps even with the risk that a human would abuse their power. But, Nissenbaum also reminded us that the scenario of an actual attack may require a *speed* of decision-making that humans are not capable of,

or rather, that they are not equipped to undertake in a manner that could be reached via the deliberation of ethics. The solution here presented by computational decision-making and speed was invoked by Nissenbaum to reference the legacy of historical notions of computer science for the public good, and also served to lay out the dense layer of questions that would unfold over the course of the next two days.

Theatre of the Future: Autonomous Vehicles as a Test Case for Designing for Speed

Wendy Ju, Cornell Tech

Next on the panel was Wendy Ju, Assistant Professor at the Jacobs Technion-Cornell Institute at Cornell Tech in the Information Science program. Ju's presentation "The Theater of the Future" introduced us to her work on autonomous vehicles and developing techniques for studying human use and interaction with them. Ju's research is guided by the premise that "in order to study the way people will interact with new technology we have to *simulate the future*."

Most of Ju's work in this area has concerned driving simulators, which are "theaters for one person" where the surroundings one might encounter on a real drive – suburban streets, pedestrians, open roads, etc. – appear on a concave screen that engulfs the test "driver's" vision. Over the course of Ju's research, the simulators themselves began to reveal unexpected information about drivers' habits and physical needs: test drivers would often fall asleep and lose interest, situations that would pose serious danger in real life. Thus, rather than testing the safety of the car, driving simulators would also have to address the safety of the driver, and account for ways to keep them entertained and awake so that they would be able to respond should a problem with an autonomous vehicle occur.

The surprise of autonomous test drivers falling asleep led Ju to a new research problem: How can we envision the unexpected? One "theater" for tackling this question for autonomous vehicles is found in monitoring pedestrian interaction. How will people react when these machines are on the road, and more critically perhaps, when they are being introduced and thus not yet familiar or normalized? How will they interact with them? And, importantly, how will the people "who didn't sign up for this future" react?

Ju's presentation concluded with a welcome moment of self-reflection on the resources required to do this kind of research, and offered a model for how less-funded universities and even lay people could create a driverless car simulator using readily available materials. Circling back to her introduction, Ju ended by stating that theater has the power to give us alternate realities, and can thus be a vital tool in testing just how that future might look.

Artifice and Intelligence

Andrea Matwyshyn, Northeastern University

The final speaker was Andrea Matwyshyn, Professor of Law and Co-Director, Center for Law, Innovation and Creativity, at Northwestern University, speaking on “Artifice and Intelligence.” Matwyshyn’s talk focused on the issue of translating technology in different ways, emphasizing that “the tools used in translation are just as important as the outcomes or policies themselves.” The presentation was segmented by five “grim fables” which framed Matwyshyn’s observations on translating technology.

The first ‘fable’ illuminated that what people mean by ‘AI’ varies, and choices in builds and governance matter. Given the inevitability of difference in interpretation, organizations and individuals developing AI tools need to ensure that they are doing so with shared goals and, as much as possible, a shared language. The second fable concerns the degree of human intervention in AI, and the fact that many products marketed as AI-driven are indeed driven by humans behind the scenes. The third fable concerned the problem of short-term fixes that can actually result in long-term problems, leading to an even more wide-ranging concern: Just because we can build something, do we need to? The fourth fable – illustrated by a pair of talking electronic hamsters caught in a recursion loop – brought us to the question of the limits of humans to predict AI’s capacities. We seek uniformity and predictability in algorithms, but forget that even our best designed plans can contain overlooked flaws, or that no two machines are truly entirely identical. The fifth tale, entitled “the infiltrating squirrel”, reminded us that factors beyond our control can always present hiccups and wrenches, as demonstrated by the high number of major technology infrastructures that have been taken out by small animals.

The cumulative lesson in these fables is that we need to maintain an awareness not only of the unpredictability of AI, but that as AI plays a role in connecting us to “commodities,” to “other humans,” and now to “the internet of things” we need to remain aware of the things that could go awry, and that get lost in translation along the way.

Panel Discussion

Moderated by Helen Nissenbaum, Digital Life Initiative Director

Audience questions were brief due to time constraints, but introduced an element of the panel title that had gone largely unmentioned: “cheap.” Each of the panelists reflected that developments in AI technologies would force changes in the supply chain, but, maintaining each of their self-critical ambivalence, recognized that we needed to think carefully and openly about the ways in which markets and manufacturers (and perhaps consumers and employees) might find unexpected ways of organizing themselves to meet new conditions within the radically shifting temporality of AI’s speed.

Panel 2 | Content Moderation

Reported by Erica Du, DLI Visiting Fellow

Halfway Around the Internet: Moderation in the Age of Memes

James Grimmelmann, Cornell Tech

Dr. James Grimmelmann spoke about how speed distorts the spread of ideas on the internet. In particular, he focused on how memes are fueled by velocity and ambiguity. He defines velocity as virality from one user to another, and ambiguity as how the same message can come across differently based upon the context and the readers.

Grimmelmann drew analogies from classic science fiction novels illustrating colonies in space, citing the inevitability of these societies evolving in ways that are unpredictable, and often unstable. Similarly, he demonstrated how natural rates of change on the internet can also evolve and take the user to strange places. For instance, companies like the New York Times now have the ability to rapidly A/B test headlines, designing their algorithms to be ruthlessly efficient in order to maximize the number of eyeballs on their content. When society combines these techniques with the desire to profit, the consequences can be quite dark.

One example arises from following the chain of recommended Peppa the Pig videos on YouTube. Though the initial videos are innocuous, pursuing the trail may eventually lead children to twisted, graphic, or violent content. Because the aim for the algorithm is to capture attention at any means possible, it will flood the user with any content with the potential to go viral, even when this very content is scarring to them.

Another example exists in e-commerce, where businesses will create and market products pandering to absurd trends. Infrastructure now exists such that the market can immediately capitalize on whatever people are into. This also allows businesses to use recommender features on e-commerce platforms like Amazon to manipulate consumers and push specific viewpoints, like anti-vaccine campaigns.

Grimmelmann also pointed out the purposeful ambiguity of internet trends. He illustrated how it quickly becomes really hard to draw lines between the thing itself, criticizing the thing, making fun of the thing, and paying attention to the thing. When content is always both a joke and not a joke, it becomes hard to criticize. He cited the Tide Pod craze in particular, where a simple caution of “Don’t eat tide pods” started trending as people began to cash in on the craze. By simply paying attention, the internet gave this idea power.

Furthermore, platforms need to make content moderation decisions fast – as quickly as users are producing content. Ambiguity causes this problem to be even more difficult than it already is.

Last, Grimmelmann discussed the role of controversy. During the 2016 election, there was a piece of fake news claiming the Pope endorsed Trump. Although it's uncertain if people really believed it, natural market dynamics revealed that because people wanted to read content of this nature, even if false, suppliers were willing to write it to monetize the ad dollars.

Grimmelmann argues that ambiguity is what gives memes the ability to move fast. These images are capable of holding multiple meanings, and often consumers do not agree on what these images mean. The images take off because consumers can decide meaning for themselves. Consumers can remix memes, taking advantage of various contexts and viewpoints as the artifact itself has no fixed meaning.

There is an enormous capacity for play and creativity in this space, though also an enormous amount of danger that can result from misunderstanding, like the pizzagate incident.

The audience was curious about the right way to interpret images. Was it moving from interpreting images in isolation to looking at a network of images? Grimmelmann agreed that we can only understand an image in the context and culture in which it comes, but was unsure if our task was to definitely understand these things. The audience was also curious as to how to reconcile the protected category of hate speech and borderline offensive speech. Grimmelmann argued that the basic problem is having a vaguely defined, general purpose platform, where the only line moderators can police is "hate speech". In these scenarios, he argued there would always be problems with how to define hate speech. He proposed a topically constrained platform, where moderators can be very strict about the focus of conversations.

Facebook v. Sullivan

Kate Klonick, St. John's University Law School

Kate Klonick unpacked problematic attempts at content moderation strategies made by courts and Facebook, in their efforts to define "newsworthiness" and "public figure."

Klonick illustrated the issue with an anecdote from right before the presidential election in 2016. Mark Zuckerberg, the CEO of Facebook, held a closed town hall meeting discussing Facebook's moderation policy, which was sparked by allegations of Trump using Facebook to post "hate speech." Zuckerberg ultimately claimed that Trump's status as a public figure deemed this speech "newsworthy", and therefore important to public interest and acceptable to keep on the platform – even though the speech violated Facebook's policies.

Klonick likened this situation to the 1960 Supreme Court case *New York Times v. Sullivan*, except the "Sullivans" of today are public figures. She cited *Gertz v. Robert Welch's* identification of 3 types of public figures: general purpose public figures who command persuasive power or influence, limited purpose public figures who intentionally thrust themselves to the forefront of a particular controversy, and involuntary public figures, which the court left undefined. Dr. Klonick shared that

when she interviewed employees of Facebook, the reasons they give about their content moderation policies strongly resembled Sullivan.

As an example, anti-cyberbullying groups advocated for more policing on Facebook in 2009, but Facebook struggled with how to define the rule for bullying. Facebook argued that the crux of bullying was in an unhealthy power dynamic, and that power dynamic was hard to derive from purely looking at content. They instead decided to take a passive approach, taking down content that was flagged for bullying. However, an exception to this policy was made for public figures. How were public figures defined? An employee would use Google news, search for the person in question, and decide based on search hits.

