AI and Critical Systems

From hype to reality

Phil Laplante, Penn State
Dejan Milojicic, Sergey Serebryakov, Hewlett Packard Labs
Daniel Bennett, NREL

Artificial intelligence will be deployed in increasingly more systems that affect the health, safety and welfare of the public. These systems will better utilize scarce resources, prevent disasters and increase safety, reliability, comfort and convenience. Despite the technological challenges and public fears these systems will improve the quality of life of millions of people worldwide.

Prediction

The use of artificial intelligence (AI) in critical infrastructure systems will increase significantly over the next five years. Critical infrastructure systems or, more simply, “critical systems” are those that directly affect the health, safety and welfare of the public and in which failure could cause loss of life, serious injury or significant loss of assets or privacy. Critical systems include power generation and distribution, telecommunications, road and rail transportation, healthcare, banking and more [Moteff].

AI and Software

AI plays an important role in some of humanity’s most complex systems, especially in safety-critical systems. In critical systems, software is generally involved in controlling the behavior of electromechanical components and monitoring their interactions [Wong], but it is also used in many other ways. AI in critical systems can involve pattern matching, and/or decision making, prognostics and predictive analytics, anomaly detection and more. In a simple scenario, AI can serve a significant benefit in automating many of the mundane tasks that in the past would have required humans (e.g. analysts) to sift through massive amounts of data in order to derive information for which decisions would need to be made off of and in many cases the AI can also make many of those decisions if properly trained. While AI can be implemented in hardware, firmware or software, the design, implementation and testing must all be concerned with very high safety, security, and reliability margins. Ultimately, AI for critical systems must combine real-time analysis with robust network communications structures to continually adapt to changing circumstances.

Today’s AI is different from general software in the following way. There is a need for training of current AI algorithms, with the possible evolution towards self-learning and understanding. The outcome of this training is used as a black box leading to a lack of “explainability” in the use of trained algorithms. Such training can be a cause of a bias (a vulnerability), because training is only as good as the data used for the training. Finally, compared to traditional software, in AI there is a more pressing need for ethical considerations.

In critical systems, internal and external interactions, timing, and general processing errors can lead the software to an unsafe state or lead to a system failure. AI can be employed to help avoid or recover from these unsafe states. When the AI in critical systems does not operate as intended (including to prevent and not contribute to system failure) there can be serious consequences. The consequences can range from minor performance anomalies to a catastrophic failure leading to significant loss of money and property, injury and loss of human life, perhaps on a large scale [Wong]. For these reasons, it is required to have AI software that provides explainable results. The recommendations made need to be predictable and repeatable across a wide variety of inputs, in terms of timing, bias, and results.

AI Advances and Challenges

Advances in data analytics, machine intelligence, deep learning and related artificial intelligence (AI) technologies have and will continue to motivate critical systems design. These technologies leveraged by more accurate image recognition and pattern matching, the Internet of Things (IoT), edge computing and security technologies such as advanced encryption, hardware accelerators and more will drive the increase in deployments and in public confidence and trust in these systems. These systems will exhibit high levels of connectivity, intelligence, automation coupled with AI/machine learning enriched cybersecurity. IoT devices integrated with robotic process automation (RPA) allow secure robust communication among sensors, actuators, and power sources [Lange].

AI is being deployed on wide ranging systems from data centers to edge devices. Systems are becoming more responsive in thinking, perceiving, and acting within time performance constraints. Designers are more confident in applying multiple technology advancements to solve volatile, uncertain, complex, and ambiguous challenges.

There are a number of AI challenges that need to be addressed in order to successfully apply AI models in critical systems. Models are only as good as the data used to train these models. This requires that the following aspects be addressed.

Model bias: Overrepresentation of one example and underrepresentation of other examples (unbalanced data) makes a model biased towards a major class or classes. In social domains such as healthcare or finance, this may result in unfair and unethical decisions. It is not uncommon to have much more unlabeled data than labeled, and some sort of mechanism that automatically validates AI models is required.

