Top Level Newsletter: Connected Vehicle

(Published by IEEE Orange County Section)

February 2020
Vol 5.1

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

Concepts (in alphabetical order)

Artificial Intelligence (AI) is the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience. Since the development of the digital computer in the 1940s, it has been demonstrated that computers can be programmed to carry out very complex tasks—as, for example, discovering proofs for mathematical theorems or playing chess—with great proficiency.

Big data has spawned the current interest and increased investment in artificial intelligence. The availability of large volumes of data—plus new algorithms and more computing power—are behind the recent success of deep learning, finally pulling AI out of its long "winter." More broadly, the enthusiasm around big data (and the success of data-driven digital natives such as Google and Facebook), has led many enterprises to invest heavily in the collection, storage, and organization of data. (170)

Building Blocks of Connected Vehicle Technology: Mimicking the many things humans do while driving requires a complex blend of technologies. Several sensors are required to analyse inputs and make decisions about steering, braking, speed…and routing. OEMs, Tier1s, start-ups etc are all involved in this disruption. Key areas are (47):

- *Processing Power*: a range of technologies from NXP. Infineon, Renesas, STMicroelectronics, Intel etc. Joined by (1) Nvidia with its GPUs which excel at performing multiple tasks simultaneously, like analyzing the many pixels streaming in from cameras. Nvidia's recent platform for SAE level 3 to 5 is the size of a car license plate delivering 320Moperations per sec. (2) Mobileye/Intel with its dedicated video processor EyeQ3. Conventional CPUs will not disappear as the need to make sequential decisions remains. Specialised parallel devices can be made from FPGAs from Xylinx, Altera/Intel, Lattice and Microsemi. The fusion of sensors and localized processors will continue. The likely progression will be that a major control unit will collect all relevant data and make decisions. That will push the demand for more cores, lower latency/higher clock rates, and lower power consumption. Another development is driver alertness monitoring for SAE 2 upwards. The driver's face angle is perceived to determine drowsiness. (155)

- *Sensors*: Radar, Video, Lidar: Humans need two eyes and some mirrors to drive well but vehicles will need upto 30 sensors to match. The radar component is dominated by NXP, STMicroelectronics, Analog Devices and Renesas who are pairing controllers and radar. Start-ups/others include Oculii, Omniradar, Artsys360, Bosch, Delphi. Cameras and radar currently work together and overlap to identify objects and it is likely that lidar will be added when solid state devices meet automotive requirements. The shift to 77 GHz radar from 24 GHz permits more bandwidth. Advanced techniques for sending and receiving signals are helping radar identify objects instead of simply providing range information. In video, 8-10 Mpixels will supersede today's 2-4 Mpixel cameras. Increased data will require innovative latency management.
Various smaller companies now make electro-mechanical lidar systems but all are developing solid state versions. Some lidar systems are ready but only for short distance usage like lane departure. High resolution at distances of 200 to 300 ft is required. (162)

- *Architecture and SW/HW interfaces*: The architecture is of prime importance in achieving quality of performance. Amongst many responsibilities, it determines how sensor data is collected and fused to create a single and dynamic view of the surroundings, how data is shared throughout the vehicle, how decisions are made, and how efficiently data is cross-checked. Software will play a key role as electronic controls determine actions in reponse to what is happening around the vehicle. Hardware must be powerful enough to make computations in time to avoid danger. Today software and hardware are often made by the same supplier, though the Automotive Open System Architecture (AUTOSAR) has enabled some separation. However, a continued necessity for abstraction between the two prevails in order to bring in third party software. One critical aspect is whether the processing power needs to be centralized or distributed (ADAS favours the latter). There is evenly divided opinion on this. Some OEMs opt for a mix of the two approaches. It is difficult to create software that will respond correctly to all the possibilities that autonomous vehicles will see while in operation. This is sparking a huge investment in creating AI programs that 'learn' as vehicles are in operation. (202)

