Top Level Newsletter: Connected Vehicle

(Published by IEEE Orange County Section)

October 2023
Volume 44.0

Editor-in-Chief: Kay Das, IEEE Life Member,
Systems Research Development and Innovation

Vol 44: The first article includes very critical and differing viewpoints on the marriage of connected vehicles with AI (Artificial Intelligence). The second article introduces reconfigurable intelligent surfaces, a new approach that 6G is slated to develop further. The article titles are:

1. Five Conclusions from an Automation Expert with Firsthand Knowledge of Highway Regulation and What Self-Driving Cars Tell Us About AI Risks.

2. Integrated Sensing and Communications With Reconfigurable Intelligent Surfaces: Integration of Signal Modeling and Processing

General Notes
This series of newsletters is intended to provide the IEEE member with a top level briefing of the many different subjects relevant to the research, development and innovation of the connected vehicle.

The objective is to provide a platform for fast learning and quick overview so that the reader may be guided to the next levels of detail and gain insight into correlations between the entries to enable growth of the technology. Intended audiences are those that desire a quick introduction to the subject and who may wish to take it further and deepen their knowledge. This includes those in industry, academia or government and the public at large. Descriptions will include a range of flavors from technical detail to broad industry and administrative issues. A (soft) limit of 300 to 600 words is usually set for each entry, but not rigorously exercised.

As descriptions are not exhaustive, hyperlinks are occasionally provided to give the reader a first means of delving into the next level of detail. The reader is encouraged to develop a first level understanding of the topic in view. The emphasis is on brief, clear and contained text. There will be no diagrams in order to keep the publication concise and podcast - friendly. Related topics in the case of Connected Vehicle technology, such as 5G cellular and the Internet of Things will be included. The terms Connected Vehicle and Automated Driving will be used inter-changeably. Articles from other published sources than IEEE that add to the information value will occasionally be included.

This newsletter forms part of the regional Advanced Technology Initiative (ATI) of which connected vehicles form a constituent part. Technical articles solely from IEEE journals/magazines are referred to by their Digital Object Identifier (DOI) or corresponding https link. The link for each article is provided. Those readers who wish to delve further to the complete paper and have access to IEEE Explore (www.ieeexplore.ieee.org) may download complete articles of interest. Those who subscribe to the relevant IEEE society and receive the journal may already have physical or electronic copies.

In case of difficulty please contact the editor at kaydas@mac.com. The objective is to provide top level guidance on the subject of interest. As this is a collection of summaries of already published articles and serves to further widen audiences for the benefit of each publication, no copyright issues are foreseen.

Readers are encouraged to develop their own onward sources of information, discover and draw inferences, join the dots, and further develop the technology. Entries in the newsletter are normally either editorials or summaries or abstracts of articles. Where a deepening of knowledge is desired, reading the full article is recommended.

Article 1. Five Conclusions from an Automation Expert with Firsthand Knowledge of Highway Regulation: What Self-Driving Cars Tell Us About AI Risks (by Mary Cummings)

Published in: IEEE Spectrum, (Volume: 60, Issue: 10, October 2023)
Page(s): 32-35.

The lack of technical comprehension across industry and government is appalling. People do not understand that the AI that runs vehicles—both the cars that operate in actual self-driving modes and the much larger number of cars offering advanced driving-assistance systems (ADAS)—are based on the same principles as ChatGPT and other large language models (LLMs). These systems control a car's lateral and longitudinal position—to change lanes, brake, and accelerate—without waiting for orders to come from the person sitting behind the wheel.

Both kinds of AI use statistical reasoning to guess what the next word or phrase or steering input should be, heavily weighting the calculation with recently used words or actions. Go to your Google search window and type in “now is the time” and you will get the result “now is the time for all good men.” And when your car detects an object on the road ahead, even if it's just a shadow, watch the car's self-driving module suddenly brake.