Adversarial attacks: With the rise of Deep Learning models, a new trend has emerged in security known as adversarial attacks. This type of an attack uses data that at a macro level looks like real data (for instance, a road sign), but at a micro level it was modified so that it has a dramatic impact on a model's decision. AI models need to be either robust enough to tampered inputs or need to be accompanied by other AI models that check if input data is from a set of expected inputs.

Data security: AI models are all about datasets that constantly grow in size. An attacker may modify a dataset by changing existing examples or introducing new ones so that a model learns that adverse behavior. Special security protocols and frameworks need to be introduced to ensure the validity of datasets.

Model security: In the future systems, AI models will be deployed everywhere – from data centers to edge and wearable devices. This is significantly different from current deployments where devices are assumed to be located in secure facilities. Systems hosting AI models need to be able to verify the validity of models, identify attempts to modify them and re-deploy them in case they are compromised.

Trust: Using systems with a high level of automation (which is not common nowadays) such as self-driving vehicles means that people need to learn to trust those systems. AI, in particular deep and reinforcement learning techniques, has and will continue to have significant impact on society in almost every area, leading to an extreme need for public trust.

Explainability and self-assessment: An AI-based model or a control system needs to be able to continuously defend/explain its decisions. These models need to be able to identify situations in which they are not confident in making the right decisions and inform human operators that they need to take control. Hence, AI models must be explainable, which is a challenge because many of the models are used now as black boxes for which outputs are hard or impossible to explain.

Legal: One of the important factors that may delay wide deployments of AI models is the legal one. Certain certification and validation procedures will need to be introduced in order to enable wide adoption of AI in critical systems.

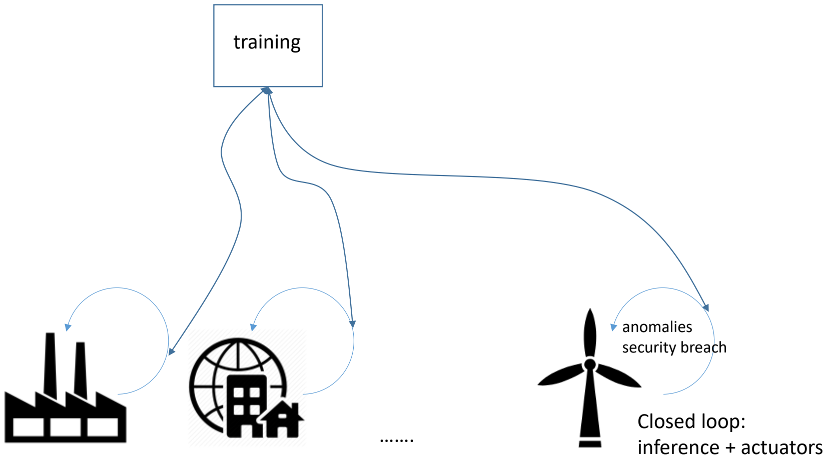

Figure 1. In addition to traditional inference and training for IoT and back-end, Critical infrastructure also has real-time loops at the edge that may require actuators to trigger actions as a result of detected anomalies or security breaches.

Predictive Analytics

In critical infrastructure systems robust cybersecurity is a requirement and this is where we expect increased application of AI [Bresniker]. Cyberattacks on an electrical grid, air traffic control system, or telecommunications infrastructure can have devastating effects. AI will continue to play an increased role in enhancing cybersecurity. AI enables analysis of large quantities of stored and streamed data to understand the threat space. Both supervised and unsupervised learning methods can play a significant role in anomaly detection to prevent and/or to help allow a system to respond quickly to possible nefarious activity and thereby helping to mitigate the effects that it may have. In post incident analysis, supervised and unsupervised learning can also be used on this data to predict the cause of the intrusion and monitor the system to prevent future incursions.

