

*Editorial Introduction to the Special Articles on Context*

# Artificial Intelligence, Autonomy, and Human-Machine Teams: Interdependence, Context, and Explainable AI

*W. F. Lawless, Ranjeev Mittu, Donald Sofge, Laura Hiatt*

■ *Because in military situations, as well as for self-driving cars, information must be processed faster than humans can achieve, determination of context computationally, also known as situational assessment, is increasingly important. In this article, we introduce the topic of context, and we discuss what is known about the heretofore intractable research problem on the effects of interdependence, present in the best of human teams; we close by proposing that interdependence must be mastered mathematically to operate human-machine teams efficiently, to advance theory, and to make the machine actions directed by AI explainable to team members and society. The special topic articles in this issue and a subsequent issue of AI Magazine review ongoing mature research and operational programs that address context for human-machine teams.*

Context supposedly explains everything in the environment that influences our perceptions, cognitions, actions (Sober 2009), and awareness, the latter setting the stage for deception (Lawless 2017a.) For example, Wells Fargo, the largest US mortgage lender and the third largest US bank, “admitted to deceiving the US government into insuring thousands of risky mortgages” (Stempel 2016). Context is the word sequence in a sentence that allows the brain to discover a missing word, unravel the meaning of handwriting, or repair faulty grammar (McClelland and Rumelhart 1988). With Bayesian inference, Marwah and colleagues (2012) applied context-specific ontologies in cancer research. An organization’s context is its management, culture, and systems (Doolen et al. 2003).

## The Definition of Context

Merriam-Webster<sup>1</sup> offers two definitions of context: “the parts of a discourse that surround a word or passage and can throw light on its meaning; or, the interrelated conditions in which something exists or occurs; for example, setting the historical context of the war.”

In a review of the theories of meaning, Speaks (2017) states that different meanings arise from different situations: “These situations are typically called *contexts of utterance*, or just *contexts*, and expressions whose reference depends on the context are called *indexicals* or *context-dependent expressions*.”

The *Oxford English Dictionary*<sup>2</sup> locates the origin of context in Late Middle English, from Latin *weave* and *text*, meaning, “The circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood.”

In *AI Magazine*, it has been common to define context as situational awareness (as, for example, by Pfautz et al. [2015, p. 42]). Authors in *AI Magazine* have also sometimes used the term *context* for internal processing, as occurs in natural language understanding; for example, “translating text from one natural language to another ... [based on a constructed] internal representation” (Brill and Mooney 1997, p. 21). What we like about the definition used in *AI Magazine* is that the awareness of a context or situation may be nonverbal and not lend itself exclusively to a “mental object with semantic properties ... [to make] sense of each other’s behavior” (Pitt 2018). After all, for quantum mechanics, no consensus exists on its meaning even after nearly a century of debate, yet it is the most predictive of sciences (see, for example, the work of Weinberg [2017]).

When is context clear? Clear context might be a chaplain giving last rites, a prearranged visit with a doctor, or an official letter from the IRS demanding back taxes.

When is context unclear? An uncertain context might be the fog of war, a shout from an airline steward to brace for impact, or the legal rules of discovery preventing objective reconstruction for the context of a crime (Felson and Eckert 2015).

When is context an illusion? In 1944, based on an expedition that found evidence for Einstein’s theory of relativity, an editorial in the *New York Times* declared that the world was illusory. The physicist Carlo Rovelli (Garner 2016) wrote that “reality is not as it appears.” Humans are prey to Adelson’s (2000) checkerboard illusion, whereas photometers are not. Cybercriminals and spies rely on illusion.

A 1983 example of computers and context comes from the RAND Corporation. Irving (2018) writes that on September 26, 1983, in a bunker near Moscow, Lt. Col. Stanislav Petrov was monitoring Russian satellite data for signs of missile launches by the United States when a siren sounded. Petrov’s screen flashed “Missile launch” five times. Petrov didn’t know that a satellite had misinterpreted sunlight reflected on

clouds, but he phoned his duty officer and informed him of a false alarm. The duty officer passed the false-alarm message up the chain of command without asking Petrov for an explanation.

Context is determined by an awareness of social reality (Pfautz et al. 2015, p. 42), interdependently, that is, change the situation, players, or social norms, and context changes (Lawless et al. 2018). For example, context changes after a sports team substitutes a novice for its star player, after a deception is uncovered in court before a jury, or after a political party changes its leadership. The context of war changed early in 2018 with news that a swarm of drones attacked a previously impregnable Russian base in Syria, reportedly killing Russian soldiers (Grove 2018).<sup>3</sup>

Already, the disruptive economic impact of machine learning (ML), a subset of artificial intelligence, has been estimated in the trillions of dollars (Brynjolfsson and Mitchell 2017). Applications of ML and other AI algorithms are having an unprecedented impact across industry, the military, medicine, finance, and society generally. But as autonomous machines become important to society, so does context. Judea Pearl (2002) warned AI scientists that they must “build machines that make sense of what goes on in their environment,” a warning that unheeded could impede further development (Pearl and Mackenzie 2018).