- *Communications:* If vehicles can communicate, they can get information their onboard sensors cannot get, such as emergency braking, blind spot activity etc. V2I communications to roadside beacons can also aid in safety and in traffic flow. The NHSTA is advocating the use of Dedicated Short Range Communication (DSRC), though the uptake has been slow in some sectors of the industry. DSRC promises low latency. There are some spectrum issues to be resolved. Cadillac has been the first to adopt V2V. On highway tests have already proven DSRC's performance. While some 5G proponents feel that cellular, already embedded in the industry, can displace DSRC, many feel that both systems can share roles and co-exist. It will take time for V2X vehicles to displace older cars. Aftermarket systems cannot offer many safety features since it is difficult to verify that messages are coming from a verified source and not a hacker! (151)

Caching: a hardware or software component that stores data so that future requests for that data can be served faster; the data stored in a cache might be the result of an earlier computation or a copy of data stored elsewhere. A *cache hit* occurs when the requested data can be found in a cache, while a *cache miss* occurs when it cannot. Cache hits are served by reading data from the cache, which is faster than recomputing a result or reading from a slower data store; thus, the more requests that can be served from the cache, the faster the system performs.
To be cost-effective and to enable efficient use of data, caches must be relatively small. Nevertheless, caches have proven themselves in many areas of computing, because typical computer applications access data with a high degree of locality of reference. Such access patterns exhibit temporal locality, where data is requested that has been recently requested already, and spatial locality, where data is requested that is stored physically close to data that has already been requested. (177)

Cloud Computing is the on-demand availability of computer system resources, especially data storage and computing power, without direct active management by the user. The term is generally used to describe data centers available to many users over the Internet. Large clouds, predominant today, often have functions distributed over multiple locations from central servers. If the connection to the user is relatively close, it may be designated an edge server. Clouds may be limited to a single organization (enterprise clouds) or be available to many organizations (public cloud,) or a combination of both. (92)

Computer Vision is an interdisciplinary field that deals with how computers can be made to gain high-level understanding from digital images or videos. Computer vision is concerned with modeling and replicating human vision using computer software and hardware. It is a discipline that studies how to reconstruct, interrupt and understand a 3D scene from its 2D images in terms of the properties of the structure present in scene.

From the perspective of engineering, it seeks to automate tasks that the human visual system can do. Computer vision is concerned with the automatic extraction, analysis and understanding of useful information from a single image or a sequence of images. It involves the development of a theoretical and algorithmic basis to achieve automatic visual understanding. As a scientific discipline, computer vision is concerned with the theory behind artificial systems that extract information from images. The image data can take many forms, such as video sequences, views from multiple cameras, or multi-dimensional data from a medical scanner. (164)

Connected Vehicle: Although many terms are in use such as Autonomous Vehicle, Connected Car, and others, we choose this nomenclature to include partially and fully autonomous vehicles of all descriptions, for example, including two-wheel machines.

Co-ordinated Multipoint (CoMP) When a UE is in the cell-edge region, it may be able to receive DL signals from multiple cell sites and the UE's UL transmission may be received at multiple cell sites regardless of the system load. Given that, if the signaling transmitted from the multiple cell sites is coordinated, the DL performance can be increased significantly. This coordination can be simple as in the techniques that focus on interference avoidance or more complex as in the case where the same data is transmitted from multiple cell sites. For the UL, since the signal can be received by multiple eNBs, if the scheduling is coordinated from the different cell sites, the system can take advantage of multiple reception to significantly improve the link performance. A UE at the edge of a cell is able to be served by two or more eNBs to improve reception / transmission and increase throughput particularly under cell edge conditions.