AI can also be used to address a frequent source of faults in critical systems, that is, the interaction between components and/or subsystems. These interactions are prone to design errors, especially when there are changes in new products. In such systems, advanced diagnostics and prognostics, driven by AI, can be employed to isolate the root cause of faults and to anticipate component failure, enabling repair or replacement pre-incident. The systemic linking of failure mechanisms to system life cycle management is referred to as prognostics and health management (PHM). PHM is used to help anticipate failure so that components can be replaced before they fail, and other precautionary measures can be taken. Using PHM is an important part of achieving ultra-high reliability and safety in critical systems [Zhang].

Extensive data collection of retrospective data and real-time data is an important part in implementing successful PHM systems. Numerous sensors need to be added to systems to collect real-time data.

Furthermore, the entire PHM data collection program must be planned early in the project life cycle and closely coordinated with field support and maintenance personnel [Janasak].

Applications

For critical systems AI models can be used in knowledge reasoning, planning, natural language processing, computer vision, robotics and artificial general intelligence, making it highly suited for a wide array of critical systems applications. AI can be used for improving performance optimization, failsafe operation, fraud and intrusion prevention and detection, health prognostics and post failure analysis.

Typical domains for critical systems include healthcare, transportation, IT and utilities infrastructure, commerce and finance. Table 1 shows application domains for critical infrastructure, which is adapted from a US Congressional study on critical infrastructures identification [Moteff], along with example systems, and typical AI uses.

Table 1. Applications domains for critical systems, example systems and possible uses for AI. Key -- FA: failure analysis, FD: fraud detection, FS: failsafe operation, ID: intrusion prevention and detection, PHM: system prognostics and health management.

Application Domain	Example Systems	Typical AI Uses
telecommunication infrastructure	public telephone network, local branch exchange	ID, FA
water supply systems	water treatment plant, dam control	FA, FS, PHM
electrical power systems	nuclear power plant, regional electrical grid	FA, FS, PHM, PO
oil and gas generation and distribution	gas pipeline, gas powered power plant	FA, FS, PHM, PO
roadway transportation systems	Smart interstate highway, traffic monitoring and control	FA, FS, PHM, PO
railway transportation systems	high speed train line, metropolitan train network control	FA, FS, PHM, PO
air transportation systems	air traffic control system network, passenger aircraft autopilot	FA, FS, PHM, PO
banking and financial services	pension fund management, stock market management	FD, ID
public safety services	air passenger screening, police dispatch	FD, ID, PO
healthcare systems	robotic surgery, healthcare record management	FD, ID, PO
administration and public services	employee personnel database, retirement management	FD, ID

Failure Analysis (FA) is determined by well-defined rules and protocols. AI-based engine uses knowledge about incident (collection of evidences) to infer new knowledge and propose next steps. Fraud (FD) and Intrusion (ID) detection problems are solved using methods of anomaly detection. Nowadays, these methods are usually either statistical based or based on methods of machine learning. Each anomaly detection model makes decision based on one small number of information sources. Future systems will be based on deep learning methods, such as autoencoders, and will be able to find anomalies analyzing a large collection of input signals simultaneously. Failsafe operations (FS) will be achieved by using ensembles of AI models to ensure that if one fails others continue making decisions. These models will be controlled by intelligent agents that will be communicating with each other according to specific protocol to coordinate their decisions. For example, if one agent detects that it is not able to make credible decision due to unexpected conditions, it will seek for help from other agents asking what their decisions are about current situation. System prognostics and health management (PHM) will be achieved by a number of AI models deployed on respective devices or systems that will constantly be monitoring operational telemetry in order to predict what components are likely to fail within next minute, hour, day etc. Performance optimization (PO) of systems will be achieved by creating “Digital Twins.” They are AI-based models, such as deep neural networks, that model response to input signal. Feeding in signals and their various combinations, it will be possible to determine the optimal configuration that minimizes specific criteria.