For example, self-driving cars have been involved in at least three fatalities, including that of a pedestrian, yet humans, not ML, had to unravel the pedestrian’s context. After the pedestrian’s death in March 2018, an article in *Commentary* (Rothman 2018) blamed the pedestrian, an article in the *New York Times* (Wakabayashi 2018) blamed Uber’s self-driving autonomy software, and the preliminary decision by the National Transportation Safety Board (2018) blamed Uber, because the car saw the pedestrian 6 seconds ahead and selected emergency braking 1.3 seconds ahead, but the emergency brakes were not operational and could not improve the car’s performance. The vehicle operator engaged the steering wheel less than a second before impact and began braking 1 second after impact. For our purposes, the car performed appropriately, and it performed faster than the human operator, but it was not programmed to share what it knew, although it could have almost 5 seconds before the human became aware of the situation.

These seemingly unrelated problems regarding context require an interdisciplinary approach to solve, but, until now, social science has been slow to contribute to a theory of interdependence for human-machine teams (Lawless 2017a; Lawless 2017b). Although social science has contributed to the development of statistical analysis, it has struggled with the replication of experiments (Open Science Collaboration 2015). Relatedly, interdependence theory indicates that the information from behavior and mental concepts of behavior are orthogonal, producing poor correlations (for example, see the work of

Zell and Krizan [2014]). Putting aside that psychological tests can be financially lucrative, personality and other tests have been discredited as invalid, but they have not been discontinued — for example, the Myers–Briggs personality test, the most lucrative test in psychology (Emre 2018); the implicit association test, used to determine implicit racism (Blanton 2009), although racism is resistant to treatment (Jussim 2017); and self-esteem tests, used in schools for decades (Baumeister et al. 2005).

## The Future Determination of Shared Context in Real Time with Human-Machine Teams

The US Department of Defense is shifting to real-time operations, which makes critical the computational determination of context. Based on the RAND Corporation's story about Petrov in 1983, the biggest concern with AI is its use to determine the context of a nuclear confrontation when humans are not in the loop, a so-called Skynet situation (Lawless 2018). Of concern, in the recent analysis by RAND, is that AI systems may undermine the stability between nations and make catastrophic war more likely (see, for example, the work of Geist and Lohn [2018]).

China has demonstrated swarm intelligence algorithms that enable drones to hunt in packs. Russia has announced plans for an underwater drone that could guide itself across oceans to deliver a nuclear warhead powerful enough to vaporize a major city. Adding urgency to the determination of context in real time, China and Russia have announced the addition of hypersonic missiles to their military arsenals. Despite this urgency, “Americans seem generally complacent about the dominance of their armed forces ... creating a crisis of national security” (Edelman and Roughead 2018, p. vi).

As with Uber and its vehicle operator in the death of the pedestrian (National Transportation Safety Board 2018), where the machine worked and the human operator failed, what if in future situations contexts change more rapidly than humans can process, so that at some point AI systems alone must determine context? Woo (2018), for example, notes that quicker human-reflex-like responsiveness is thought to be likely with 5G.

How can we arrange human-machine teams to make the best possible decisions in real time, not only to protect national defense, to respond to medical emergencies, or to warn other cars while riding inebriated in a self-driving car, but also to accomplish these tasks more productively, efficiently, and safely than now? For example, can user interventions improve the learning of context for autonomous machines operating in unfamiliar environments or experiencing unanticipated and rapid events? Can autonomous machines be taught to explain contexts by reasoning, with inferences about causality, and with decisions to humans relying on comprehensible

explanations (Kambhampati 2018). And for mutual context, can AI machines interdependently affect human awareness, teams, and society, and how might these machines be affected in turn? In short, in real time, can situational awareness of context be mutually constructed, mutually shared, and mutually trusted among machines and humans and thus be productive, safe, efficient, and a benefit to society?

To address these questions, we need to know more about the effects of interdependence, which Jones (1998, p. 33) said characterized social interaction but was bewildering theoretically. Nonetheless, our knowledge about interdependence is growing (Lawless 2017a; Lawless 2017b). It not only determines context (Lawless et al. 2018), but it is a social state very sensitive to changes in context, exemplified by instability when two adversaries angrily express their two-sided stories, but once adversaries compromise, their context is determined. (See, for example, the bipartisan legislation passed overwhelmingly in response to the 2018 nuclear posture review [Mattis 2018; Payne 2018].) The universal motivation is for convergence to a single story (however, removing an alternative interpretation increases uncertainty and risk [Lukianoff and Haidt 2018]) and nonfactorability — for example, the struggle to write a successful screenplay that dramatizes a courtroom scene, to direct a winning political battle, or to describe protectively an engineering innovation in a patent.

According to the National Academy of Sciences (Cooke and Hilton 2015), teams are interdependent, and the best teams are highly interdependent (Cummings 2015). Interdependence is associated with innovation (Lawless 2017a; Lawless 2017b). However, to maintain a state of interdependence, a leader must train or quickly replace poorly performing team members (Hackman 2011), such as when Verizon removed the architect of its struggling online advertisement business (FitzGerald and Ramachandran 2018); keep a complex technology composed of numerous parts fully integrated, such as the US Army's recent successful missile defense system (Freedberg 2018); reduce uncertainty by sustaining an active competition among self-interested, two-sided perspectives not only to reach the best decisions — for example, the “informed assessment of competing interests”<sup>4</sup> — but also to reduce human error (Lawless et al. 2017); and, finally, to keep a team focused on collecting and analyzing the objective and statistical evidence that guides the search for vulnerabilities in a team and its opponents without becoming overly confident (Massey and Thaler 2005), such as when an overconfident CBS was defeated by Viacom (James 2018).