Neither the AI in LLMs nor the one in autonomous cars can “understand” the situation, the context, or any unobserved factors that a person would consider in a similar situation. The difference is that while a language model may give you nonsense, a self-driving car can kill you.

Section 1: Human Errors in Operation Get Replaced by Human Errors in Coding
Proponents of autonomous vehicles routinely assert that the sooner we get rid of drivers, the safer we will all be on roads. They cite the NHTSA statistic that 94 percent of accidents are caused by human drivers. But this statistic is taken out of context and inaccurate. As the NHTSA itself noted in that report, the driver's error was “the last event in the crash causal chain…. It is not intended to be interpreted as the cause of the crash.” In other words, there were many other possible causes as well, such as poor lighting and bad road design

Moreover, the claim that autonomous cars will be safer than those driven by humans ignores what anyone who has ever worked in software development knows all too well: that software code is incredibly error-prone, and the problem only grows as the systems become more complex.

Consider these recent crashes in which faulty software was to blame. There was the October 2021 crash of a Pony.ai driverless car into a sign, the April 2022 crash of a TuSimple tractor trailer into a concrete barrier, the June 2022 crash of a Cruise robotaxi that suddenly stopped while making a left turn, and the March 2023 crash of another Cruise car that rear-ended a bus.

These and many other episodes make clear that AI has not ended the role of human error in road accidents. That role has merely shifted from the end of a chain of events to the beginning-to the coding of the AI itself. Because such errors are latent, they are far harder to mitigate. Testing, both in simulation but predominantly in the real world, is the key to reducing the chance of such errors, especially in safety-critical systems. However, without sufficient government regulation and clear industry standards, autonomous-vehicle companies will cut corners in order to get their products to market quickly.

SECTION 2: AI Failure Modes are Hard to Predict

A large language model guesses which words and phrases are coming next by consulting an archive assembled during training from preexisting data. A self-driving module interprets the scene and decides how to get around obstacles by making similar guesses, based on a database of labeled images— this is a car, this is a pedestrian, this is a tree—also provided during training. But not every possibility can be modeled, and so the myriad failure modes are extremely hard to predict. All things being equal, a self-driving car can behave very differently on the same stretch of road at different times of the day, possibly due to varying sun angles. And anyone who has experimented with an LLM and changed just the order of words in a prompt will immediately see a difference in the system's replies.

One failure mode not previously anticipated is phantom braking. For no obvious reason, a self-driving car will suddenly brake hard, perhaps causing a rear-end collision with the vehicle just behind it and other vehicles farther back. Phantom braking has been seen in the self-driving cars of many different manufacturers and in ADAS-equipped cars as well.

The cause of such events is still a mystery. Experts initially attributed it to human drivers following the self-driving car too closely (often accompanying their assessments by citing the misleading 94 percent statistic about driver error). However, an increasing number of these crashes have been reported to NHTSA. In May 2022, for instance, the NHTSA sent a letter to Tesla noting that the agency had received 758 complaints about phantom braking in Model 3 and Model Y cars. This past May, the German publication *Handelsblatt* reported on 1,500 complaints of braking issues with Tesla vehicles, as well as 2,400 complaints of sudden acceleration. It now appears that self-driving cars experience roughly twice the rate of rear-end collisions as do cars driven by people.

Clearly, AI is not performing as it should. Moreover, this is not just one company's problem-all car companies that are leveraging computer vision and AI are susceptible to this problem.

As other kinds of AI begin to infiltrate society, it is imperative for standards bodies and regulators to understand that AI failure modes will not follow a predictable path. They should also be wary of the car companies' propensity to excuse away bad tech behavior and to blame humans for abuse or misuse of the AI.

SECTION 3: Probabilistic Estimates Do Not Approximate Judgment Under Uncertainty

Ten years ago, there was significant hand-wringing over the rise of IBM'sAI-based Watson, a precursor to today's LLMs. People feared AI would very soon cause massive job losses, especially in the medical field. Meanwhile, some AI experts said we should stop training radiologists.