The idea of newsworthiness, too, has been debated at Facebook. One controversial decision followed the Boston Marathon bombings in 2013, after gruesome photos of victims were posted on Facebook. Facebook initially decided to take these photos down because they met the definition of “gore,” where the “insides were on the outside.” However, Facebook later put these photos back up because of their newsworthiness. This contrasted with the decision to take down photos of a horrible beheading in Mexico that had happened shortly before. The inconsistency of these ad hoc decisions led to questions and concerns around Facebook’s process in determining what was “newsworthy” and what was not.

Klonick left us with three main points. First, that the problem with talking about Facebook news algorithms is that they will always over and under censor. Second, social media has led to the development of the involuntary public figure, and it is counterproductive that the court does not recognize their existence. Last, looking at court systems and public figure doctrine, coupled with Facebook’s behaviors in the past years, Facebook has emerged as both the court and the New York Times, acting as both players simultaneously. Its moderation policies are inconsistent, taking down videos from some political extremists and not others. In order to create these changes, Klonick advocated discussing what we, as a society, want these policies to be.

The audience noted that Facebook seemed to be using the First Amendment in a perverse way to protect the speaker, and those with more power, which Klonick agreed with. An audience member also asked about validation for these policies, to which Klonick replied that Facebook, as a private company, technically does not require validation or approval from anyone external to the business.

A History of the Instant in Media and Message Exchange

Jason Farman, University of Maryland, College Park

Jason Farman discussed the history of delays in message exchange, focusing predominantly on the pneumatic tube system in New York City in the early 1900s. He is particularly interested in interpretations of wait times in “instant” messaging, arguing that time, speed, and delay are all content.

Farman began with a historical overview of the pneumatic tube mail delivery system. Pneumatic tubes were placed underground and connected people by giving them a seemingly instant way to send packages and mail to recipients across the city. This took around 18 minutes by tube, faster than any truck at the time. Their placement underground was both practical and magical, avoiding disrupting anything above ground, while also allowing imaginations to cultivate a sort of mysticism. Farman proposed that the invisibility of these technologies led to the feeling that communication needed no medium, and feelings that sender and recipient could connect mind to mind. At the time, these tubes were the technology of the future, and a symbol of cosmopolitan life.

The pneumatic tube system was decommissioned in 1984, as popular perceptions of the technology's alignment with expectations of time and speed shifted. Farman argued that the definition of *instant* is contrasted with the definition of *delay*, and those definitions are specific to a single point in history, which often do not have staying power. He proposed that what it means to *wait* changes from era to era, that waiting is experienced in context, and not experienced as the actual amount of time one spends waiting, whether it be seconds, minutes, or longer.

Technology constantly adapts. During the Dot Com boom, the internet became the newest medium of instant communication, and DSL provided a mechanism to connect to the internet, even if quality degraded as people moved away from the connection. This illustrated how even if tech does continuously accelerate our pace of communication, there will still always be a time lag, or times where users will be waiting. Farman argued that this context would always shape how we view and interpret wait times.

Public Pauses: Sociotechnical Dynamics of Temporal Whitespace in the Networked Press

Mike Ananny, University of Southern California

Mike Ananny shared a mashup of two main ideas he has been investigating: first, listening and hearing as part of a good and healthy press; and second, whitespace press, where pauses, silences, and absences act as sociotechnical constructions.

To begin, Ananny used three anecdotes to illustrate how media pauses are made and where these constructions come from. First, media pauses can come from elite journalists deciding that nothing needs the public's attention. For example, on April 18th, 1930, a headline declared that "There is no news." Second, private organizations can use their marketing environments for pauses, much like how Lifetime Fitness removed news broadcasting from their workout environment due to complaints. Third, groups can temporarily hide from media to protest or build communities, similar to how students at the University of Virginia organized a social media blackout as a response to violent racial conflict on campus in August 2017.

In this industry, pauses may seem as incidental and unimportant, as privilege goes to production of content, not to its absence. However, Ananny claims that collective self-governance needs more than just a marketplace of speakers: it also requires a public right to hear, and a value of listening.

To create these pauses, Ananny looked to past sources of “pauses” in journalism. He first discussed censorship, the removal or replacement of specific content, which essentially prevents authors from reaching their intended audiences. He also mentioned self-censorship. Examples include Israeli journalists working with state representatives to limit investigative reports, U.S. journalists embedded within the military being selective about what they publish at certain times, or even Bill Keller agreeing to then-President Bush’s requests to delay reporting on CIA prisons. Other times, reporters will self-censor their content if they know a competitor is covering similar content. Reporters may also take into account rituals or norms, which explains underreporting of HIV/AIDS and lower frequencies of suicide reports. Hiding source identities in radio, as well as silence on social media due to a shared tragedy (like the 2014 world cup shootouts), were also described as sources of journalistic pauses.

Last, Ananny advocated for the importance of these pauses, because of their ability to act as socio-technical constructions and socio-technical diagnostics. He proposed that understanding and designing for pauses may help create listening publics.

The audience questioned the plausibility of pauses in an age of pervasive social media, never-ending twitter streams, and technology. Ananny cited a few examples where the press purposely has been using pauses and constrained content in their user experience design to differentiate themselves. While the web version of the L.A. Times employs continuous scrolling, publications like the Economist and the Wall Street Journal do not, to emphasize a sense of completion and finality. These are both moves by news organizations to control time.

Panel Discussion

Moderated by Mason Marks, DLI Research Fellow

There was time for a final round of questions for the panel after the last speaker finished. First, the audience challenged the panelists to think normatively – what *should* happen, what do we want, and how do we get there? The panelists discussed the current paradigm of parties exercising power by controlling time and rhythm, and how the very ability to determine at what time one speaks is power. They also highlighted the power content platforms have to rank and prioritize content. Klonick noted that there would always be some sort of content curation, but that companies should focus on understanding and being intentional about the values they put in place. Last, an audience member pointed out the multiplicative impact of combining scale and speed, asking if a lot of harms would disappear if they were to choose between breaking things up or slowing things down. The panel countered with the example of internet shaming, which felt mostly like a speed issue. They also argued that diversity was as important as breaking things up when counteracting scale problems.

Panel 3 | Warfare and Policing

Reported by Marijn Sax, DLI Visiting Fellow

Decelerating Destruction: Speed, Stability, and Fragility in Military Arms Races

Frank Pasquale, University of Maryland

In this talk, Pasquale uses social theories of speed and acceleration to address AI warfare. He will discuss existing social and legal theory before addressing its implications for what can and should be done to regulate AI warfare.

Pasquale begins with a reference to *the* classic social theorist in this domain, Paul Virilio. He explains how speed is essential to warfare, as, for instance, the role of speed in nuclear warfare exemplifies. As states attempt to create nuclear weapons that can strike as fast as possible, the windows for retaliation become increasingly small, “thus finally abolishing the head of state’s power of reflection and decision in favor of pure and simple automation of defense systems.”

Nowadays, autonomous weapon systems (AWS) are defended on multiple grounds. By deploying AWS, the promise is that the military is better able to respect basic moral norms and international humanitarian law by targeting attacks to minimize casualties. It is still an open question how lethal autonomous weapon systems (LAWS) will develop. Pedro Domingos has hypothesized a Jurassic Park-like Robotic Park, where many different types of LAWS are pitted against each other to decide what the ultimate LAWS is. (According to Domingos, he suggested the idea to someone in the US military who replied “that’s feasible.”)

The idea of speed plays an important role in this new development and exploration of LAWS. In line with Virilio, it is suggested that the faster the US can develop AI tools for warfare, the more power it affords them. Similarly, Vladimir Putin has said that “Artificial intelligence is the future, not only in Russia, but for all mankind [...] Whoever becomes the leader in this sphere will become the ruler of the world.” Understanding the role of speed in this manner will lead nations to (try to) develop AI technology for warfare as fast as possible.

In response to this looming arms race, there are Abolitionists who call for the need to define categories of weapons for which development and proliferation should be stopped. Abolitionist initiatives have already started a multilateral process at the UN to define and ban killer robots (as has been done with biological weapons). Reformers take a somewhat different approach. Reformers emphasize that killer robots are already here, and that we are too late to stop them. Instead of working on bans like the Abolitionists, Reformers insist on the need to *regulate* (the use of) killer robots.

In response to both the Abolitionists and Reformers, Realists insist on the futility of strong controls on a (potential) arms race between great powers (i.e., nation states like the US and Russia). The standard Realist story is as follows: Abolitionists and Reformers promote international agreements. It is close to impossible to verify and enforce such agreements. Great powers thus have a strong incentive to publically propagate adherence to international agreements, while secretly continuing to develop LAWS. Because of a fundamental lack of trust at the international level, an unstoppable arms race is bound to ensue. In this arms race, speed is very important (as the word race already implies).

Inspired by Hartmut Rosa's book *Social Acceleration: A New Theory of Modernity*, Pasquale asks whether speed *itself* is a condition, or a causal agent. Put differently, is speed itself part of the environment (in which AI warfare is developed and pursued), or something that changes the very environment? To tease out this question, Pasquale discusses four implicit binaries found in the AI warfare context: competition vs. cooperation, race vs. stability, speed & acceleration vs. slowness & deceleration, automation vs. deliberation. Why, Pasquale asks, is the first half of each of these binaries (i.e., competition, race, speed & acceleration, automation) so compelling in the military context? For an answer, he turns to Rosa. Many military doctrines such as (1) Chinese 'Unrestricted Warfare,' (2) Russian information war (e.g., on Facebook), and (3) the US Defense Advanced Research Projects' theories of dominance via encryption and quantum computing, all appear to have both slow and fast elements. They can be understood as slow wars because they are perpetual and drag on forever. They can be understood as fast because "each small conflict or advance presages the general conflagration, which reaches back from the future unto us." The key insight here is an interesting paradox: everything seems to be changing and developing faster, but the structures that guide the change seem to become more rigid.