Theoretically (Lawless 2017b), it has been difficult to explain why information obtained from observing the performance of the best teams seldom generalizes (for example, even veteran movie studios with past successes can fail at the box office with a movie sequel [Fritz 2017]). One reason is that, mathematically, by reducing the degrees of freedom, the

structure of the best teams minimizes the entropy produced (Lawless 2017a; Lawless 2017b), allowing the best teams to increase their maximum entropy production, but becoming at the same time an impediment to determining how each member of a team contributes to a team's context or its performance, the loss of structural information accounting for one of the bewildering effects caused by interdependence.<sup>5</sup> Inverting the problem to put a finer point on what we know about human teams, in contrast to the dearth of information from well-run teams, dysfunctional teams produce too much information (for example, divorce; the costly CBS–Viacom merger standoff; and the US Department of Energy's mismanagement of military nuclear wastes indicated by remediation costs in the many tens or hundreds of billions of dollars [Lawless et al. 2014]).

We also need to explore what happens to shared context when machines begin to think (Gershenfeld 1999) or, like humans, develop subjective states that allow them to monitor and report on their interpretations of reality (Dehaene, Lau, and Kouider 2018), forcing scientists to rethink the general model of human social behavior (see, for example, the work of Lawless et al. [2018]). If dependence on AI and ML continues or grows, we and the public will also be interested in what happens to context shared by users, teams of humans and machines, or society when these machines malfunction (Kissinger 2018). As we “think through this change in human terms” (Shultz 2018), our ultimate goal is for AI to advance the performance of autonomous machines and teams of humans and machines for the betterment of society wherever these machines interact with humans or other machines.<sup>6</sup>

## Summary

If a computer program for computational context automatically knows the situation that improves performance, it may not matter whether context is real, uncertain, or illusory. This idea agrees with the Department of Defense need for automatically having a common perception of the surrounding world and being able to place it into context. Yet, hybrid systems able to share awareness among the members of an autonomous human-machine or machine-machine team is a computational challenge that will increase when “things begin to think” (Gershenfeld 1999). Instead of a challenge, however, the computational determination of shared context may offer an opportunity to advance the science of teams, where, for example, until recently (Lawless 2017a; Lawless 2017b), the size of teams was considered to be an open problem (Cooke and Hilton 2015).

## Safety

If the boundary of a well-trained human-machine team includes its human operators, in the context where a human operator threatens the human-machine team or human life, AI can place a system

into a safe mode, as might happen in the future if another copilot attempts to commit suicide, as occurred with Germanwings Flight 9525 in 2015, killing all 150 people aboard (Lawless et al. 2017). Gill Pratt, Toyota's top research executive, said that Toyota is pursuing a semiautonomous track that would rescue the human driver when the human user becomes distracted or inebriated. Per Pratt, the system (Bigelow 2016) would amount to an “autonomous guardian angel ... that would allow humans to maintain control of their vehicles in almost all cases except when it can help them avoid poor decisions or imminent dangers.”

## Emotion

From an individual perspective, a poorer interpretation of the context arises when new information strengthens a user's confirmation bias, increasing disagreement between the two sides of an issue, potentially leading to polarization.<sup>7</sup> The result should be an emotional response arising from the disagreement between two central attitudes or beliefs, as may happen for an individual when an action compromises a central belief (Cooper 2007, p. 182). From the perspective of a team, however, while more research is needed, we have linked low states of emotion to high-performing teams (a ground state) and excited states of interdependence to dysfunctional teams (Lawless 2017b).

Finally, computationally, what shared context even means and whether it affects the performance of systems is not yet settled. However, from a social science perspective, context is a fundamental concept: “any communicative exchange is situated in a social context that constrains the linguistic forms participants use. How these participants define the social situation, their perceptions of what others know, think and believe, and the claims they make about their own and others' identities will affect the form and content of their acts of speaking” (Krauss and Chiu 1998).

## Future Research: Theoretical Issues

The study of context has been the province of social scientists for decades. Recently, however, social scientists have been struggling with a replication crisis (Open Science Collaboration 2015). Moreover, there is money to be made in discovering deficiencies in human cognition by using questionnaires designed to discover inferior cognitive skills (for example, self-esteem [Baumeister et al. 2005]), heretofore unknown biases (such as implicit racism [Blanton et al. 2009]), and social personalities (for example, the Myers–Briggs type inventory [Grant 2013]), despite their lack of validity. For our purposes, these questionnaires are based on individual reports as the *sine qua non* of social science, diminishing the value of context and the most important element of teamwork: interdependence (Lawless, Mittu, and Sofge 2018).

Making this topic even more relevant, swarms of drones in the air, under the sea, and over the land are

entering military service. Nuclear weapons are again a concern. Hypersonic missiles are being developed. As decision-making times are further shortened, this situation underscores the need for machines that can process rapidly changing contexts much faster than can humans — what the Uber car did last March — and to share those changes with their human teammates — what the Uber car did not do. This motivation means we must be able to address the effects of interdependence computationally, which until now has bewildered social scientists (Jones 1998, p. 33). The National Academy of Sciences (Cooke and Hilton 2015), however, reported the presence of interdependence in human teams, especially in the best-performing scientific teams (Cummins 2015), giving it the respectability it needs to become the focus of new theory.