But the use of AI In critical systems is not necessarily limited to these domains. Other public facing critical systems where failure could lead to injury or even death can be found in consumer systems (e.g. hot food vending machines, smart car wash) entertainment (e.g. interactive amusement park rides, virtual reality games, electronic toys) and in the home (e.g. smart homes, smart appliances).

Impact

Various impactful scenarios that will be seen in the next five years include:

In the telecommunication and power systems infrastructure, AI enhancements in machine learning will bring new levels of security including threat anticipation and mitigation. Networks will be fully self-reconfigurable to adapt to outages and peak load patterns. Robotic inspection, powered by enhanced AI, will enable better maintenance programs for remote and hard-to-reach assets.

In water supply systems we expect AI to improve water quality through enhanced monitoring algorithms. Rural and poorly served areas will see increases in the availability of fresh water due to more reliable, AI-enhanced filtration, purification and pumping systems. These systems can also promote better stewardship and sustainability of natural resources through advanced data analytics. The same is true for monitoring sewage for the epidemics and drug usage [Pretz].

Electrical power and oil and gas generation and distribution systems will continue to be improved through the use of AI. Already significant improvements can be seen in AI driven electrical supply load balancing and diagnosis of faults. But even in “advanced” countries like the US, components in the electrical power grid and in oil and gas distribution can be more than 100 years old. These aging infrastructures can strongly benefit from PHM to prevent critical failures and enable timely retrofitting [Daniels]. The majority of the grid is reliably and inertially driven largely up to this point via the use of fossil fuel or even hydro (e.g. dams) driven generation. This power is then distributed in most cases over long distances and the inherent inefficiencies with transmission line losses as well as the reliance on the use of fossil based (hydro excluded). Distributed Energy Resources (DERs) at the grid edge serve the benefit of being able to provide power more proximate to the demand and without relying on fossil fuels but are obviously not as consistent in terms of the amount and frequency for which they are delivering power. Batteries are then used in concert with these renewable resources to store power from these resources in order to be able to deliver power particularly at times when the resources are not generating power (e.g. when the sun is not shining in the case of a photo-voltaic cell). Batteries can also be leveraged during peak power consumption times and their charge augmented by bulk grid provided power as necessary during off-peak times in order to minimize costs and to even out demands and efficiencies throughout the day. A key challenge with these systems is the synchronous control of DC to AC conversion needed for the DERs and/or battery storage. Additionally, extra power generated by the DERs and/or provided by the battery storage can be leveraged back to the grid for broader use and credited to producers/consumers via some bartering or market type mechanism. All of these aspects can and should be managed and optimized via the use of artificial intelligence [Khan].

Autonomous and semi-autonomous vehicles are already being deployed in many cities. More fully autonomous vehicles will proliferate worldwide. But we also expect to see rapid increases in numbers of connected vehicles, and correspondingly, we expect pilot and small-scale applications for intelligent roads and highways, which are already deployed, to expand throughout the world. Intelligent transportation systems provide advanced services through analytics such as monitoring traffic and identifying speeders. Smart highways, interoperating with other systems such as drones [Frachtenberg] and traffic awareness services (e.g., Waze), can improve traffic flow, prevent traffic and accidents, protect the safety of bikers, pedestrians, and crossing wildlife, and help drivers find parking spaces [Laplante]. The advent of 5G and other high bandwidth mechanisms will facilitate the ability of autonomous vehicles to negotiate machine-to-machine the flow of traffic at intersections. AI will make these developments possible and we can foresee the end of traffic lights and even human drivers in the future. However with these possibilities we need to be cognizant of the cyber-security implications of these capabilities and be sure to incorporate this as these technological advances continue to evolve and be developed. Vehicles will continue to evolve to be more and more electrical which serve as mobile devices to be charged and provide charge at different locations tied in to the grid. The management of how and when to charge and discharge them as necessary provides further AI managed optimization for the future.