An advantage of intra-site CoMP is that significant amount of exchange of information is possible since this communication is within a site and does not involve the backhaul (connection between base stations). Inter-site CoMP involves the coordination of multiple sites for CoMP transmission. Consequently, the exchange of information will involve backhaul transport. This type of CoMP may put additional burden and requirement upon the backhaul design. (221)

Dedicated Short Range Communications DSRC allows high-speed communications between vehicles and the roadside, or between vehicles. It has a range of up to 1,000 meters. DSRC applications for public safety and traffic management include intersection collision avoidance and many others. DSRC applications now in use include electronic toll collection, and electronic credentialing and monitoring of commercial vehicle operations. Current applications operate at 915 MHz utilizing a bandwidth of 12 MHz and primarily use proprietary technology,although some standards- compliant devices have been developed. The new 5.9 GHz frequency permits much higher data-transmission rates than the lower-frequency 915 MHz bandand provides 75 megahertz of spectrum for DSRC applications. In the case of 5.9 GHz, other users in the band include military radars and satellite communications systems.

DSRC has low latency and high reliability, is secure, and supports interoperability. It receives very little interference, even in extreme weather conditions, because of the short range that it spans. This makes it ideal for communication to and from fast-moving vehicles. DSRC technology can be used in either a vehicle-to-vehicle (V2V) or vehicle-to-infrastructure (V2I) format, and communicates using transponders known as on-board units (OBUs) or roadside units (RSUs). (192)

Deep Learning: This is a class of machine learning algorithms that use multiple layers to progressively extract higher level features from raw input. For example, in image processing, lower layers may identify edges, while higher layers may identify human-meaningful items such as digits/letters or faces. In deep learning, each level learns to transform its input data into a slightly more abstract and composite representation. In an image recognition application, the raw input may be a matrix of pixels; the first representational layer may abstract the pixels and encode edges; the second layer may compose and encode arrangements of edges; the third layer may encode a nose and eyes; and the fourth layer may recognize that the image contains a face. Importantly, a deep learning process can learn which features to optimally place in which level *on its own*.

It is predicted that by 2020, deep learning will have reached a fundamentally different stage of maturity. Deployment and adoption will no longer be confined to experimentation, becoming a core part of day-to-day business operations across most fields of research and industries. This is because advancements in the speed and accuracy of the hardware and software will have made it both viable and cost-effective. Much of this added value will be generated by deep learning inference - that is, using a model to infer something about data it has never seen before. Models can be deployed in the cloud or data center, but more and more we will see them on end devices like cameras and phones. (254)

Device-to-Device (D2D) This mode of communication is expected to play a significant role in upcoming cellular networks as it promises ultra-low latency for communication among users. This new mode may operate in the licensed or unlicensed spectrum bands. It allows UEs in close proximity to communicate using a direct link rather than having their radio signal travel all the way through the base station or the core network. One of its main benefits is ultra-low latency in communication due to a shorter signal traversal path. Various short-range wireless technologies like Bluetooth, WiFi Direct and LTE Direct can be used to enable D2D communication.

In a typical cellular network, UEs achieve time and frequency synchronization using periodic broadcasts from the BS. Devices in D2D communication can also synchronize with the same broadcasts so long as they belong to the same BS. The situation gets complicated in the following cases: (1) UEs belong to different BSs that may not be themselves synchronized, or (2) some of the UEs are in the coverage of the network and some outside the coverage, and (3) all UEs lie outside network coverage [9]. Synchronization among UEs is beneficial for D2D communication because it helps a UE use the right time slot and frequency for discovering and communicating with its peer and thus engage in more energy-efficient communication. Note that *global* synchronization among all UEs in a network may not be required for D2D communication; rather *local* synchronization among neighboring devices is sufficient. (243)

Fog Computing : An architecture that uses edge devices, which provide entry points into enterprise or service provider core networks, to carry out a substantial amount of computation, storage, with local communication and routed over the internet backbone. It can be perceived both in large cloud systems and big data structures, making reference to the growing difficulties in accessing information objectively, resulting in a lack of quality of the obtained content. Fog networking consists of a control plane and a data plane. On the data plane, fog computing enables computing services to reside at the edge of the network as opposed to servers in a data-center. Fog computing emphasizes proximity to end-users and client objectives, dense geographical distribution and local resource pooling, latency reduction, and backbone bandwidth savings. Fog networking supports the Internet of Things (IoT) concept, in which most of the devices used by humans on a daily basis will be connected to each other. (156).