Building on Zygmunt Bauman's *Liquid Modernity*, Pasquale suggests that speed can also be understood as the intensification of an activity into its disappearance. There is, for instance, the suggestion that college education must be more efficient; shorter and faster. But how short and fast? There comes a point where the insistence on an ever shorter college education undermines the very idea of a college education. Acceleration can *drive* change, but also *break* things. So what would intensification into disappearance mean in the context of warfare? Pushbutton war where an entire war is reduced to the push of one button?

Bauman's and Rosa's ideas can thus be applied to the AI warfare arms race. The paradox that the increase in speed and change also leads to more rigid structures can be reinterpreted in the arms race context: by keeping some persons, firms, and nations locked in certain competitions, some powerful actors benefit. So what can this paradox of rigidification learn from the Realists? An arms race may not only be a strategy for domination by the strongest, but at the very same time also lock us in a rigid system of wasting resources. Pasquale explains the metaphor of fighting stags with big antlers. The stag with the biggest antlers has the upper hand in a direct battle with another stag. At the same time, very big antlers are unproductive for everything else a stag normally does (e.g., navigating through a forest is much harder with very big antlers.) Referring to Thomas C. Schelling

(*Arms and Influence*), Pasquale says that no country can spend 100% of its budget on defense and war. So even for this practical reason, we need to think more about getting out of the (unproductive) arms race. Understanding the paradox of acceleration and rigidification is a good beginning.

Question from audience: Can't the arms race also be understood as productive (as opposed to wasteful) because it spurs the development of new technology?

Answer: Pasquale acknowledges that the question exposes the one major weakness of the talk. He agrees that many useful technologies are actually spin-offs from technology developed by the military. For example, your smartphone is filled with military spin-off technology. So the real challenge is to think harder about how we can create research into (military) technology that is most likely to result in productive spin-offs.

Speed Kills: Time, Space and Meaningful Human Control in Autonomous Weapons

Peter M. Asaro, The New School

Asaro begins his presentation with two quotes from Paul Virilio's *Speed and Politics*:

"In fact, history progresses at the speed of its weapons systems." (p. 90)

"With the realization of *dromocratic* [dromos means "racetrack" in Ancient Greek] type progress, humanity will stop being diverse. It will tend to divide only into hopeful populations (who are allowed to hope that they will reach, in the future, someday, the speed that they accumulating, which will give them access to the possible – that is, to the project, the decision, the infinite: speed is the hope of the West) and despairing populations, blocked by the inferiority of their technological vehicles, living and subsisting in a finite world." (p. 70)

As both quotes show, Virilio puts speed at the center of concepts of power. The importance that is attributed to speed can also be seen in Futurist art (e.g., *Dynamism of Train Ship Airplane*, Giulio D'Anna, 1930), and in the work of Futurists more generally, where speed is *fetishized*. For example, Filippo Tommaso Marinetti wrote in 1909 in his *Fondazione e Manifesto del Futurismo*, "We declare that the splendor of the world has been enriched by a new beauty: the beauty of speed." Even around 500 BCE, the importance of speed was acknowledged by Sun Tzu, who wrote that "Speed is the essence of war." Asaro claims that this acknowledgement of speed's importance throughout history is, nowadays, still one of the core ideas behind autonomous warfare technology. As a case in point, Asaro refers to a passage from four-star General John R. Allen's book *On Hyperwar*, where the importance in speed is also emphasized: "Until the present time, a decision to act depended on human cognition. With autonomous decision-making, this will not be the case. While human decision-making is potent, it also has limitations in terms of speed, attention, and diligence." Here, again, one sees the Virilioean notion of speed being connected to (military) power.

To analyze the role of speed in autonomous weapons systems (AWS), Asaro first explains how AWS can be defined. To understand the role automation plays, it helps to consider the Kill Chain: Find, Fix, Track, Target, Engage, Assess. If *every* part of the Kill Chain is automated, including the targeting and engaging, one has an AWS. Often, these are also referred to as *lethal* autonomous weapons systems (LAWS) since the lethal decisions are automated. Most of the discussions in this domain revolve around the idea of ‘humans-in-the-loop,’ and the question of how much (if at all) humans should be involved in especially the target and engage stage of the Kill Chain – the parts of the chain where the actual lethal decisions are made.

AWS raise a range of potential problems, many of which are rather obvious. Asaro mentions some of these, such as: mistakes and accidents, conflicts as a result of unintended AWS attacks, the lowering of thresholds for conflicts because warfare is cheaper and poses less risk, the rapid proliferation of AWS (including to non-state actors), the risk of an arms race, the (potential) unattributability of attacks, the risk of a swarm of AWS turning into a weapon of mass destruction, the vulnerability of AWS to hacking, spoofing, and cyberattacks. There are also some less obvious, but important Kantian moral and legal questions to be asked: how to deal with accountability gaps when no human has literally pushed the button leading to a kill, but we still want to think of human beings being accountable for killing. There are also important questions regarding human dignity: what does it say about the value of humans and a human life when we let robots make decisions about which lives end?

Next, Asaro asks *why* we would like to automate lethal decisions. Often, the following reasons are mentioned: AWS do not need constant communication, meaning they also do not suffer from threats of jamming of communications by enemies. AWS can also be deployed underwater, where communication is much harder, even without the presence of jamming attempts by enemies. The performance of AWS is also thought to be “better” in term of precision, accuracy, information ingestion, and because they do not depend on human psychology which is often faulted for errors.

However, such rather practical considerations tend to overlook fundamental moral and legal considerations. To start, decisions to kill are exceptional. Killing is prohibited, unless we can offer very good justification. However, to be a moral agent, one must have mental access to those justifications. Machines lack such mental access. We can of course try to *express* moral rules through them, but that does not make the machines *themselves* moral agents. Moreover, AWS cannot *recognize* (in the moral sense) humans *as humans* with dignity and as part of a moral community. If AWS are used to kill nonetheless, we could be dealing with what Asaro calls ‘improper delegation.’ Moral and legal agents are subject to legal and moral accountability. But because AWS do not have mental access to the justifications required for moral agency, they cannot be subject to moral and legal accountability (at least not in the way that humans as moral agents are).

At this point, proponents of AWS often point to the idea of humans-in-the-*wider*-loop, i.e., including humans as moral agents through their coding (which determines how the AWS go about their killing). This, however, is not really the same as letting humans make the *actual* decisions to kill;

preprogramming rules based on which someone might get killed is not the same as actually being a moral agent who, at the moment it matters, makes the actual decision to kill. Being a moral agent is not simply transferable in this proposed way. On top of this, there is the more practical but important problem of the inherent complexity of predicting actions and effects, especially in the chaotic and messy context of war.

Asaro wants to emphasize, however, that he does not wish to romanticize human deliberation. Our current military practices are certainly not perfect, and humans are certainly part of the explanation of why they are not perfect. There is also a long history of atomization and routinization in the military, as attempts to deal with the inherent fallibility and biased nature of human deliberation. Despite these faults, humans do bring something unique to the table: situational awareness and the ability to reason about contextual meaning.

In thinking about ways forward, there is a lot of talk about *meaningful* human control. This term has become a vessel for many different ideas, many of which focus on the general question of what kind of morality we want for AWS. Asaro's most important suggestion is that warfare technology must allow people to *exercise* moral agency, instead of improperly delegating moral agency. More specifically, it is important that the technology provides moral agents with the (type of) information necessary for the moral judgements that are required for justifying killing. From the perspective of responsibility, the warfare technology must make sure that conditions are created where operators can feel confidence in actions that are sufficient for accountability. Lastly, creating conditions that allow the exercise of moral agency and actions that are sufficient for accountability does not only come down to the technology itself, but is a property of the larger socio-technical system which also includes the chain of command, training, and maintenance of the system.

The reason this discussion is so challenging is that technology is developing so incredibly fast that our norms cannot keep up. So how can we innovate our norms in a time where the "speed of tech" is unparalleled? This is a very difficult question because normally norms get settled *after* a practice emerges. Now we are asked to innovate our norms before a practice emerges or settles. Additionally, *whose* values should be reflected in norms? Who are the stakeholders and what role should they play?

Question from audience: How does your argument relate to *defensive* systems, which also sometimes inflict collateral damage?

Answer: Asaro explains that he makes a distinction between systems that are explicitly designed to kill, and systems that are designed to defend (and which are *not* explicitly designed to kill) such as missile defense systems. Precisely because they are not designed to kill persons, they are not as morally problematic as AWS.

The Influence of Speed on Human Decision Biases in Supervising and Designing Autonomy

Mary "Missy" Cummings, Duke University

Before receiving her PhD in 2004, Cummings was a fighter jet pilot for the US military for over 10 years. This affords her with not just an academic perspective (as Professor in the Department of Mechanical Engineering and Materials Science) on automated and autonomous warfare. Because she has experienced the use of (advanced) technology in warfare situations first hand as a pilot, she also has a lot of practical knowledge on the subject.

In this talk, Cummings wants to provide a framework to help us better understand the role AI plays in warfare. She encounters a lot of hysteria around the topic, mostly created by people who do not have intimate knowledge of the technology (as Cummings herself does have). The framework can help provide clarity amidst the hysteria.