The effect of interdependence on the aggregation of humans befuddled Jones (Lawless 2017a). Interdependence is more than Shannon information. It is the constructive and destructive interference of common social experience. Shannon information can be used to command an aggregation of slaves, like drones, but to accomplish a task, their degrees of freedom are unaffected; however, a human-machine team multitasks to perform its work, reducing its degrees of freedom, an effect we have known about for some time (Kenny, Kashy, and Bolger 1998) but have just begun to exploit. For example, the human brain operates as a unified organ, which is not true of surgically split brains (Gazzaniga 2011); a happily married couple operates as a unified team, which is not true if it becomes public that one partner is cheating on the other (Viteilli 2014); and two well-merged businesses are fully integrated, which is not true if the resistance to merge becomes a public fight, as with CBS and Viacom (Hagey and Flint 2018). However, if we assume that a perfect team is in a ground state, and a dysfunctional, distressed team is in an excited state (computational emotion [Lawless 2017b]), practical applications for AI begin to appear.

When a machine has been trained to operate in a human-machine team, it knows not only what it is supposed to do but also what its human teammate is supposed to do (Lawless, Mittu, and Sofge 2018). If, and only if, we humans let a machine take over when the human part of a team becomes distressed, lives may be saved (Lawless et al. 2017) — for example, let the train take over when its human operator allows it to reach unsafe speeds; let the airliner take over when its commercial copilot attempts to commit suicide; or let the fighter plane take over when its pilot loses consciousness from excessive gravitational forces (Allison 2018). This last is already becoming operational.

## The Special Topic Articles

In this issue of *AI Magazine*, contributors discuss the meaning of shared context and its effect on performance for systems of autonomous human-

machine teams. They address how the interdependence between perception, cognition, and behavior determines the context, whether clear, uncertain, or illusory. Our goal is to use AI to advance the computational construction of context to improve the autonomy of hybrid teams consisting of arbitrary combinations of humans, machines, and robots. We expect that explainable AI is an affordance of computational context.

Lauro Snidaro (University of Udine), Jesus Garcia (Universidad Carlos III de Madrid), Jim Llinas (University at Buffalo), and Erik Blasch (Air Force Office of Scientific Research) have been immensely successful with information fusion (IF) over many years. In their article, “Recent Trends in Context Exploitation for Information Fusion and AI,” they propose to revise IF and AI approaches to exploit context to gain a better understanding of the world and to better adapt fusion tools to specific situations. After reviewing the intricacies of IF, they review their ongoing efforts to incorporate context as a critical part of human reasoning to make for a better fusion product that is generalizable to specific situations. They review the role of context along with the main phases (or waves) of research in AI since 1960 and its uncertainty with relations and ML, and they review fusion, AI, and general intelligence. Their goal is to build an explainable and adaptive model with perception and reasoning generalizable from contexts by the flow of data to develop the next generation of context-sensitive IF systems. Llinas (2016) was an invited speaker at the 2016 AAAI Fall Symposium on AI and the Mitigation of Human Error: Anomalies, Team Metrics and Thermodynamics (Lawless et al. 2017).

In the second article, Integrating Context into Artificial Intelligence: Research from the Robotics Collaborative Technology Alliance, Kristin Schaefer and her team address how to determine the social context in an unstructured, uncertain environment for humans and robots operating in a human-machine team. Schaefer and her team contributed to our past AAAI symposia (Schaefer et al. 2017; Lawless et al. 2018, Chapter 4). In this article, they identify research from the Army Research Laboratory’s Robotics Collaborative Technology Alliance to address the gaps in knowledge that arise in figuring out a robot’s contributions to its team, what the robot knows about the environment and its teammates, and the robot’s intentions as it navigates autonomously through the environment. Operating in the field as part of a team means that a robot’s inferences have to be efficient, drawn quickly, and bidirectional, that is, collaborative yet understandable by all of its teammates. Moreover, the robot’s inferences must be based on what the team is doing in its environment, must be mission specific, must be able to use natural language to construct a model of its view of the world, and must be able to contribute its novel findings to build a shared context.

“Context-Driven Proactive Decision Support for Hybrid Teams,” by Manisha Mishra, Pujitha Mannaru, David Sidoti, Adam Bienkowski, Lingyi Zhang, and Krishna R. Pattipati, was written specifically for the maritime domain. In the article, they describe their key challenge as identifying the context within which humans interact with a smart Internet of Things. They define context interdependently as the evolving multidimensional feature space consisting of a ship’s mission and its goals, assets, threats and tasks, and the cognitive states of its commander and human operators working as a team in uncertain environments while at sea, including hybrid human-machine teams. They have created and validated an operational system for proactive decision making amid a host of technical challenges posed by the integration and allocation of assets and tasks for an Internet of Things using AI to determine context and achieve superior performance. In addition, they provide more details, mathematics, and descriptions about a user test bed in supplementary material online.

In their article “Identifying Critical Contextual Design Cues Through a Machine Learning Approach,” Missy Cummings and Alex Stimpson of Duke University review the safety and productivity benefits of autonomous technologies with a goal of understanding how autonomous systems can be better designed to improve the interactions between humans working with or around autonomous systems. These safety critical systems generate immense amounts of data. The authors review and use ML to design a human-user interface. They evaluate a proposed pedestrian signaling display mounted on a driverless car through traditional inferential statistics that looked at broad population characteristics, finding no significant relationships. Instead, by paying attention to individual user characteristics using a ML clustering approach, they uncover critical contextual cues that have led to improved reaction times for one variant of the pedestrian signaling display.