In railway transportation systems AI is already used to improve on-time and safe operation, and autonomous railway lines are readily found throughout the world (e.g. at airports and amusement parks). But these systems, powered by advanced AI, will need to further integrate with other smart infrastructure to provide optimal benefits, including enhanced cybersecurity.

In air transportation systems AI can improve safety, on-time arrivals, reduced fly times, and can protect against cyber-intrusion. Even regional airports can use advanced AI to increase security. Improvements in autopilot systems can improve on-time arrival, reduce accidents and improve fuel consumption statistics.

In banking and financial services AI provides enhanced security and fraud detection and prevention. Certain populations, particularly the elderly, are very susceptible to scams, and machine learning algorithms are key to protecting their assets. We expect these technologies to continue to advance and increase public confidence in all financial systems [Qi].

In public safety the benefits of AI are already being seen in threat identification through facial recognition and behavioral pattern matching. Increasingly, deep learning algorithms will be used to monitor and identify threats in public spaces. Law enforcement will increasingly use AI to cope with the large scale of information involved in criminal investigations, making them more effective and the public safer. In public safety services AI can be used for more efficient and resilient operation of first responder and ambulance dispatch systems. Security robots for monitoring and surveillance will find increasing use in public spaces, in stores (e.g. for cleanup, restocking), malls, and even in private settings. Significantly more deployments will occur, despite possible reluctance from the public about privacy issues.

In healthcare there are numerous deployed AI applications in patient diagnosis and in robotics. Robotics applications include surgery, hospital supply delivery and robotic companions. There are also robots in use in home care to assist in lifting patients to transfer from bed to wheelchair, lead patients in exercise, and engage in simple conversations. Companion robots can also help vulnerable populations, such as the elderly or disabled by providing comfort, emotional support and information. We expect to see a significant increase in all of these applications of AI in healthcare and many new ones to emerge in the next five years.

Finally, in public administration and public services AI can provide citizens with more trustworthy, transparent, and reliable services (e.g. tax collection, real estate assessment) through enhanced cybersecurity, anticipatory responses, resilience and fault-tolerance and recovery.

Technology Challenges

Critical systems are very expensive to develop, test and deploy and AI enabled critical systems are even more expensive because of the complexity of the algorithms involved. But aside from cost, there are very significant challenges ahead. One such challenge is the creation of new systems of systems due to unanticipated interaction. With critical infrastructure systems, the potential for *ad hoc* connection to insecure systems mean that a significant amount of effort needs to be invested in planning use and misuse cases [Laplante].

Another challenge is to determine an appropriate software architecture. There are many standard architectures to consider. For example, PROMISE provides a "secure infrastructure for the exchange and processing of life cycle management data throughout all life cycle phases, places a particular emphasis on improving the accessibility and usability of life cycle data during the middle-of-life and end-of-life phases" [PROMISE].

Complex legacy hardware and software architectures can make the implementation of even the "simplest" AI enabled feature difficult. Legacy systems that were not built for interoperation with other systems, or which contain old security vulnerabilities, can create system integration problems. AI enabled features for critical infrastructure must be introduced very carefully, beginning with simple experiments. But this slow move to AI can create impatience and frustration amongst sponsors, public officials and users, creating new pressures and risks in the rush to deployment.

Finally, recent disasters, such as several Boeing 737 Max crashes may serve to curtail public acceptance of AI technologies in critical infrastructure [Hawkins]. Thus, the future needs of AI in critical systems will require not only significant advances in hardware and software technology (hardware, software, legal frameworks and human factors) but verified trust along numerous dimensions. The US National Standards Institute (NIST) identified 18 trust-related concerns for any 'network of things', and many of these concerns apply to AI for critical systems including: control and ownership; composability, interoperability, integration and compatibility; specifications and requirements; synchronization; predictability; testing and assurance approaches; certification; security; reliability; data integrity and insurability and risk measurement [Laplante2]. In order to foster such trust, AI for critical systems will require trusted and qualified (possibly licensed) individuals to build these systems.