Cummings first discusses the OODA loop, which stands for Observe, Orient, Decide, Act. The OODA loop is designed for (and often just in the) military context, but applies more broadly to everything we do. Roughly speaking, one first has to observe a situation, then orient the situation by recognizing all the elements relevant to a decision, then have all the relevant elements inform the decision before actually acting. Another important framework, more specifically focused on combat is the Kill Chain. (Although Cummings emphasizes that, again, the Kill Chain is not *only* applicable to combat situations. The Kill Chain would also work in the context of, for instance, pesticide use for farming.) The Kill Chain consists of six steps: Find -> Fix -> Track -> Target -> Engage -> Assess (-> Find -> ...). Importantly, *time* compresses the kill chain in military situations: when the pressure is high and there isn't much time for making a decision, it causes decision stress. People working in such situations are affected by decision stress, as a forthcoming paper by Cummings illustrates. In the paper, she details studies that show how under decision stress, automation bias leads people to not (properly) assess the suggestions made by automated systems which causes them to make errors.

Cummings argues that people often fail to realize that drones are so attractive precisely because they can improve (parts of) the kill chain. Simply put, the deployment of drones can take (parts of) the stress away, allowing for better informed decision-making in combat situations. Here, Cummings can refer to her own experience of flying a fighter jet in combat situations: it is very stressful and comes with many things that can make errors more likely. For example, before engaging in combat, one often has to fly for hours before reaching a target. This costs energy (both mental and physical) that cannot be used later on, during the actual combat. She also explains how prior to missions, she would be shown pictures of targets during briefings which she then would have to memorize (for hours) before actually reaching the target. These are just two examples of things which can introduce faulty judgments and errors.

Given these human limitations, the introduction of drones allows for much calmer decision-making. For example, because the flying is (largely) automated, the pilot flying the drone is not exhausted (as

much) prior to the actual combat. The pilot also does not have to memorize targets, since the pilot is not restricted to the very small cockpit of a fighter jet without space for documents or the possibility to review documents. Another advantage is the fact that the actual final decision-making does not have to come down to the one pilot in the fighter jet. When a drone is deployed, an entire committee can in principle be involved in the decision-making. It is important to emphasize that in the example of drone deployment, humans are still very much in the loop. When it comes to *targeting* and *engaging*, the drone will not make autonomous decisions. Cummings thus concludes that the introduction of drones has, in this sense, *slowed down* the kill chain. Based on her own experiences, Cummings strongly believes that this results in significantly less collateral damage during missions. Drones allow for a better informed decision loop.

So how do the OODA Loop, the Kill Chain, and our current use of drones inform our thinking about the future of automated and/or autonomous technology in warfare? In order to strike the right balance when introducing AI in warfare technology, Cummings introduces another framework. This framework differentiates between skills-based, rules-based, knowledge-based, and expertise-based tasks. Roughly speaking, the introduction of automation through AI is most likely to succeed in skills-based scenarios, becoming increasingly challenging to implement as we move through rules-based, knowledge-based, and expertise-based tasks.

Skills-based tasks are those that take years of training to master; through training they become in a sense ingrained in us. Once we master them, we can (almost automatically) perform the tasks. For example, when one has trained and practiced flying for years, one can (more easily) change between different cockpits and still be able to fly. At this level, automation through AI can certainly be implemented. Think of drones that can (almost completely) fly themselves: there is nothing inherently human about being skilled at flying. Another example are commercial airplane pilots who barely touch the stick while flying and are also encouraged *not* to touch the stick unnecessarily because autopilot is more energy efficient.

At the next level – rules-based tasks – things become more complicated. Humans can internalize rules and then apply them. In this basic sense, automation through AI could similarly apply the rules it has been given. However, humans also internalize rules, freeing up cognitive resources for reasoning about (the application of) rules. Here, the knowledge-based tasks immediately become relevant, since rules introduce the need to reason about them. Cummings refers to an example of driving a car where a ball suddenly rolls over the road in front of the car. As humans, we know how to quickly create a narrative (e.g., a child was playing with the ball, lost control over the ball which caused the ball to roll on the road, meaning that there is probably going to be a child running after the ball onto the road) which allows us to guide and structure our thinking. Such reasoning about situations where rules are internalized, but also require application are challenging to automate properly. Cummings also refers to the flying of a fighter jet, which is an incredibly dynamic activity where there are constantly new situations emerging which require new exercises of reasoning.

This last example brings Cummings to the last level of expertise-based tasks, which are tasks that have to be performed at the highest level of uncertainty. Here, she refers to the commercial airplane that was landed on the Hudson by its pilots. This situation defied regular applications of skills, rules, and knowledge. At this highest level of uncertainty, automation is also most unlikely to work.

Cummings thus concludes that there are certain categories of action where automation is preferable. However, as we ‘move up the curve of uncertainty,’ automation becomes increasingly less feasible.

Panel Discussion

Moderated by Dan Lee, Professor, Cornell Tech

Question for Asaro: Given the fact that we need norm innovation in times of technological change, and given the fact that norms tend to change slower than technological innovation, how can we think about the need to innovate our norms *quickly*? Can norms change quickly? Don’t we need to change the rules first and allow the norms to catch up later?

Answer: One could take an international law perspective and think of norms that evolve in state practice, which is in turn guided by rules and laws. Norms are more generally about what society should look like, so norms should be understood as concrete things/tools to use in our thinking about society.

Question for Pasquale: In discussing Putin’s vision on the importance of AI, you suggested Putin was not broadcasting (all of) his true view. Why?

Answer: Pasquale refers to Strauss’ distinction between esoteric and exoteric writing and speech (see *Persecution and the Art of Writing*, where Strauss explains how the very same words can have different meanings to different audiences, and how writers and speakers can knowingly formulate their message in a way that addresses both audiences differently with the very same utterances/sentences, MS.) For Putin, there is an advantage in setting the rules of the arms race, of discrediting experts. With the statement, Putin does not attempt to describe the world, but rather attempts to *create* the world!

Question for Cummings: Have you come across pilots who tinker with the settings of the plane before taking off? Is this evolving (due to technological developments), or not?

Answer: To answer the question, Cummings breaks it down in two separate questions. There is a difference between *automation* (determinate reasoning embedded in the system, of the form ‘if -> then’) and *autonomous* aspects (which incorporate probabilistic algorithms which make determinations and guesses). Pilots in fighter jets can, and also will, change the *automation*-related settings. In commercial aviation this does not really happen, for instance for reasons of legal liability. Pilots do not tinker with the *autonomous* aspects.

Panel 4 | With All Deliberate Speed (Security)

Reported by Nirvan Tyagi, DLI Doctoral Fellow

Fast, Slow, and In-Between: Speed & its Discontents

Steven Jackson, Cornell University

Prof. Jackson began his talk by describing the need to separate the notion of speed from "speed stories" (stories about speedy technologies). Typically speed stories, as they are currently told, are used to mark distinctions, (speedy vs unspeedy) becomes (fast, slow) or (modern, traditional) or (smart, dumb), etc. He raised three high level points as problems with how speed stories are currently being told:

- 1) Neglect multiplicity of times
- 2) Neglect system failures
- 3) Neglect stories of necessary human effort

As an example of neglect of multiplicity of times, Jackson described the construction of a cell phone and the life of the constructing materials. While the final product is a speedy cell phone, the chemical elements (gold, silicon, etc.) that make up the phone were extracted over long periods of time (decades), and formed over even longer periods of time (millenia).

To describe neglect of system failures, Jackson gave the example of high frequency trading (HFT) platforms. While on the surface, these platforms are very speedy, there are edge cases where the system breaks down entirely. Microwave links are used to connect Wall Street to servers in New Jersey, however high heat/humidity slows down these links such an extent that it has a substantial impact on financial markets. In this way, speed can create systems with little tolerance.

To explain the final proposition, neglect of stories of human effort, Jackson gave two examples. The first is how the speed of a taxi system overshadows the enormous amount of time taxi drivers spend idle, waiting for passengers. In the second example, he describes tedious manual work that goes into cell phone repair. The high level idea is that slow, human processes are essential to maintaining speedy systems.

As a final note, Jackson spoke about how different groups of people have different relationships with speed. People in power, people on the receiving end of technology use, and imprisoned people are each differently impacted by speed. How then should we tell speed stories? The first step is to realize "speed" does not always equal "fast"; to begin to describe the limitations and human effort behind speed. Second, we must realize the "tragedy of the slowest": the differential impact of speed on different groups.

Flow Control: Internet Traffic Management & the Uneven Distribution of Delay

Fenwick McKelvey, Concordia University

Prof. Fenwick McKelvey spoke about the power of transmission. He presented three different short video clips to illustrate this power.

The first video, called "Allot Communications - Optimize, Monetize, Personalize", described a commercial bandwidth proposal by Allot Communications. In this proposal, Allot identifies the problem that some users have unusually high bandwidth applications (e.g., P2P and video streaming) which clog up the pipeline for other users. As a remedy, they suggest different personalized plans for users with different internet usage habits or needs, with different pricing policies for each.

Questions about how to manage delay in internet applications started at the very beginning: the proposal for Arpanet in 1969 specified 1/2 second delay. Extending these kinds of specifications and minimum service requirements to today's internet is complex. What should be the "optimal" distribution of delay? Different distributions of delay can be seen as unfair, and have sparked the recent "net neutrality" debates. For example, Comcast was discovered to be slowing down peer-to-peer (P2P) applications in 2007.

The second video, titled "Rogers Internet Turboboost", depicted an internet user becoming increasingly frustrated with the slow loading times of his browser. After switching to Rogers' Turboboost internet service, his internet connection suddenly became very fast, loading quickly and leading to him being very happy and getting the girl. McKelvey made the point that slow loading speeds (high delay) can be frustrating and can create desirable and undesirable forms of communication. For example, it's been shown that the government of Kazakhstan nudges internet users away from rival political parties by throttling internet speeds. In this manner, speed and delay can be culturally constructed.