These fears didn't materialize. While Watson could be good at making guesses, it had no real knowledge, especially when it came to making judgments under uncertainty and deciding on an action based on imperfect information. Today's LLMs are no different: The underlying models simply cannot cope with a lack of information and do not have the ability to assess whether their estimates are even good enough in this context.

These problems are routinely seen in the self-driving world. The June 2022 accident involving a Cruise robotaxi happened when the car decided to make an aggressive left turn between two cars. As the car safety expert Michael Woon detailed in a report on the accident, the car correctly chose a feasible path, but then halfway through the turn, it slammed on its brakes and stopped in the middle of the intersection. It had guessed that an oncoming car in the right lane was going to turn, even though a turn was not physically possible at the speed the car was traveling. The uncertainty confused the Cruise, and it made the worst possible decision. The oncoming car, a Prius, was not turning, and it plowed into the Cruise, injuring passengers in both cars.

Cruise vehicles have also had many problematic interactions with first responders, who by default operate in areas of significant uncertainty. These encounters have included Cruise cars traveling through active firefighting and rescue scenes and driving over downed power lines. In one incident, a firefighter had to knock the window out of the Cruise car to remove it from the scene. Waymo, Cruise's main rival in the robotaxi business, has experienced similar problems.

SECTION 4: Maintaining AI is Just as Important as Creating AI

Because neural networks can only be effective if they are trained on significant amounts of relevant data, the quality of the data is paramount. But such training is not a one-and-done scenario: Models cannot be trained and then sent off to perform well forever after. In dynamic settings like driving, models must be constantly updated to reflect new types of cars, bikes, and scooters, construction zones, traffic patterns, and so on.

In the March 2023 accident, in which a Cruise car hit the back of an articulated bus, experts were surprised, as many believed such accidents were nearly impossible for a system that carries lidar, radar, and computer vision. Cruise attributed the accident to a faulty model that had guessed where the back of the bus would be based on the dimensions of a normal bus; additionally, the model rejected the lidar data that correctly detected the bus.

This example highlights the importance of maintaining the currency of AI models. “Model drift” is a known problem in AI, and it occurs when relationships between input and output data change over time. For example, if a self-driving car fleet operates in one city with one kind of bus, and then the fleet moves to another city with different bus types, the underlying model of bus detection will likely drift, which could lead to serious consequences.

Such drift affects AI working not only in transportation but in any field where new results continually change our understanding of the world. This means that large language models can't learn a new phenomenon until it has lost the edge of its novelty and is appearing often enough to be incorporated into the dataset. Maintaining model currency is just one of many ways that AI requires periodic maintenance, and any discussion of AI regulation in the future must address this critical aspect.

SECTION 5: AI Has System-Level Implications that Can't be Ignored

Self-driving cars have been designed to stop cold the moment they can no longer reason and no longer resolve uncertainty. This is an important safety feature. But as Cruise, Tesla, and Waymo have demonstrated, managing such stops poses an unexpected challenge.

A stopped car can block roads and intersections, sometimes for hours, throttling traffic and keeping out first-response vehicles. Companies have instituted remote-monitoring centers and rapid-action teams to mitigate such congestion and confusion, but at least in San Francisco, where hundreds of self-driving cars are on the road, city officials have questioned the quality of their responses.

Self-driving cars rely on wireless connectivity to maintain their road awareness, but what happens when that connectivity drops? One driver found out the hard way when his car became entrapped in a knot of 20 Cruise vehicles that had lost connection to the remote-operations center and caused a massive traffic iam.

Of course, any new technology may be expected to suffer from growing pains, but if these pains become serious enough, they will erode public trust and support. Sentiment toward self-driving cars used to be optimistic in tech-friendly San Francisco, but now it has taken a negative turn due to the sheer volume of problems the city is experiencing. Such sentiments may eventually lead to public rejection of the technology if a stopped autonomous vehicle causes the death of a person who was prevented from getting to the hospital in time.