Brian Jalaian, Michael Lee, and Stephen Russell address how different sources of uncertainty affect the interpretations of contexts differently when using different ML methods in “Uncertainty Quantification in Machine Learning.” They start with a review of basic statistics, observational errors, models and errors, and optimizations under uncertainty. Then they review an autonomous mission command architecture, statistical learning and stochastic optimization, and uncertainty in ML. Finally, they address four sources of uncertainty — noise, parameter uncertainty, the uncertainty in model specification, and the uncertainty from extrapolation — providing readers with a nonparametric Bayesian model that graphically portrays the uncertainty in the model. They address the motivation to resolve these uncertainties as critical to the application of ML in the field, where erroneous forecasts may put lives at risk.

The final article, written by Erik Blasch, Robert Cruise, Alex Aved, Uttam Majumder, and Todd Rovito, proposes a method of decisions to data that provides

a path to establish the value of data foraging, collecting data, and sense making, using AI with human reasoning to assess the context for complex sets of data. Their model addresses data in various states (rest, motion, fusion, transition, and use). They review AI dynamic data-driven applications systems and IF and how these paradigms align with AI; reasoning contexts; types of ML; and applications for situational understanding that serve to achieve human-machine awareness. To determine the context of a dynamic target, they provide an example with AI and deep multimodal image fusion where the data are collected in multiperspectives from a command-guided swarm.

## Conclusions

We hope that readers enjoy all six of the articles contributed on the topic of artificial intelligence, autonomy, and human-machine teams: interdependence, context, and explainable AI. We also hope that readers will join us at a future AAAI symposium on the topic. The advent of human-machine teams has created a time of intellectual ferment, extraordinary technological advances, and the introduction of interdependence to mathematicians, physicists, and AI theorists and practitioners.

## Notes

1. Merriam-Webster, s.v. “context,” [www.merriam-webster.com/dictionary/context](http://www.merriam-webster.com/dictionary/context).
2. [en.oxforddictionaries.com/definition/context](http://en.oxforddictionaries.com/definition/context).
3. Russian deaths have been denied, and American involvement has been denied.
4. *American Electric Power Co., Inc., et al., Petitioners v. Connecticut et al.*, 564 U.S. 10-174 (2011), [www.supremecourt.gov/opinions/10pdf/10-174.pdf](http://www.supremecourt.gov/opinions/10pdf/10-174.pdf). Justice Ginsburg wrote the unanimous opinion.
5. For example, Nick Saban, the coach of the University of Alabama football team, has never lost a game against his former assistant coaches who had taken coaching jobs at competing schools (Kirshner 2018); there have been many failed attempts to clone Silicon Valley (Lucky 2014); and despite studies on how to recreate the innovation culture at Bell Labs (Kelly and Caplan 1993) existing in a facility where the work for several Nobel awards was completed, the facility has since closed (Martin 2006) and the lab has been renamed Nokia Bell Labs.
6. For more on these issues, see [aaai.org/Symposia/Spring/sss19symposia.php#ss01](http://aaai.org/Symposia/Spring/sss19symposia.php#ss01).
7. Personal communication with C. Sibley, May 26, 2009.

## References

- Adelson, E. H. 2000. Lightness Perceptions and Lightness Illusions. In *The New Cognitive Sciences*, 2nd ed., edited by M. Gazzaniga. Cambridge, MA: MIT Press.
- Allison, G. 2018. F-35 to Incorporate Automatic Ground Collision Avoidance System ‘Five Years Earlier Than Planned.’ UK Defence Journal (February 2). [ukdefencejournal.org.uk/f-35-incorporate-automatic-ground-collision-avoidance-system-five-years-earlier-planned/](http://ukdefencejournal.org.uk/f-35-incorporate-automatic-ground-collision-avoidance-system-five-years-earlier-planned/)