The final video was titled "Living with Lag". It depicted a person wearing a visual and auditory mask that delayed real-world inputs delayed by some amount of time (say 1 second), so the person sees and hears things 1 second after they occur in real time. It was a humorous video showing the person attempting to play table tennis and partake in a group dance class. McKelvey noted that in order for communication to occur, a common consensus of temporality is needed. Uneven delay distributions work against this commonality.

Desirable Inefficiency

Paul Ohm & Jonathan Frankle, Georgetown University Law Center

Prof. Paul Ohm (professor of law) and Jonathan Frankle (2nd year PhD in computer science) presented together on the idea of desirable inefficiency. They observed a common thread across

some high technology systems: they purposefully introduce inefficiency into their design. This seems somewhat counterintuitive to the normal direction of technological advancement. They give three examples of technologies where desirable inefficiency has been introduced.

First, Bitcoin's Proof-of-work algorithm is designed intentionally to require millions of useless hashes before one is found that results in Bitcoin being created. These computations are not useful in maintaining the system itself, but instead are design choices that expend energy and computational power to ensure that the price of Bitcoin is not unnecessarily inflated.

The second example is iOS's mandated delay after incorrect passcode attempts. If a user enters a passcode incorrectly, iOS requires a short delay before the next attempt. Subsequent incorrect attempts increase the amount of delay before the next attempt is allowed. As a response, Ohm and Frankle pose the following question: How do you reason about choosing delay times?

Ohm and Frankle's last example is the "Magic shoebox" at the IEX stock exchange. This box, located at IEX's trading center in New Jersey, contains 38 miles of coiled fiber-optic cable, through which all new trades flow. The long physical link through which computations are run delays them by 315 microseconds. The delay controls the speed advantage that is available to high frequency traders.

Ohm and Frankle suggest that these techniques are used as a way to inject human values back into computer systems: they create "digital speed bumps". For example, the magic shoebox intentionally implements the delay by running computation through a long physical link, rather than use a software- implemented delay. This is an important design decision, as it binds the delay to a physical limitation rather than a software implementation that can be bypassed.

In each of the systems that Ohm and Frankle mentioned, inefficiency adds human values back to technological systems. In the case of Bitcoin, the value injected is trust – extra computation ensures that mined bitcoins are legitimately owned. In the case of iOS's passcodes, legitimacy is returned – delays allow the phone to ensure that the entity accessing the phone is its owner and not an automated system. Finally, IEX's magic shoebox re-introduces fairness to the trading system, ensure that high-frequency traders cannot unfairly profit from delays.

The talk ended with Frankle reporting on how the researchers' work can be connected to machine learning. Frankle advocated for using a similar methodology: looking for patterns that exist in how the technology is currently being used. Instead of working from idealistic definitions of fairness, it may be helpful to observe what ad-hoc techniques machine learning practitioners are currently putting into place.

Panel Discussion

Moderated by Laura Forlano, DLI Visiting Fellow

The following is a brief transcript of the questions and discussion following the presentations. Questions and answers are paraphrased.

SPEED CONFERENCE REPORT

Question: What is the relation between desirable inefficiency to economic idea of inefficiency?

Ohm: Economists are not good at tech policy. Economists do not think about fairness in robust terms.

Frankle: Decentralization. Systems impose some cost and do not make an effort to recover that cost.

Question: What is the role of stories of slowness? Meditation, slow food, and hand-made crafts are all popular. Do they act as a counterweight?

Jackson: This is, in part, a cultural/aesthetic response. I would prefer to focus on other non-aesthetic experiences like cell phone repair.

Ohm: There are differences in authenticity

Jackson: Not a huge fan of the term "authenticity". They are different experiences, wouldn't want to weigh one as more valuable than another.

McKelvey: There are ways to push back against selfish notion of speed.

Question: What is the role for inefficiency in process/lifecycle of innovation/creation/deployment of systems? And not just inefficiency of system, by itself.

Ohm: We sometimes think of innovation as a force of nature that tramples forward, and then societal norms work to catch up. I am in favor of heavy intervention.

Frankle: This was not something we thought about. The systems we looked at were brought to market very quickly, and part of the reason the system succeeded.

McKelvey: It might be possible to govern speed properly, but how to govern the governor of speed?

Question: What about governance of speed? Are regulators open to hearing about these ideas?

McKelvey: Excited Canada now has a latency standard. Regulators are aware, but not much traction. Shift towards quality of service/ minimum service guarantees, which is good.

Question: One way I've seen many people think about these systems is as: AI is doing its thing, ethics is separate and must intervene. Not a good way to think about it. I like how Jonathan describes a different approach: finding patterns that already exist.

Frankle: I have a negative reaction to fairness being thought of as a way to rehabilitate ML systems or fix at the end, instead of as a part of the process from the beginning.

Question: Considering the situation where some artifact of the system is desirable, but not realized until after it is removed. Is "speed" uncovering some of these desirable properties?

Ohm: Absolutely.

Frankle: Agreed, example: Aggregating information over multiple courthouses is a privacy violation without paying the cost of physically acquiring it.

Panel 5 | Labor and Manufacturing

Reported by Dr. Elizabeth O'Neill, DLI Visiting Fellow

The Labor and Manufacturing session featured two speakers and was moderated by Solon Barocas, Assistant Professor in Information Science at Cornell University. The first speaker, Argyri Panezi, a postdoctoral research fellow in political science at the Stanford Center on Philanthropy and Civil Society, spoke on “Fast Labor and Faster Machines in the Gig Economy.” The second speaker, Rory Van Loo, Associate Professor of Law at Boston University School of Law, spoke on “The Problem with Digital Market Perfection.”

Fast Labor and Faster Machines in the Gig Economy

Argyri Panezi, Stanford Center on Philanthropy and Civil Society

To address the topic of the gig economy, speed, and labor law, Panezi began by discussing two books: *Humans as a Service* by Jeremias Prassl, and *Re-Engineering Humanity* by Brett Frischmann. She highlighted some lessons from these books, including Frischmann’s point that in addition to humans engineering technology, technology engineers us back. She argued that there needs to be more communication between labor law scholarship and other scholarship on the relation between humans and technology.

Panezi then turned to the topic of the reality that lies behind the automated, easy way in which we experience the services supplied by the gig economy. She showed a clip from *The Prestige*, in which a magician appears to crush a bird, only to make it reappear. In the clip, a small child is upset by the apparent death of the bird, and shouts that the magician has killed the bird. It turns out that for this magic trick there are two birds, and the one that appeared to be crushed in fact was. The bird is a casualty of the magic. Panezi said she sees herself as the child in this scenario, shouting that the magician has killed the bird. For instance, behind the absence of gory images on social media platforms is outsourcing—workers in precarious jobs identifying pictures as disturbing so that they never appear on the users’ screens. Panezi asks: What do we sacrifice in terms of our labor force for the magic to happen?

Panezi stressed that she is not arguing that we are in an extreme dystopia, but she does think there is a trick behind the gig economy. Among other things, the gig economy has used the language of entrepreneurship in a clever way. Amazon’s Mechanical Turk plays with this trick—the original mechanical Turk was an apparent machine that was in fact operated by a human; Amazon uses this term for its own human-run system, playing with the irony.

Panezi argued that labor law regulators need to be careful when people use disruption and innovation in a way that evades labor law. Similar arguments have already been made about tax law. We should also worry about doublespeak, e.g. in the case of whether a person is an entrepreneur or is a worker as they were before the algorithmic version of the same job. We need to regulate invisible

labor as labor. The issues associated with the gig economy may not only be a matter of labor but also of dignity. These jobs may take away some of our humanity. For instance, algorithmic control plays an important role in some gig economy jobs. We also need to consider where the burden of entrepreneurship and speed falls—the social cost. Currently these costs are borne by workers, but it doesn't need to be that way; the costs could be borne by the platforms.

In the movie *Prestige*, part of the message is that the audience wants to be fooled—so they don't look too hard at the bird being killed. There may be parallels in the gig economy, Panezi suggested.

Interlude: I Love Lucy

[“Chocolate Factory” clip from I Love Lucy, 1952](#)

The Problem with Digital Market Perfection

Rory Van Loo, Boston University School of Law

Van Loo sketched a picture of forty years in the future, when people will look back on our current online shopping practices as archaic. The search for the best deal, sorting through Amazon results, navigating complex pricing schemes—eventually, Van Loo argued, we will have *digital butlers* who can do all of these tasks for us. They may even execute transactions for us. In the future, the practice of finding the best product on one's own will become like writing out directions by hand or printing a map today—infrequently done. Many major tech companies are aiming for this kind of outcome.

On the upside, digital butlers could save time and money for those who try to find the best deals and save money for those who currently don't spend their time searching for the best deals. Even sophisticated consumers overpay because they can't find the best option. Other potential benefits include the promotion of innovation and benefits to economy.

When considering downsides, Van Loo highlighted that it is important to note that the digital butler futuristic scenario doesn't fit into the usual paradigms we use to think about regulation. One can imagine a number of new problems that could result from a scenario in which as many people are using butlers as now use flight search engines. What would be the effects of a large group of people getting the advice to switch away from a product to one with a lower price? The competitor may respond by lowering its price, but if it can't, one can imagine a market response in which investors anticipate a company losing its customers and they respond by selling. The result could be market turbulence. Widespread collapse of companies outside the financial sector could have repercussions for financial sector. There is an analogy to runs on the bank; we now have the FDIC to address that kind of issue. But there is nothing comparable in place to address this problem associated with movement of customers. Van Loo described several possible responses we might see, like banks

hoarding cash. Companies could have a reason for conglomerization, due to the risk of one product failing. There could be consequent problems for employment and other social costs.