***Software code is incredibly error-prone, and the problem only grows as the systems become more complex*!**

So what does the experience of self-driving cars say about regulating AI more generally? Companies not only need to ensure they understand the broader systems-level implications of AI, they also need oversight—they should not be left to police them-selves. Regulatory agencies must work to define reasonable operating boundaries for systems that use AI, and issue permits and regulations accordingly. When the use of AI presents clear safety risks, agencies should not defer to industry for solutions and should be proactive in setting limits.

AI still has a long way to go in cars and trucks. I'm not calling for a ban on autonomous vehicles. There are clear advantages to using AI, and it is irresponsible for people to call for a ban, or even a pause, on AI. But we need more government over-sight to prevent the taking of unnecessary risks.

And yet the regulation of AI in vehicles isn't happening yet. That can be blamed in part on industry overclaims and pressure, but also on a lack of capability on the part of regulators. The European Union has been more proactive about regulating artificial intelligence in general and in self-driving cars particularly. In the United States, we simply do not have enough people in federal and state departments of transportation that understand the technology deeply enough to advocate effectively for balanced public policies and regulations. The same is true for other types of AI.

This is not anyone administration's problem. Not only does AI cut across party lines, it cuts across all agencies and at all levels of government. The Department of Defense, Department of Homeland Security, and other government bodies all suffer from a workforce that does not have the technical competence needed to effectively oversee advanced technologies, especially rapidly evolving AI.

To engage in effective discussion about the regulation of AI, everyone at the table needs to have technical competence in AI. Right now, these discussions are greatly influenced by industry (which has a clear conflict of interest) or Chicken Littles who claim machines have achieved the ability to outsmart humans. Until government agencies have people with the skills to understand the critical strengths and weaknesses of AI, conversations about regulation will see very little meaningful progress.

Recruiting such people can be easily done. Improve pay and bonus structures, embed government personnel in university labs, reward professors for serving in the government, provide advanced certificate and degree programs in AI for all levels of government personnel, and offer scholarships for undergraduates who agree to serve in the government for a few years after graduation. Moreover, to better educate the public, college classes that teach AI topics should be free.

We need less hysteria and more education so that people can understand the promises but also the realities of AI. (2797 words)

DOI: https://doi.org/10.1109/MSPEC.2023.10271347

Article 2. Integrated Sensing and Communications With Reconfigurable Intelligent Surfaces: the Integration of Signal Modeling and Processing (by Sundeep Chepuri et al)

Published in: IEEE Signal Processing Magazine , (Volume: 40, Issue: 6, September 2023)
Page(s): 41-62.

Abstract
Integrated sensing and communications (ISAC) is envisioned to be an integral part of future wireless networks, especially when operating at the millimeter-wave (mm-wave) and terahertz (THz) frequency bands. However, establishing wireless connections at these high frequencies is quite challenging, mainly due to the penetrating path loss that prevents reliable communication and sensing. Another emerging technology for next-generation wireless systems is reconfigurable intelligent surface (RIS), which refers to hardware-efficient planar structures capable of modifying harsh propagation environments. In this article, we provide a tutorial-style overview of the applications and benefits of RISs for sensing functionalities in general, and for ISAC systems in particular.

We highlight the potential advantages when fusing these two emerging technologies, and identify for the first time that 1) joint sensing and communications (S&C) designs are most beneficial when the channels referring to these operations are coupled, and that 2) RISs offer the means for controlling this beneficial coupling. The usefulness of RIS-aided ISAC goes beyond the obvious individual gains of each of these technologies in both performance and power efficiency. We also discuss the main signal processing challenges and future research directions that arise from the fusion of these two emerging technologies.