- Baumeister, R. F.; Campbell, J. D.; Krueger, J. I.; and Vohs, K. D. 2005. Exploding the Self-Esteem Myth. *Scientific American* 292(1): 84–91. doi.org/10.1038/scientificamerican0105-84
- Bigelow, P. 2016. Toyota Wants Autonomous Guardian Angel to Help Drivers Instead of Autonomy Taking the Driver out of the Picture. *Autoblog* (blog). April 8. www.autoblog.com/2016/04/08/toyota-autonomous-driving-guardian-angel-featured/
- Blanton, H.; Jaccard, J.; Klick, J.; Mellers, B.; Mitchell, G.; and Tetlock, P. 2009. Strong Claims and Weak Evidence: Reassessing the Predictive Validity of the IAT. *Journal of Applied Psychology* 94(3): 567–82. doi.org/10.1037/a0014665
- Brill, E., and Mooney, R. J. 1997. An Overview of Empirical Natural Language Processing. *AI Magazine* 18(4): 13–24.
- Brynjolfsson, E., and Mitchell, T. 2017. What Can Machine Learning Do? Workplace Implications. *Science* 358: 1530–4. doi.org/10.1126/science.aap8062
- Cooke, N. J., and Hilton, M. L., editors. 2015. *Enhancing the Effectiveness of Team Science*. Washington, DC: National Academies Press.
- Cooper, J. M. 2007. *Cognitive Dissonance: 50 Years of a Classic Theory*. London: SAGE Publishing.
- Cummings, J. 2015. Team Science Successes and Challenges. Paper presented at the National Science Foundation Sponsored Workshop on Fundamentals of Team Science and the Science of Team Science. Bethesda, MD, June 2.
- Doolen, T. L.; Hacker, M. E.; and Van Aken, E. M. 2003. The impact of organizational context on work team effectiveness: a study of production teams. *IEEE Transactions on Engineering Management* 50(3): 285–96. doi.org/10.1109/TEM.2003.817296
- Dehaene, S.; Lau, H.; and Kouider, S. 2017. What Is Consciousness and Could Machines Have It? *Science* 359(6374): 400–2. doi.org/10.1126/science.aan8871
- Edelman, F., and Roughead, G. 2018. *Providing for the Common Defense. The Assessment and Recommendations of the National Defense Strategy Commission*. Washington, DC: United States Institute of Peace. www.usip.org/sites/default/files/2018-11/providing-for-the-common-defense.pdf
- Emre, M. 2018. *The Personality Brokers: The Strange History of Myers-Briggs and the Birth of Personality Testing*. New York: Doubleday.
- Felson, M., and Eckert, M. A. 2015. *Crime and Everyday Life*. Fifth edition. Thousand Oaks, CA: SAGE Publications.
- FitzGerald, D., and Ramachandran, S. 2018. AT&T Boss Writes Script for HBO: More Data, More Money. *Wall Street Journal* (September 12). www.wsj.com/articles/as-it-absorbs-time-warner-at-t-expects-profitability-next-year-1536760417
- Freedberg, Jr., S. J. 2018. New Tests Prove IBCS Missile Defense Network DOES Work: Northrop. *Breaking Defense* (August 15). breakingdefense.com/2018/08/new-tests-prove-ibcs-missile-defense-network-does-work-northrop/
- Fritz, B. 2017. Hollywood Hopes 'It' Can Pull Box Office out of the Gutter. *Wall Street Journal* (September 4). www.wsj.com/articles/hollywood-hopes-it-can-pull-box-office-out-of-the-gutter-1504488416
- Garner, D. 2016. Seven Brief Lessons on Physics Is Long on Knowledge. *New York Times* (March 22). www.nytimes.com/2016/03/23/books/review-seven-brief-lessons-on-physics-is-long-on-knowledge.html?\_r=0
- Gazzaniga, M. S. 2011. *Who's in Charge? Free Will and the Science of the Brain*. New York: Ecco.
- Geist, E., and Lohn, A. J. 2018. *How Might Artificial Intelligence Affect the Risk of Nuclear War?* Santa Monica, CA: RAND Corporation. doi.org/10.7249/PE296
- Gershensfeld, N. 1999. *When Things Begin to Think*. New York: Henry Holt.
- Grant, A. 2013. Goodbye to MBTI, the Fad That Won't Die. *Give and Take* (blog), *Psychology Today*, September 18, 2013. www.psychologytoday.com/us/blog/give-and-take/201309/goodbye-mbti-the-fad-won-t-die
- Grove, T. 2018. Drone Attacks on Russian Bases in Syria Expose Security Holes. *Wall Street Journal* (January 15). www.wsj.com/articles/drone-attacks-on-russian-bases-in-syria-expose-security-holes-1516017608
- Hackman, J. R. 2011. Six Common Misperceptions About Teamwork. *Harvard Business Review* (June 7). hbr.org/2011/06/six-common-misperceptions-about
- Hagey, K., and Flint, J. 2018. Once Allies, Two Media Chiefs Go to War over the Future of CBS. *Wall Street Journal* (May 28). www.wsj.com/articles/once-allies-two-media-chiefs-go-to-war-over-the-future-of-cbs-1527530599
- Irving, D. 2018. How Artificial Intelligence Could Increase the Risk of Nuclear War. *The RAND Blog* (blog). April 28, 2018. www.rand.org/blog/articles/2018/04/how-artificial-intelligence-could-increase-the-risk.html
- James, M. 2018. With Leslie Moonves Out, Shari Redstone Emerges Winner in Fight over Control of CBS. *Los Angeles Times* (September 11). www.latimes.com/business/hollywood/la-fi-ct-shari-redstone-moonves-20180911-story.html
- Jones, E. E. 1998. Major Developments in Five Decades of Social Psychology. In *The Handbook of Social Psychology*, edited by D. T. Gilbert, S. T. Fiske, and G. Lindzey, 3–57. New York: McGraw-Hill.
- Jussim, L. 2017. Mandatory Implicit Bias Training Is a Bad Idea. *Rabble Rouser* (blog), *Psychology Today*, December 3, 2017. www.psychologytoday.com/us/blog/rabble-rouser/201712/mandatory-implicit-bias-training-is-bad-idea
- Kambhampati, S. 2018. Challenges of Human-Aware AI Systems. Address presented at the 32nd AAAI Conference on Artificial Intelligence. New Orleans, LA, February 4.
- Kelly, R., and Caplan, J. 1993. How Bell Labs Creates Star Performers. *Harvard Business Review* (July–August). hbr.org/1993/07/how-bell-labs-creates-star-performers
- Kenny, D. A.; Kashy, D. A.; and Bolger, N. 1998. Data Analyses in Social Psychology. In *Handbook of Social Psychology*, edited by D. T. Gilbert, S. T. Fiske, and G. Lindzey, 233–65. 4th edition. New York: McGraw-Hill.
- Kirshner, A. 2018. Kirby Smart Got Close Again, but Nick Saban's Former Assistants Still Can't Beat Him. *SBNation* (December 1). www.sbnation.com/college-football/2018/12/1/18121541/nick-saban-kirby-smart-sec-championship
- Kissinger, H. A. 2018. How the Enlightenment Ends. *The Atlantic*. (June) www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124
- Krauss, R. M., and Chiu, C.-Y. 1998. Language and Social Behavior. In *Handbook of Social Psychology*, edited by D. T. Gilbert, S. T. Fiske, and G. Lindzey, 41–88. 4th edition. New York: McGraw-Hill.
- Lawless, W. F.; Akiyoshi, M.; Angjellari-Dajcic, F.; and Whitton, J. 2014. Public Consent for the Geologic Disposal of Highly Radioactive Wastes and Spent Nuclear Fuel. *International Journal of Environmental Studies* 71(1): 41–62.