What are the potential social implications of this world? The law will largely determine the costs and benefits. Van Loo proposed that what we need is a radar gun for markets—to gauge their speed. This is not something people have worried about before because of the slowness of markets. People are slow to switch between products, overcome their habits, etc. If we develop a mechanism for measuring the speeds of markets, we would also want to know if markets are going too slowly. This could indicate a problem—e.g. perhaps contracts are locking people in, or termination fees are a problem; slow switching could be evidence of collusion, inadequate consumer access to data, etc.

Panel Discussion

Moderated by Solon Barocas, Professor, Cornell University

In the Q&A there was discussion of what the market positioning might be for digital butlers: whether they would be most likely provided by the current big players like Amazon and Google; whether such companies would be serving the interests of users. Some doubt was raised about whether the digital butlers would actually serve the interests of consumers. Van Loo suggested that the model for digital butlers might resemble that of Google Maps or search engines; the provider of a digital butler would not charge the customer but it would need to earn their trust. Questions were raised about how merchants and providers might respond to digital butlers, and also whether there might be potential for price discrimination and honeypot pricing, as there is now in financial markets. Van Loo agreed that there would be tremendous effort to undermine digital butlers; he drew an analogy to the resistance that airlines put up against websites like Expedia; in this case the intermediary eventually got the advantage over the companies.

In response to a question about what is distinctive about information markets, Van Loo proposed that in contrast to goods and services, information markets might have some potential to withstand something like a digital butler. For instance, it is hard to imagine the butler advising that one should leave Facebook. On the other hand, perhaps a digital butler could sort through all the factors that determine the worth of a search engineer and tell you which is best. In response to the same question, Panezi noted that every sector affects the labor sector, and she highlighted the role of labor in the information sector; there are now jobs devoted to reducing the harm that information (images) can cause on social media.

Asked about what could be done to make the metaphorical bird behind the magic—the sacrifice behind the gig economy—more visible, Panezi observed that we have laws in place to ensure that people are not enslaved—we have labor laws—and so it is noteworthy that much of the language behind the gig economy was sending the message that *this is not work*. Thus, there is a potential strategy in recognizing gig economy work as work and talking about it as work.

Panel 6 | Finance

Reported by Sabiha Sadeque, Purdue University

Man vs. Machine: The Past and Future of Human Collaboration with Machines in Financial Markets

Simona Abis, Columbia University

Simona Abis, a professor of business at Columbia University who takes a functional perspective on the past and future of human collaboration with machines in financial markets. In her research she depicts the classical trade-offs and bottlenecks of using technology – including AI and machine learning – for investment management. She also studies changes to investment management techniques brought about by new technologies.

Abis’s research highlights how natural language processing and big data can be used to answer difficult questions about asset management, and suggest regulatory best practices. Her present work focuses on using natural language processing to assign US investment funds into two different types. The first type, quantitative funds, rely on computer models to pick investments. The second, discretionary funds, are under the direction of skilled human investment managers.

Quantitative and discretionary funds differ in terms of several different characteristics. Abis’s research supports these findings using empirical evidence.

- Information processing capacity: Discretionary funds tend to hold stocks for which less information is available, while quantitative funds hold stocks with more information available, as such information is more easily processed.
- Assets held: Discretionary funds have more dispersed holdings than quantitative funds, but quantitative funds tend to hold more stocks overall.
- Performance: quantitative funds have better stock-picking ability, but discretionary funds perform better during recessions. In addition, because quantitative funds tend to pick the same stocks, they are prone to “overcrowding”. This causes declines in performance.

These findings show that while quantitative funds often are able to execute transactions at higher speeds, this does not always cause increases in performance. Because quantitative and discretionary funds learn about the market differently, they display different investment and risk management strategies. Abis also emphasizes that the number of quantitative investors operating at once has an impact on funds’ performance.

Abis concludes with suggestions for future research. She notes the increasing importance of big data in creating quantitative models – which may help these funds diversify, but could also worsen overcrowding. However, computational techniques such as deep learning and neural networks have potential to overcome the limitations of current quantitatively guided investors.

High Speed Trading Algorithms & Human Manipulations: Agency & Accountability in Complex Financial Markets

Katie Brennan, SUNY Polytechnic Institute

Kathleen P.J. Brennan is a political theorist interested in questions of non-human and distributed agency and how they impact regulations and regulatory environments. Her presentation, “High Speed Trading Algorithms and Human Manipulations: Agency and Accountability in Complex Financial Markets”, uses the example of high-frequency trading to explore questions of algorithmic agency. In her talk, Brennan focuses on the Flash Crash of May 2010, in which algorithmically driven fluctuations in a niche futures market caused a momentary drop in the Dow Jones that was of greater magnitude than the drop associated with the 2008 financial crisis.

Brennan begins by noting that human agency, which we tend to prize over other forms of agency, is already distributed, allowing for more careful consideration of non-human agency. She uses theories that promote a non-anthropocentric view of agency, including work on assemblages. Distributed agency, she says, provides a foundation for thinking about algorithmic agency in a concrete way.

There are several characteristics of HFT algorithms that suggest they can be considered as active agents, and therefore subject to accountability and regulation. Most notable is the ability of HFT algorithms to coalesce into crowds. They work together - often in “packs” or in ways that could be considered predatory – to impact financial markets. Though these ecologies work on timescales to minute or fast for humans to perceive, the collaborative operation of HFT algorithms can change events in ways that matter on human timescales, often in a way that seems to defy traditional linear causality. Also notable is the fact that HFT algorithms have become entities that produce data rather than just consume or analyze data. They transform data and react to macro-level events, suggesting they can both anticipate and impact oncoming events.

So, who or what was responsible for the Flash Crash? How can we avoid this incident in the future? Opinions differ. With regard to accountability, Western legal practice demands that responsibility is assigned after the crisis, rather than before. After the crash, an individual trader faced charges of market manipulation for the incident. But the market was already unstable, and the specific consequences of his actions are unclear enough that there’s not an easy regulatory answer to the question of accountability.

Why is this important? Brennan states that incidents that happen on the stock market impact everyone with a pension or retirement fund, not just people who actively choose to trade stocks. So a solution needs to be found to this issue of “packs of predatory HFT algorithms” that involves meaningful human participation and moves toward a better world for all participants.

Here, Brennan likens the actions of amalgamations of algorithms to “magic”. While not suggesting that their effects are wholly unknowable, she suggests that assemblages of algorithms, such as HFT algorithms, should be investigated for emergent behavior, non-human affect, or other unanticipated consequences of algorithms working as crowds.

Finally, Brennan concludes by noting that that issues of agency and accountability for HFT algorithms are also issues of regulation. Can regulation stop the “predatory” nature of these algorithms, or does their nature – to seek profit – make predation unavoidable? Should regulation place accountability on individual humans, individual companies, or larger groups or entities?

While humans have great impact on global financial market, it’s unclear that our agency is more impactful than that of algorithms. To Brennan, this means that accountability and regulation for financial markets should take assemblages and ecologies into account. These ecologies involve human traders and companies, HFT algorithms themselves, physical infrastructures, resource availability, and U.S. and global political and legal structures. Regulations can’t just focus on newest tech, they need to take a broader view of the context in which markets operate. Importantly, regulations for algorithms that operate on faster-than-human timescales would also need to operate on those timescales in order to be effective.

One takeaway from the Flash Crash was the introduction of the consolidated audit trail, allowing the SEC to track trades from their inception, through all traders and brokers. This isn’t yet functional, but could be a powerful tool to better understand the networks of agents involved in transactions, which could yield still better regulation.

Speed II: Have We Reached a Tipping Point?

Larry Tabb, TABB Group

The third speaker was Larry Tabb, founder of Tabb Group, who by his own description has been dealing with issues of speed in finance since 1980. He developed fintech systems for major companies, and has been an analyst looking at financial markets. His talk focuses on the current and future state of speed in equity markets. Tabb notes that some of his work contradicts points made by some of the other presenters, stating that worries about people taking advantage of speed for potentially nefarious reasons are unfounded.

Tabb begins by citing some figures on the importance of speed in the current market: 37% of orders for large cap stocks are executed within 1 second, 73% within 10 seconds. This means that there’s little time for people to protect themselves from bad trades. In addition, the U.S. stock market generates a huge amount of data – millions of messages per second. Execution time for trades has dropped to fractions of seconds, as the market becomes more efficient, investors are keeping more of the money.

In terms of location, major U.S. based players are largely located in New Jersey, with the exchange in Chicago also playing a large role. The NYC-Chicago connection is vital, and the race to shave seconds, and then fractions of seconds, off of this connection has been a major competition (and the subject of a book). Firms have fought over the advantages that this connection provides, eventually moving the link from cable to microwave to a microwave-laser combination.

Tabb notes several implications of speed in financial markets. First, speed is vital for profit: market makers need to make sure they aren't taken advantage of, so they post transactions as fast as possible. But in "high rebate" markets, fees for transactions are larger, so posting them quickly ensures traders the best price. Speed bump regulation comes in here: in some exchanges, such as IEX and NYSE, intentional delays are introduced to obfuscate prices and create a fairer market.

For financial transactions that involve large amounts of stock, selling is more difficult in faster markets. Here, small orders are more efficient, favoring smaller buyers. For larger buyers (like mutual funds or larger asset managers), selling large quantities of stock, speed is problematic. These companies push to slow the market down – but this makes it more expensive for others to trade.

Both exchanges and regulatory bodies have proposed "speed bumps" to keep markets fair. Mechanisms like different delays for buying vs. selling, auction mechanisms, and long-life orders have all been proposed. Here, Tabb points out, we have moved from a market that is unintentionally slow, to a market that is intentionally fast. With these new regulations, we are moving toward a market that is intentionally slow.