Coverage of Article
Recent years have witnessed growing research and industrial attention in RISs. An RIS is an array of elements with programmable scattering properties. Although this definition accommodates a broad range of technologies, the common treatment of RISs considers 2D arrays whose elements can be tuned independently to generate desirable reflection patterns in a nearly passive fashion, without utilizing active radio-frequency (RF) chains to process the impinging signals. The core benefits of RISs stem from their ability to shape the propagation profile of information-bearing electromagnetic (EM) waves in a flexible, low-cost, and energy-efficient manner.

One of the most common applications of RISs is in wireless communications. This metamaterial-based technology brings forth the vision of smart programmable environments, where RISs are expected to improve coverage, energy efficiency, reliability, and EM field exposure of wireless communications. The role of RISs in future wireless communications in modifying harsh propagation environments and establishing reliable links for communication via multiple-input multiple-output (MIMO) systems is widely studied in the literature. Consequently, they are expected to play a key role in 5G-Advanced as well as in 6G wireless networks.

Another application area where the use of RISs is the focus of considerable research attention is RF localization. The conventional task in such systems considers a mobile device that determines its position based on the impinging signals received from several terminals whose locations are known. The commonly studied capability of RISs in facilitating RF localization follows from their reliance on forming multiple signals received from known locations. RISs can thus create additional signal propagation paths for their impinging signals without increasing the number of transmitting terminals as well as overcome nonline-of-sight (NLoS) conditions in an energy-efficient manner. The growing popularity of RIS-enabled/-aided RF localization has also been attributed to the fact that this application area is naturally associated with wireless communications devices, where RIS research has recently been highly established.

The article focuses on another emerging technology: the ISAC paradigm. This technology unifies wireless communications and RF sensing as both applications are associated with radiating EM waves, and are expected to be simultaneously employed by a multitude of mobile devices. As such, ISAC is envisioned to be an integral part of future wireless systems, especially when operating at the mm-wave and THz frequency bands. However, signal propagation at these high frequencies is quite challenging due to the penetrating path loss, which can be so severe that the NLoS paths may be too weak to be of practical use, preventing reliable communication and, in certain cases, sensing. ISAC systems share many similar aspects, in both operation and challenges, with the more established domains of wireless communications and RF localization.

In particular, data communication is one of the two functionalities incorporated in ISAC systems. Yet, although RF localization can be viewed as a form of sensing, the way the propagation of EM waves is used typically differs from that in ISAC systems. In particular, ISAC commonly focuses on a wireless receiver (Rx) (or transceiver) utilizing impinging RF signals to sense an unknown environment, similar to what happens in radar systems. Nonetheless, the ability of RISs to generate controllable, desired reflection patterns was recently shown to facilitate sensing applications. This indicates that the combination of RISs with ISAC systems may contribute to the development of future wireless networks and their mobile devices.

We review the applications of RISs for ISAC systems as well as the signal processing challenges and the offered gains that arise from the integration of these two emerging technologies. We present up-to-date research directions of such synergies, contribute new fundamental insights, and highlight exciting research opportunities arising from the use of the RISs. For this goal, we begin by reviewing the fundamentals of the RIS technology that are relevant for ISAC applications. We briefly describe RIS hardware architectures, focusing mostly on conventional passive architectures as well as recent hybrid passive-active technologies.

We then discuss the key capabilities of RISs in shaping wireless environments, and their associated use cases, based on which we establish a generic model for the signals involved in RIS-aided ISAC systems that we specialize in the subsequent sections. We continue by overviewing the existing methodologies for utilizing RISs for sensing applications, focusing specifically on radar sensing, which has the closest functionality to ISAC systems.

The article concludes with a discussion on the open challenges and future research directions associated with RIS-empowered ISAC systems. Exploring these directions is expected to pave the way to a successful fusion of the concepts of smart wireless environments and ISAC, further bringing forth the gains of joint holistic designs of RISs and ISAC. (961 words)

DOI: https://doi.org/10.1109/MSP.2023.3279986