- Lawless, W. F. Mittu, R., Sofge, D. and Russell, S., editors. 2017. *Autonomy and Artificial Intelligence: A Threat or Savior?* Berlin: Springer.
- Lawless, W. F. 2017a. The Entangled Nature of Interdependence. Bistability, Irreproducibility and Uncertainty. *Journal of Mathematical Psychology* 78: 51–64. doi.org/10.1016/j.jmp.2016.11.001
- Lawless, W. F. 2017b. The Physics of Teams: Interdependence, Measurable Entropy and Computational Emotion. *Frontiers in Physics* 5. doi.org/10.3389/fphy.2017.00030
- Lawless, W. F. 2018. Resilience, Teamwork (Interdependence) and Context: Defeating Skynet. Paper presented at Resilience Week 2018. Denver, CO, August 20–23.
- Lawless, W. F., R. Mittu, and D. A. Sofge, editors. 2018. *Computational Context: The Value, Theory and Application of Context with AI*. Boca Raton, FL: CRC Press. doi.org/10.1201/9780429453151
- Llinas, J. 2016. Argumentation-Based System Support Toolkit for Intelligence Analyses. Paper presented at the AAAI 2016 Symposium on AI and the Mitigation of Human Error: Anomalies, Team Metrics and Thermodynamics. Palo Alto, CA, March 21–23.
- Lucky, R. W. 2014. Why Cloning Silicon Valley Is So Hard. *IEEE Spectrum* (blog). April 28, 2014. spectrum.ieee.org/at-work/innovation/why-cloning-silicon-valley-is-so-hard).
- Lukianoff, G., and Haidt, J. 2018. *The Coddling of the American Mind: How Good Intentions and Bad Ideas Are Setting Up a Generation for Failure*. New York: Penguin Press.
- Martin, A. 2006. Pastoral Site of Historic Inventions Faces the End. *New York Times* (June 14). www.nytimes.com/2006/06/14/realestate/commercial/14bell.html?\_r=1&n=Top/News/Business/Small%20Business/Innovation&oref=slogin
- Marwah, K.; Katzin, D.; Zollanvari, A.; Noy, N. F.; Ramoni, M.; and Alterovitz, G. 2012. Context-Specific Ontology Integration: A Bayesian Approach. In *AMIA Joint Summits on Translational Science Proceedings AMIA Summit on Translational Science*, 79–86. Bethesda, MD: American Medical Informatics Association.
- Massey, C., and Thaler, R. 2005. Overconfidence vs. Market Efficiency in the National Football League, NBER Working Paper 11270. Cambridge, MA: National Bureau of Economic Research.
- Mattis, J. 2018. Nuclear Posture Review, Secretary of Defense, Final Report. Washington, DC: Department of Defense.
- McClelland, J. L., and Rumelhart, D. E. 1988. *Explorations in Parallel Distributed Processing — A Handbook of Models, Programs, and Exercises*. Cambridge, MA: MIT Press.
- Mishra, M.; Sidoti, D.; Avvari, G. V.; Mannaru, P.; Fernando Martinez Ayala, D.; and Pattipati, K. R. 2017. Context-Driven Proactive Decision Support: Challenges and Applications. In *Computational Context: Why It's Important, What It Means, and Can It Be Computed? Papers from the 2017 AAAI Spring Symposium*. Technical Report SS-17-03. Palo Alto, CA: AAAI Press. aaii.org/Library/Symposia/Spring/ss17-03.php
- National Transportation Safety Board. 2018. *Preliminary Report Released for Crash Involving Pedestrian, Uber Technologies, Inc., Test Vehicle*. Washington, DC: NTSB. www.nts.gov/news/press-releases/Pages/NR20180524.aspx
- Open Science Collaboration. 2015. Estimating the Reproducibility of Psychological Science. *Science* 349 (6251): aac4716. doi.org/10.1126/science.aac4716
- Payne, K. B. 2018. Deterrence: The 2018 NPR, Deterrence Theory and Policy. *Real Clear Defense* (August 13). www.realcleardefense.com/articles/2018/08/13/deterrence\_the\_2018\_npr\_deterrence\_theory\_and\_policy\_113700.html
- Pearl, J. 2002. Reasoning with Cause and Effect. *AI Magazine* 23(1): 95–111. doi.org/10.1609/aimag.v23i1.1612
- Pearl, J., and Mackenzie, D. 2018. AI Can't Reason Why. *Wall Street Journal* (May 18). www.wsj.com/articles/ai-cant-reason-why-1526657442
- Pfautz, S. L.; Ganberg, G.; Fouse, A.; and Schurr, N. 2015. A General Context-Aware Framework for Improved Human-System Interactions. *AI Magazine* 36(2): 42–9. doi.org/10.1609/aimag.v36i2.2582
- Pitt, D. 2018. Mental Representation. In *The Stanford Encyclopedia of Philosophy* (Winter 2018). plato.stanford.edu/archives/win2018/entries/mental-representation
- Rothman, R. 2018. Driverless Cars and the Perpetual Moral Panic. *Commentary* 20 (March). www.commentarymagazine.com/american-society/driver-less-cars-moral-panic/
- Schaefer, K. E.; Chen, J. Y. C.; Wright, J.; Aksaray, D.; and Roy, N. 2017. Challenges with Incorporating Context into Human-Robot Teaming. In *Computational Context: Why It's Important, What It Means, and Can It Be Computed? Papers from the AAAI 2017 Spring Symposium*. Palo Alto, CA: AAAI Press. aaii.org/Library/Symposia/Spring/ss17-03.php
- Shultz, G. P. 2018. America Can Ride the 21st Century's Waves of Change. *Wall Street Journal* (June 27). www.wsj.com/articles/america-can-ride-the-21st-century-s-waves-of-change-1530139914?shareToken=st16d75090420344f3816-f775f61f7e249&oref=article\_email\_share
- Sober, E. 2009. The Power of Context. Getting Beyond Labels and Limitations. *The 99th Monkey* (blog), *Psychology Today*, July 10, 2009. www.psychologytoday.com/blog/the-99th-monkey/200907/the-power-context
- Speaks, J. 2017. Theories of Meaning. In *The Stanford Encyclopedia of Philosophy* (Fall 2017). plato.stanford.edu/archives/fall2017/entries/meaning/
- Stempel, J. 2016. Wells Fargo Admits Deception in \$1.2 Billion U.S. Mortgage Accord. *Reuters* (April 8). www.reuters.com/article/us-wellsfargo-settlement-idUSKCN0X52HK
- Viteilli, R. 2014. After Infidelity. How Can Couples Learn to Move on Following One Partner's Infidelity? *Media Spotlight* (blog), *Psychology Today*, April 28, 2014. www.psychologytoday.com/us/blog/media-spotlight/201404/after-infidelity
- Wakabayashi, D. 2018. Uber's Self-Driving Cars Were Struggling before Arizona Crash. *New York Times* (March 23). www.nytimes.com/2018/03/23/technology/uber-self-driving-cars-arizona.html
- Weinberg, S. 2017. Steven Weinberg and the Puzzle of Quantum Mechanics. *The New York Review of Books* (April 6). www.nybooks.com/articles/2017/04/06/steven-weinberg-puzzle-quantum-mechanics/
- Woo, S. 2018. Why Being First in 5G Matters. The U.S., China, South Korea and Japan All See a Big Payoff from Winning the Battle for the Wireless Future. *Wall Street Journal* (September 12). https://www.wsj.com/articles/why-being-first-in-5g-matters-1536804360
- Zell, E., and Krizan, Z. 2014. Do People Have Insight into Their Abilities? A Metasynthesis. *Perspectives on Psychological Science* 9(2): 111–25. doi.org/10.1177/1745691613518075
- In 1983, **William Lawless** blew the whistle on Department of Energy (DOE) mismanagement of military radioactive wastes. After his PhD, he joined DOE's citizen advisory board at its Savannah River Site where he coauthored over 100 recommendations on its cleanup. His research today is