Here is where Tabb comes to the idea of the "tipping point" of speed. He notes that high-speed trading is approaching the physical limits of speed. Information is traveling at almost at the speed of light, so traders are less able to create profit. The new focus must be about intelligence or "meta speed": measurements about the quality of speed. How fast is the market? What's the traffic? Is it at capacity? Is there jitter? What other speed conditions are there?

As markets become more speed diverse, firms hoping to profit in these markets have to identify themselves by characteristics other than speed, notably their awareness of the surrounding information of each market.

Speed and Reliability: Cryptocurrency's Irreconcilable Opposites?

Michael O'Connor, Pennsylvania State University

Dr. Michael O'Connor began the last presentation of the conference by giving a background in internet currency and blockchain tech. In sum: The two fundamental needs for internet currency are the need to prove that you own your cash, and need to prove that you have as much cash as you say you do. The first problem is solved with public key cryptography, the second is more difficult. A unified ledger can keep track of transactions, but these are a single point of failure and are vulnerable to many different types of attacks. Early attempts at single ledger internet currency were shut down for allowing bad actors to conduct transactions, or failed because of lack of user trust.

Another option is to use multiple ledgers, on the assumption that if one has a bad record, another would be better. Disagreements are solved, and untrustworthy intermediaries deterred, using Satoshi Nakamoto's solution to the Byzantine Generals Problem: Transactions are assigned to blocks and hashed repeatedly, attempting to get particular result. If correct answer found, the transaction

is added to the block and transmitted to the network. This becomes the new end of the blockchain and is linked to the previous block. Thus, neither currency nor transactions can be faked. Users have incentive to spend electricity and computing power on this process because they receive currency for participating.

O'Connor notes that there are a few paradoxes inherent in this process. An example is that it would seem that having more computers participating in this process would make the network more redundant, and therefore more secure. However, the fact that there are hundreds of cryptocurrencies being traded means that users with more mining power can use that power to execute attacks on smaller, thinly traded currencies. Once mining power is used to create currency, it is then sold for cash and the computing power used to rewrite the ledger, allowing the attacker to sell the currency again. These attacks are surprisingly affordable.

How can these attacks be stopped? Can blockchain tech be scaled, stay secure, and not kill the environment? O'Connor proposes that regulation is the answer, however governments and corporations need to work together on regulation. Private action to in cryptocurrency crises has already been effective, as in the case of Coinbase restoring consumer accounts from its own treasury after a crash. The combination of government and corporate regulation will be needed to ensure trust in cryptocurrency.

Other ways to accelerate transactions while maintaining reliability and reducing environmental impact involve moving or reducing the proof of work burden. There are three main proposals for achieving this. First, currency systems could allow "off-chain" or "side chain" transactions which are later settled with the main chain, like a credit card. Second, the proof of work burden could be reduced or eliminated. This is difficult, because proof of work is the only thing that conclusively eliminates untrustworthy actors. However, proposals such as proof of stake, in which participants offer some of their own currency as collateral when completing transactions and lose the collateral if transactions are false, can make malicious actions too costly to be practical. Third, the fundamental structure of the blockchain could be changed. Some blockchain projects need fewer nodes to approve a transaction, or use "sharding" to break transaction verification across the network. Nodes don't know in advance what transactions they will be assigned, so they can't game the system.

O'Connor concludes with a reminder that, in its current state, the blockchain system cannot sustain the growth of cryptocurrency. He foresees a few possible futures for cryptocurrency as it currently exists: either governments will demand change and punish projects that allow bad actors or flout regulations, users who lose money on cryptocurrency will sue, or markets will just fail for lack of trust.

Panel Discussion

Moderated by James Grimmelmann, Professor, Cornell Tech

Question: Regulation seems to be a theme here. What if regulators are automated systems?

Answer: Automated/autonomous regulation might not work here. Human regulators would need to take the time to understand the tech they're regulating, and could implement things themselves to better understand what's happening (although people with this type of skill would get swallowed up by the industry and not be regulators anymore). AI are not independent regulators yet. Algorithms will find patterns independently, so regulators or practitioners are necessary to stop errors.

Also, technologies like the consolidated audit trail will make fintech easier to secure. It increases transparency and understandability, and the ability to understand the underlying players across the industry will help the industry regulate itself without outside action. In some cases, automated action might not be feasible, as in cases where FTC access to underlying fintech algorithms was denied.

Question from Helen: Simona and Larry say we've hit a wall with speed, so a "Speed 2" is needed. What's next? People can integrate surprising datasets to find regularities. Does Speed 2 need better datasets?

Answer: There are lots of interesting datasets being analyzed. New tech like autonomous cars will increase this. Any dataset with value is being analyzed for economic value, everything is being automatically investigated – everything is data mined. Large companies are all trying to find data sources nobody else has found. Everyone wants to be at the cutting edge of data/analysis.

Question: We think of regulators as being behind, but is always them that's behind or is the tech industry understanding that there are advantages to working within regulatory frameworks? Do people building tech approach their own products differently as their products mature?

Answer: There's a lot of enthusiasm for tech that's not well founded. There's not a lot of blockchain stuff in finance that's worked. Regulators need to work closely with industry to make sure blockchain works. What about putting equities on blockchain, when they're tied to specific companies? Might be more efficient on traditional models rather than on blockchain. On the institutional side, companies are really investing in new tech for improving markets. Data is low-hanging fruit, what's not done so much is integrating new tech into the core of their trading systems.

Question. The panel is commended for diversity. When it comes to solving problems with tech, is it a matter of finding data, or having diverse people? What helps discern new connections within data?

Answer: The panelists acknowledge that diversity is important everywhere. We get examples like racist facial recognition that people don't notice for a long time. These come in regardless of the kind of data you're talking about. There's a bias that if data is big enough it will overcome its own problems. This is demonstrably not true; we need diversity at every step. Another panelist notes that humans are still needed in order to create value using algorithms. We need "translators" to relate between finance, data, context, tech, and applications. This is where humans can help.

Appendices

Speaker Biographies

Simona Abis

Professor Simona Abis joined Columbia Business School in 2017. She holds a PhD from INSEAD. Before joining the PhD program Simona worked as a quantitative researcher for a systematic hedge fund. Her research interests span the fields of information economics, empirical and theoretical asset pricing, financial econometrics, microeconomics, Bayesian learning, machine learning, mutual funds and hedge funds. Overall Simona is interested in the impact of technology on financial markets. Her current research focuses particularly on the impact of technological change on investment management through the rise of quantitative investment.

Mike Ananny

Mike Ananny is an Assistant Professor at USC's Annenberg School for Communication and Journalism where he researches the public significance of systems for networked journalism. He studies how institutional, social, technological, and normative forces both shape and reflect the design of the online press and a public right to hear. He is also an Affiliated Faculty with USC's Science, Technology and Society research cluster, and a past Faculty Associate with Harvard's Berkman Center for Internet & Society. He holds a PhD from Stanford University's Department of Communication. Ananny was a founding member of the research staff at Media Lab Europe, a founding member of Expresto Software Corp, and a Postdoctoral Researcher at Microsoft Research's Social Media Collective.

Peter Asaro

Peter Asaro is a philosopher of science, technology and media. His work examines artificial intelligence and robotics as a form of digital media, the ethical dimensions of algorithms and data, and the ways in which technology mediates social relations and shapes our experience of the world. His current research focuses on the social, cultural, political, legal and ethical dimensions of robotics and artificial intelligence, from a perspective that combines media theory with science and technology studies. Asaro's research also examines agency and autonomy, liability and punishment, and privacy and surveillance as it applies to consumer robots, industrial automation, smart buildings, and autonomous vehicles. He is currently an Associate Professor in the School of Media Studies at The New School.

Kathleen P.J. Brennan

Dr. Brennan received her PhD from the Department of Political Science at the University of Hawai'i at Mānoa. She is now a Postdoctoral Research Associate in Interdisciplinary Studies at SUNY Polytechnic Institute. Her research stems from the intersection of political theory, media studies, popular culture, and international relations, with a focus on the ways in which the interactions of online and offline worlds impact events and individuals. Her dissertation examined online and offline worlds as horizontal spaces separated by porous borders. Both her doctoral and postdoctoral work focuses on digital nonhuman actors like memes, hashtags, and algorithms.

Mary "Missy" Cummings

Mary "Missy" Cummings received her B.S. in Mathematics from the US Naval Academy in 1988, her M.S. in Space Systems Engineering from the Naval Postgraduate School in 1994, and her Ph.D. in Systems Engineering from the University of Virginia in 2004. A naval officer and military pilot from 1988-1999, she was one of the Navy's first female fighter pilots. Cummings is currently a Professor in the Duke University Pratt School of Engineering, the Duke Institute of Brain Sciences, and is the director of the Humans and Autonomy Laboratory and Duke Robotics. Her research interests include human-unmanned vehicle interaction, human-autonomous system collaboration, human-systems engineering, public policy implications of unmanned vehicles, and the ethical and social impact of technology.

Jason Farman

Jason Farman is the Director of the Design Cultures & Creativity Program, an Associate Professor in the Department of American Studies, and a faculty member with the Human-Computer Interaction Lab at the University of Maryland, College Park. He is author of the book *Mobile Interface Theory: Embodied Space and Locative Media* (Routledge, 2012 — winner of the 2012 Book of the Year Award from the Association of Internet Researchers), which focuses on how the worldwide adoption of mobile technologies is causing a reexamination of the core ideas about what it means to live our everyday lives: the practice of embodied space. He has published scholarly articles on such topics as mobile technologies, digital maps and cultural geography, locative and site-specific art, videogames, digital storytelling, performance art, social media, and surveillance.