Fronthaul and Backhaul are an essential component of the 5G radio access infrastructure. They provide connectivity and data transport in the 5G radio access network. The fronthaul connects 5G antennas with base stations and the backhaul connects base stations with the operator's core network. The term Xhaul refers to the integrated combination of fronthaul and backhaul. (56)

Internet of Things (IoT) is the extension of Internet connectivity into physical devices and everyday objects. Embedded with electronics and other forms of hardware (such as sensors), these devices can communicate and interact with others over the Internet, and they can be remotely monitored and controlled. The definition of the IoT has evolved due to convergence of multiple technologies such as real-time analytics, machine learning, commodity sensors, and embedded systems. In the consumer market, IoT technology is most synonymous with products pertaining to the concept of the "smart home", covering devices and appliances (such as lighting fixtures, thermostats, home security systems and cameras, and other home appliances) that support one or more common ecosystems, and can be controlled via devices such as smartphones and smart speakers like Alexa from Amazon.(130)

Levels of automation (SAE): The U.S Department of Transportation National Highway Traffic Safety Administration (NHTSA) provided a standard classification system in 2013 which defined five different levels of automation, ranging from level 0 (no automation) to level 4 (full automation). Since then, the NHTSA updated their standards to be in line with the classification system defined by SAE International. SAE International defines six different levels of automation that ranges from 0 (no automation) to 5 (full automation)

Level 0 – No automation. The driver is in complete control of the vehicle and the system does not interfere with driving. Systems that may fall into this category are forward collision warning systems and lane departure warning systems.

Level 1 – Driver assistance. The driver is in control of the vehicle, but the system can modify the speed and steering direction of the vehicle. Systems that may fall into this category are adaptive cruise control and lane keep assist.

Level 2 – Partial automation. The driver must be able to control the vehicle if corrections are needed, but the driver is no longer in control of the speed and steering of the vehicle. Parking assistance is an example of a system that falls into this category along with Tesla's autopilot feature. It is important to note the driver must not be distracted in Level 0 to Level 2 modes.

Level 3 – Conditional automation. The system is in complete control of vehicle functions such as speed, steering, and monitoring the environment under specific conditions. Such specific conditions may be fulfilled while on fenced-off highway with no intersections, limited driving speed, boxed-in driving situation etc. A human driver must be ready to intervene when requested by the system to do so. If the driver does not respond within a predefined time or if a failure occurs in the system, the system needs to do a safety stop in ego lane (no lane change allowed)]. The driver is only allowed to be partially distracted, such as checking text messages, but taking a nap is not allowed.

Level 4 – High automation. The system is in complete control of the vehicle and human presence is no longer needed, but its applications are limited to specific conditions. An example of a system being developed that falls into this category is the Waymo self-driving car service. If the actual motoring condition exceeds the performance boundaries, the system does not have to ask the human to intervene but can choose to abort the trip in a safe manner, e.g. park the car.

Level 5 – Full automation. The system is capable of providing the same aspects of a Level 4, but the system can operate in all driving conditions. The human is equivalent to "cargo" in Level 5 Currently, there are no driving systems at this level.

Machine Learning is the scientific study of algorithms and statistical models that computer systems use in order to perform a specific task effectively without using explicit instructions, relying on patterns and inference instead. It is seen as a subset of artificial intelligence. Machine learning algorithms build a mathematical model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to perform the task. Machine learning algorithms are used in a wide variety of applications, such as computer vision, where it is infeasible to develop an algorithm of specific instructions for performing the task. (102)

Near Far Effect : Consider a receiver and two transmitters, one close to the receiver, the other far away. If both transmitters transmit simultaneously and at equal powers, then due to the inverse square law the receiver will receive more power from the nearer transmitter. Since one transmission's signal is the other's noise, the SNR for the further transmitter is much lower. This makes the farther transmitter more difficult to understand. Possible solutions to the near–far problem:

1. Increased receiver dynamic range - Use a higher resolution ADC. This can increase the dynamic range of receiver stages that are saturating.
2. Dynamic output power control – Nearby transmitters decrease their output power so that all signals arrive at the receiver with similar signal strengths.
3. TDMA – Transmitters use a time division scheme to avoid transmitting at the same time.