## Third Conference on AI, Ethics, and Society

*Colocated with AAAI-20 from February 7–8, 2020 in New York, NY, USA*

[www.aies-conference.com](http://www.aies-conference.com)

on interdependence for teams (human-machine teams). He advised the Naval Research and Development Enterprise's Applied Artificial Intelligence Summit in 2018; co-edited four books; published widely; and co-organized eight AAAI symposia at Stanford with a ninth in 2019 on shared context.

**Ranjeev Mittu** is the Branch Head for the Information Management and Decision Architectures Branch within the Information Technology Division, US Naval Research Laboratory. His research expertise is in multi-agent systems, artificial intelligence, machine learning, data mining, pattern recognition and anomaly detection. He has won an award for transitioning technology solutions to the operational community, and has coauthored one book, coedited four books, and written numerous book chapters, articles, and conference publications. He has served on scientific exchanges, as a subject matter expert and Technology Evaluation Boards. He has an MS in Electrical Engineering from The Johns Hopkins University.

**Donald Sofge** is a computer scientist and roboticist at the U.S. Naval Research Laboratory (NRL) with more than

30 years of experience in AI and Control Systems R&D. He has served as PI or Co-PI on dozens of federally funded R&D programs, and has numerous publications on autonomy, intelligent control, quantum computing, including 5 books and one patent. He leads the Distributed Autonomous Systems Group at NRL where he develops nature-inspired computing solutions to problems in sensing, AI, and autonomous robotic systems control, including autonomous teams or swarms of robotic systems for Navy missions.

**Laura Hiatt** is a research scientist at the U.S. Naval Research Laboratory. She received her BS in symbolic systems from Stanford University and her PhD in computer science from Carnegie Mellon University. Hiatt's work has primarily focused on ways in which humans and robots can effectively work together as teammates. The research involves issues of planning and reasoning, human situational awareness, and team-based task communication strategies. Much of her work has also involved developing computational cognitive models of human cognition, and leveraging them to improve the ability of robots to team with humans and accomplish their tasks.