Jonathan Frankle

Jonathan Frankle is a second-year PhD student in computer science at MIT, where he researches topics in artificial intelligence, applied cryptography, and technology policy. His current research focus involves exploring the basic properties of neural networks by developing techniques that make it possible to find small neural networks that train successfully. On the side, he also works on the effort to design and teach a programming course for lawyers at the Georgetown University Law Center and several other law schools.

James Grimmelmann

James Grimmelmann is a professor of law at Cornell Tech and Cornell Law School. He helps lawyers and technologists understand each other, applying ideas from computer science to problems in law and vice versa. He studies how laws regulating software affect freedom, wealth, and power. He writes about search engines, social networks, data havens, hackers, trolls, copyright-infringing robots, and magical 3D printers, among other things. He is the author of the casebook *Internet Law: Cases and Problems*, now in its fifth edition, and of over forty scholarly articles and essays. He holds a J.D. from Yale Law School and an A.B. in Computer Science from Harvard College.

Steven Jackson

Steve Jackson is an Associate Professor in the Department of Information Science and Department of Science and Technology Studies at Cornell University, with additional graduate field appointments in Communication and Public Affairs. He holds a Ph.D. in Communication and Science Studies from the University of California, San Diego. His research connects contemporary questions in information science to theoretical and methodological traditions in the critical, interpretive, and historical social sciences. He is especially interested in places where new computing forms and practices meet the material world, with implications for sustainability, development, and inequality.

Wendy Ju

Wendy Ju is an Assistant Professor at the Jacobs Technion-Cornell Institute at Cornell Tech in the Information Science program. Dr. Ju comes to Cornell Tech from the Center for Design Research at Stanford University, where she was Executive Director of Interaction Design Research, and from the California College of the Arts, where she was an Associate Professor of Interaction Design in the Design MFA program. Her work in the areas of human-robot interaction and automated vehicle interfaces highlights the ways that interactive devices can communicate and engage people without interrupting or intruding. Dr. Ju has innovated numerous methods for early-stage prototyping of automated systems to understand how people will respond to systems before the systems are built. She has a PhD in Mechanical Engineering from Stanford, and a Master's in Media Arts and Sciences from MIT. Her monograph on *The Design of Implicit Interactions* was published in 2015.

Kate Klonick

Kate Klonick is an Assistant Professor at Law at St. John's University Law School and an Affiliate Fellow at the [Information Society Project](#) at Yale Law School and New America. She holds a JD from Georgetown University Law Center, where she was a Senior Editor at *The Georgetown Law Journal* and the Founding Editor of the [The Georgetown Law Journal Online](#); and a PhD from Yale Law School where she studied under Jack Balkin, Tom Tyler, and Josh Knobe. Klonick applies principles of cognitive psychology to the study of emerging issues in law and technology. She focuses on the Internet's effect on freedom of expression and private platform governance, as well as online shaming, artificial intelligence, robotics, content moderation, algorithms, and privacy.

Andrea Matwyshyn

Professor Andrea Matwyshyn is an academic and author whose work focuses on technology and innovation policy, particularly information security, consumer privacy, intellectual property and technology workforce pipeline policy. She is a professor of law at Northeastern University and co-director of the school's Center for Law, Innovation, and Creativity. She has served as the senior policy advisor at the US Federal Trade Commission, worked as a corporate attorney in private practice, and been a faculty member at the Wharton School, Northwestern University School of Law, and the University of Florida.

Fenwick McKelvey

Fenwick McKelvey is an Associate Professor in Information and Communication Technology Policy in the Department of Communication Studies at Concordia University. He studies the machines, bots, artificial intelligence, algorithms, and daemons that make up the Internet's infrastructure. His recent studies have focused on the daemons that manage Internet flows and their role in Net Neutrality debates, the new software and social media platforms that mediate political engagement and the algorithms and AIs that govern the discoverability of online content.

Helen Nissenbaum

Helen Nissenbaum is Professor of Information Science at Cornell Tech. Her research takes an ethical perspective on policy, law, science, and engineering relating to information technology, computing, digital media and data science. Topics have included privacy, trust, accountability, security, and values in technology design. Her books include [Obfuscation: A User's Guide for Privacy and Protest](#), with Finn Brunton (MIT Press, 2015) and [Privacy in Context: Technology, Policy, and the Integrity of Social Life](#) (Stanford, 2010). Recipient of the 2014 [Barwise Prize](#) of the American Philosophical Association, Nissenbaum has contributed to privacy-enhancing software, including [TrackMeNot](#) and [AdNauseam](#). Nissenbaum holds a Ph.D. in philosophy from Stanford University and a B.A. (Hons) in philosophy and mathematics from the [University of the Witwatersrand](#), South Africa.

Michael O'Connor

Michael O'Connor is a Visiting Assistant Professor of Law at Pennsylvania State University. He teaches in the areas of cyber law, including data security and privacy, cybercrime, and emerging technologies. His scholarship focuses on cryptocurrency and blockchain technology, including securities regulation, money laundering, and other topics. In previous private practice, O'Connor advised clients on legal planning for data breaches, obligations for safe handling of identifiable information, and the legal implications of blockchain and cryptocurrencies.

Paul Ohm

Paul Ohm is a Professor of Law at the Georgetown University Law Center. He specializes in information privacy, computer crime law, intellectual property, and criminal procedure. He teaches courses in all of these topics and more and he serves as a faculty director for the Center on Privacy and Technology at Georgetown. In his work, Professor Ohm tries to build new interdisciplinary bridges between law and computer science. Much of his scholarship focuses on how evolving technology disrupts individual privacy. Ohm blogs at Freedom to Tinker.

Argyri Panezi

Argyri Panezi is a Postdoctoral Fellow at the Digital Civil Society Lab. Argyri's work explores the effects that disruptive technologies have on citizens, on institutions, and the law. In the Digital Civil Society Lab, she is examining digital civil society interactions with libraries and other cultural heritage institutions, and the relevant legal frameworks incentivizing distributed methods for building a content infrastructure accessible online. She specializes in Internet law and policy, intellectual property law, with an emphasis on digital copyright, as well as data protection, intellectual goods management, and law and economics. Ultimately her work aspires to illuminate the effects that algorithmic speed and AI might have on domains such as labor and consider policy responses to scenarios of technology dystopia.

Frank Pasquale

Frank Pasquale is a Professor of Law at the University of Maryland Francis King Carey School of Law, and author of *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press, 2015). Pasquale has researched and taught in law and policy fields ranging from health care finance to privacy and intellectual property. His widely cited work has influenced both policymakers and academics. His research agenda focuses on challenges posed to information law by rapidly changing technology. Frank has been a Visiting Fellow at Princeton's Center for Information Technology, and a Visiting Professor at Yale Law School and Cardozo Law School.

Larry Tabb

Larry Tabb is the founder and research chairman of TABB Group, the research and strategic advisory firm focused exclusively on capital markets. Launched in 2003 and based on the interview-based research methodology of "first-person knowledge" that Larry developed, TABB analyzes and quantifies the investing value chain, from the fiduciary, investment manager and broker, to the exchange and custodian, helping senior business leaders gain a truer understanding of financial markets issues. Before founding TABB Group and serving as CEO through 2016, Larry was vice president of TowerGroup's Securities & Investments practice, where he managed research across the capital markets, investment management, retail brokerage and wealth management segments.

Rory Van Loo

Rory Van Loo is a second-year law professor whose research focuses on consumer transactions, with a particular interest in the intersection between technology and regulation. In each of his first two years of eligibility, he was selected through blind peer review for the Yale/Stanford/Harvard Junior Faculty Forum. His articles have been published or are forthcoming in the *University of Pennsylvania Law Review*, *Duke Law Journal*, *UCLA Law Review*, *Michigan Law Review*, and *Columbia Law Review*. Prior to joining BU Law, Professor Van Loo served on the implementation team that set up the Consumer Financial Protection Bureau, helping to build the framework for supervision of large banks. He also lectured at Harvard Law School and conducted empirical studies at McKinsey & Co. for multinational consumer companies in mergers and acquisitions, organizational design, and sales.



SPEED
CONFERENCE

DLi@ CORNELL
TECH

Conference Flyer

Speed Conference / Sep 28-29, 2018 / Cornell Tech / New York City

Speed is an overlooked aspect of algorithmic oversight: Intelligent systems and human actors operate at vastly different speeds. These differences present challenges for managing and responding to algorithmic decisions. As critical societal functions are handed off to AI control, the speed differential raises fundamental questions about the role and relevance of human oversight.

Isolated discussions about speed have occurred in areas such as robotics, finance, warfare, and online communication. Speed Conference brings together speakers from these and other fields to connect isolated efforts and to identify common themes and potential solutions in areas of speed, AI, and algorithmic oversight.

Hosted by the **Digital Life Initiative**, we welcome researchers, practitioners, thinkers, makers, and others interested in speed to join the discussion.

Confirmed speakers include: Simona Abis, Mike Ananny, Peter Asaro, Katie Brennan, Missy Cummings, Jason Farman, Jonathan Frankle, James Grimmelmann, Steven Jackson, Wendy Ju, Kate Klonick, Andrea Matwyshyn, Fenwick McKelvey, Helen Nissenbaum, Michael O'Connor, Paul Ohm, Argyri Panezi, Frank Pasquale, Larry Tabb, and Rory Van Loo.

Conference Chairs: James Grimmelmann and Helen Nissenbaum.

Speed Conference is generously supported by a gift from **Microsoft Corporation**.

CORNELL
TECH
HOME OF THE
JACOBS
INSTITUTE



Conference Registration: speedconference.eventbrite.com

Registration is free, but we ask that all attendees submit the registration form. For more information, including event contacts and the full conference program, kindly visit dli.tech.cornell.edu/speed