The use of AI in critical infrastructure will require new design methodologies tailored to systems-of-systems and will demand broad flexible standards on an unprecedented scale. Across the many domains of critical systems there are many applicable standards and their effectiveness depends on accurately identifying: all relevant standards in an ever-changing domain space, all systems interactions. It also assumes compliance by all parties and dealing with the standards harmonization problem across these interactions.

The new IEEE P7009 standard for "Fail-Safe Design of Autonomous and Semi-Autonomous Systems" establishes a "practical, technical baseline of specific methodologies and tools for the development, implementation, and use of effective fail-safe mechanisms in autonomous and semi-autonomous systems." The standard is expected to set forth "clear procedures for measuring, testing, and certifying a system's ability to fail safely on a scale from weak to strong, and instructions for improvement in the case of unsatisfactory performance" [P7009]. But also, domain specific standards, regulatory compliance, environmental, industry standards and more that have to be harmonized.

Smart infrastructures have greatly increased attack surfaces due to insecure connected devices or those only equipped with conventional security measures. However, AI based computer security can significantly improve the systems' resistance to attack, its ability to recover from the attack and the likelihood of identifying the attackers [Confluence].

Risks to Prediction

The main risks to rapid AI deployment are slow realization of benefits and societal and regulatory pushback. Exaggeration and falsification of real capabilities and risks based on science fiction movies and books, social media posts, can create unwarranted fear and uncertainty. Fictional depictions of sentient and self-aware AI systems gone awry (for example, HAL from the movie *2001, A Space Odyssey*) give the public an unrealistic perception of AI. We may be one tragic disaster away from calls to severely restrict the use of AI in critical systems.

Another problem will be finding suitable training data such as biometrics, behavioral information, patterns of use (e.g. for utilities) for these enhanced AI systems. Privacy issues, for example, involving the training data (behavioral patterns, facial recognition, and other biometrics) needed to make the AI work could thwart progress and deployment.

Finally, legacy systems integrations problems, and standards overload, and confusion may slow progress. AI for critical systems will also require focused coordination between industry and regulatory authorities. And there will need to be increased attention to government investment, community, university and industry partnerships and professional responsibility and assurance that those who build these systems are trusted in order for the true potential to be realized [Mynatt].

Summary

We are at an inflection point where the public may be ready to accept fully autonomous critical systems and not just semi-autonomous (human supervised) ones. AI applications will permeate across all areas of critical infrastructure systems, bringing significant benefit including enhanced use of limited resources, reduced injury and fatality due to accident. These benefits will be particularly noticeable in rural and underserved communities.

Once challenges with public acceptance, various technologies and ad hoc interactions of systems are overcome more ambitious applications will go live within the next five years. However, within next year or two at most, we expect some initial deployments that will trace the path for initial setting of regulatory compliance.

References

[Breniker] Kirk Bresniker, Ada Gavrilovska, James Holt, Dejan S. Milojicic, and Trung Tran, “Grand Challenge: Applying Artificial Intelligence and Machine Learning to Cybersecurity,” to appear in IEEE Computer, December 2019.

[Confluence] Artificial Intelligence and Machine Learning Applied to Cybersecurity, The result of an intensive three-day IEEE Confluence, 6-8 October 2017, https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/industry/ieee_confluence_report.pdf?utm_source=lp-link-text&utm_medium=industry&utm_campaign=confluence-paper

[Daniels] Jeff Daniels, Saman Sargolzaei, Arman Sargolzaei, Tareq Ahram, Phillip A. Laplante, Ben Amaba, "The Internet of Things, Artificial Intelligence, Blockchain, and Professionalism," IT Professional, Nov./Dec., 2018, pp. 15-19.

[Frachtenberg] Eitan Frachtenberg, “Practical Drone Delivery,” to appear in IEEE Computer, December 2019, special issue on Technology Predictions.