(138)

Network Slicing: Use cases such as mobile broadband (MBB), the Internet of Things (IoT), and enterprise—put very different requirements on the network, but you do not want to build separate physical networks for the different services. Network slicing enables you to build multiple logical networks—network slices—on top of a common shared physical infrastructure. These network slices, one for each kind of service, are separate and independent to the extent that if something goes wrong in one slice it will not affect the other slices. This separation and independence also enable you to add new slices without impacting the rest of the network. You can use network slicing for any access technology, and the building blocks for both building and managing a sliced network are available. Network slicing in 5G is expected to open lucrative new business opportunities for mobile operators and other newer entrants. For example, a mobile operator will be able to split its physical network resources into multiple logical slices and lease these slices out to multiple parties. Network slicing is expected to be a key component of 5G networks.

In November 2016, a white paper that explores an end-to-end 5G systems framework for the creation of customised network slices was released. It also considers the application of network slicing to air-interface technologies and the long-term technology roadmap. (221)

Non-Orthogonal Multiple Access (NOMA) is a new multiple access scheme proposed. It is one of the many technologies that promise greater capacity gain and spectral efficiency than the present state of the art, and as such, is a candidate technology for 5G. Each generation of cellular technology is usually characterized by a specific multiple access scheme (FDMA/ TDMA/ WCDMA/ OFDMA). NOMA is a likely candidate in the line of succession for these technologies. NOMA goes against the present trend of transmitting information in orthogonal carriers or subcarriers. Here, the multiple users are multiplexed in the power domain, either in downlink or uplink.

In the downlink version, the Base Station (BS) transmits two signals with different powers in the same frequency band simultaneously. One message will be for a User Equipment (UE) located near the BS, while the other will be intended for a UE located a fair distance away. The UE near the base station will be allotted less power, while the far UE will be allotted more power. The UE located at the high Signal to Noise Ratio (SNR) region first decodes the message intended for the far receiver. This message is then subtracted from the received signal, and we are left with the interference-free signal that the near UE is supposed to decode. This approach requires a Successive Interference Cancellation (SIC) receiver. The non-orthogonally improves resource utilization and is found to increase the user throughput by about 38%. The tradeoff is receiver complexity.

NOMA can be extended by usage in conjunction with Massive MIMO systems. The power allocation to the antennae is based on distance and direction of the user. As a result, beams with possible slight overlaps can be directed towards different users. (A. Mohan, NOMA seminar, Govt Engg College, Trichur). (289)

QUIZ

1. What are the limits (in metres) to wireless localization accuracy of moving vehicles and how can these be improved?
2. What are some of the new techniques for channel coding in 5G compared to 4G/LTE?
3. A century ago what was the big problem with horses on the roads?
4. What is the name of Google's self-driving car unit?
5. What are some of the techniques that can be utilized to minimize latencies of information flow in connected vehicles?
6. How many electronic system sensors does it take to match human performance?
7. How does platooning of connected vehicles make for reduction in system capacity demand?
8. How will network slicing help improve reliability and performance in 5G networks?
9. What is the forecast percentage in fall of urban car ownership by 2050?
10. How does the ISO 26262 standard help in the design of reliable connected vehicle systems?
11. What are the three most effective sensor technologies for automated driving?
12. What are the five levels of automation defined by the SAE?

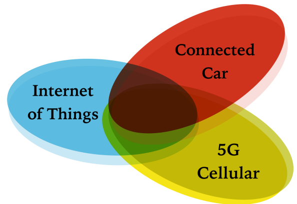
Connected Car
Internet of Things
5G Cellular