[Hawkins] Andrew J. Hawkins, “Deadly Boeing Crashes Raise Questions About Airplane Automation,” *The Verge*, https://www.theverge.com/2019/3/15/18267365/boeing-737-max-8-crash-autopilot-automation, March 15, 2019.

[Janasak] Keith M. Janasak and Raymond R. Beshears. "Diagnostics to Prognostics-A product availability technology evolution." In *Reliability and Maintainability Symposium*, 2007, pp. 113-118.[Khan] Shahzad Khan, Devashish Paul, Parham Momtahan and Moayad Aloqaily. “Artificial Intelligence Framework for Smart City Microgrids: State of the art, Challenges, and Opportunities.” In *2018 Third International Conference on Fog and Mobile Edge Computing (FMEC)*, pp. 283-288.

[Lange] Danny Lange, “Cognitive Robotics,” to appear in IEEE Computer, December 2019.

[Laplante2] Phillip Laplante and Sophia Applebaum. "NIST's 18 Internet of Things Trust Concerns." *Computer*, vol. 52, No. 6, 2019, pp. 73-76.

[Laplante] Phil Laplante, Smarter Roads and Highways, IoT Magazine, vol. 1, no. 2, December 2018, pp. 7-13.

[Moteff] J. D. Moteff and P. Parfomak , Critical Infrastructure and Key Assets: Definition and Identification, Congressional Research Service, Library of Congress, 2004.

[Mynatt] Elizabeth Mynatt, Jennifer Clark, Greg Hager, Dan Lopresti, Greg Morrisett, Klara Nahrstedt, George Pappas et al. "A national research agenda for intelligent infrastructure." *arXiv preprint arXiv:1705.01920* (2017).

[P7009] IEEE Standard P7009, Fail-Safe Design of Autonomous and Semi-Autonomous Systems ***http://sites.ieee.org/sagroups-7009/***.

[Pretz] Kathy Pretz, “Combating the Opioid Crisis, One Flush at a Time,” THE INSTITUTE, IEEE Spectrum, 03 Jul 2019.

[PROMISE] “The Information Exchange for Closed Loop Life cycle Management: The PROMISE Architecture,” http://cl2m.com/wiki/114.

[Qi] Yuan Qi and Jing Xiao. "Fintech: AI powers financial services to improve people's lives." *Communications of the ACM* 61.11 (2018): 65-69.

[Wong] Eric Wong, Xuelin Li and Phillip Laplante, "Be more familiar with our enemies and pave the way forward: A review of the roles bugs played in software failures," Journal of Systems and Software, vol. 133, October 2017, pp.68-94.

[Zhang] Yilu Zhang, Gary W. Gantt, Mark J. Rychlinski, Ryan M. Edwards, John J. Correia, and Calvin E. Wolf, “Connected vehicle diagnostics and prognostics, concept, and initial practice," *IEEE Transactions on Reliability*, vol. 58, no. 2, 2009, pp. 286-294.

BIOs

PHIL LAPLANTE is a Professor of Software and Systems at Penn State. His research interests include the Internet of Things, mission critical systems and project management. Laplante received a PhD from Stevens Institute of Technology. Contact him at plaplante@psu.edu.

DEJAN MILOJICIC is a distinguished technologist at Hewlett Packard Labs. His research interests include OSes, distributed systems, security, and systems management. Milojicic received a PhD from University of Kaiserslautern. Contact him at dejan.milojicic@hpe.com.

SERGEY SEREBRYAKOV is as senior research engineer at Hewlett Packard Labs. His research interests include machine learning, deep learning and its applications. Sergey received a PhD from Saint-Petersburg Institute of Informatics and Automation (SPIIRAS). Contact him at serge.serebryakov@hpe.com.

DANIEL BENNETT is a recently retired U.S. Army Colonel and senior technical advisor at NREL. His research interests include network infrastructure and engineering as well as wireless communications. His PhD is from the University of Colorado at Boulder. Contact him at Daniel.Bennett@nrel.gov.