Data & Safety: challenges and opportunities

Hugues Bonnin, Olivier Flebus

Data & Safety: challenges and opportunities

Hugues Bonnin, Olivier Flebus
Intelligent Transportation Systems
Continental
Toulouse, France
hugues.bonnin@continental-corporation.com, olivier.flebus@continental-corporation.com

***Abstract*— This article presents our first analysis exploring the relationships of safety and data in the context of improving the safety of road vehicles, which will be connected and more automated in the future.**

Keywords*— *Holistic Safety, Road Traffic Safety, Data, Data-driven approach, Data Engineering, Data Governance, System Engineering, Data vs. Software, Systems of Intelligence.

I. INTRODUCTION

"Safety First" has always been a core principle in transport industries, even if it has been analyzed and concretely declined into different safety approaches and standards in the aeronautics, railway and automotive domains. Several key standards have emerged over time, mainly focusing on how to engineer safe systems [2][4]. With the increased role of software in dependable systems, the question "How to engineer safe software?" has been addressed as well (in most cases from the point of view of the system integrating the software)[2][3].

Data-driven approaches have emerged with the rise of the Internet and the digital transformation in almost every domain. Even if the Big Data real innovations (e.g. easier development and operations of distributed systems at scale, releasing sometimes unnecessary constraints with *eventual consistency*…) have been turned into a hype, there is no doubt that in many domains data is considered as a key asset for the organizations that can collect them. A key asset that can be leveraged as a competitive advantage and a defensibility against competitors.

We propose to address the following basic questions:

- What is (and will be) the role (and the importance) of data in the future automotive applications? For sure data could harm those systems. But could data also bring benefits?
- How to address the rise of data in dependable systems? Are existing safety standards for systems and software ready to integrate the new data-driven stakes?

II. DEFINITIONS

In our analysis, we use the following definitions for Information and Data (according to [1]):

- Information: knowledge concerning objects, such as facts, events, things, processes, or ideas, including concepts, that within a certain context has a particular meaning.
- Data: a reinterpretable representation of information in a formalized manner suitable for communication, interpretation, or processing.

That notably means that in our analysis information concern both humans and machines, while data apply to information that has been digitalized to be managed by machines (including in a human-understandable way).

III. CONTEXT AND ASSUMPTIONS

We list here the assumptions that we have chosen neither to describe with details nor to justify in this short paper. They could of course still be discussed and formalized in a longer version of this article.

(1) The future of mobility (and thus the future of automotive) implies more automated driving and autonomous vehicles – and a key motivation for that is the expected improved Road Traffic Safety.

(2) The future of mobility implies connected vehicles. Beyond single vehicles, connectivity drastically increases the complexity of the involved systems, dealing with interactions between thousands or even millions of vehicles and infrastructures systems.

(3) Beyond the automation and the connectivity, the (digital) cooperation between people, vehicles and infrastructures is key for Intelligent Transportation Systems. With regards to the impacts for safety, it brings unprecedented complexity levels on the challenge side, but also redundancy and other opportunities.

IV. STATE OF THE ART IN DATA SAFETY

We have initiated an analysis of current standards that could contribute to the definition of data safety. Multiple topics are covered, for instance:

- Standards dedicated to the connected vehicles [9] or Intelligent Transportation Systems (ITS) (ISO TC 211, ISO TC 204, CEN TC 278, CEN TC 226)
- Standards dedicated to maps, e.g. from the Open AutoDrive Forum (OADF)
- Standards addressing data quality [6][7][8]
- The data safety guidance [11]
- Data standards used in other domains (i.e. [10])

A very (very!) rough conclusion of this first analysis could be expressed as follows: existing standards for data don't address safety stakes beyond data quality and existing safety standards for automotive (and more generally dependable) systems don't address modern data stakes. For these standards, we identify as a limit the fact that the *data safety* is only considered as a result of its processing by systems, and that nothing is established on the *information* that it contains, or on which it is based.

V. A SIMPLE INFORMATION FLOWS MODEL FOR ROAD TRAFFIC SAFETY

A. Road Traffic Safety is not only about road vehicles

It would be a mistake to consider that vehicles should be the unique way of improving the road traffic safety. As increasing automation necessarily involves that humans do less things, it seems natural to consider the safety as a whole, including the way we currently deal with human faults (using a mobile phone while driving for instance) but also the human intelligence that in some cases is the last and unique way to prevent accidents. More generally Road Traffic Safety depends on multiple actors and information, not only vehicles! For instance:

- Road users: drivers in/with a vehicle, but also pedestrians, animals
- Regulations concerning drivers: driving license, restrictions related to alcohol, drugs or mobile phones.
- Regulations concerning vehicles for their safety, emissions, etc. Attached to mandatory vehicle inspections in many countries.
- Road Infrastructures: road signs, safety barriers, speed bumps, etc.

B. *A very simple Information (data) flows model as a starting point*

In our context of more automated vehicles and holistic road traffic safety, we considered a deliberately simple model to identify the information flows that are involved in driving activities (Figure 1).

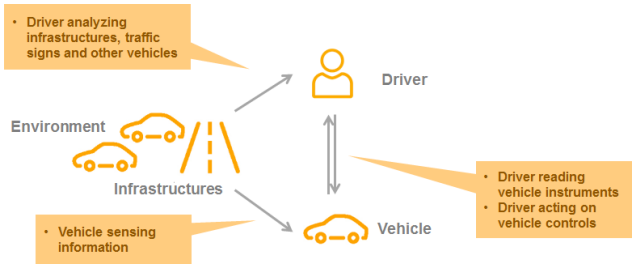

Figure 1: A simple model for Information flows in driving activities

The driver and the vehicle working together form a *man-machine system* that is critical for the safety of the transportation. In his/her driving activities, the driver mainly analyses the environment (including the other vehicles), monitors information produced by his/her vehicle, and acts on the vehicle to adapt his/her trajectory and more generally behavior.

Historically, the vehicle has been helping the driver to sense the environment (with wipers and lights for instance) and to enforce the relevance and the compliance of his/her behavior (comply with speed limits, stop when a fault is detected). More recently, vehicles have developed the ability to sense information from the environment, providing support to the driving activities, and in some cases targeting "suprahuman" sensing capabilities in specific domains (e.g. radar sensing).

C. *Can we separate (human) intelligence from information?*

Exploring the dynamics of the previous model, and still considering that the human driver keeps making the main decisions regarding the driving activities, it appears that there's a whole class of *knowledge* (information) that are involved into the human decision-making process. Example of information that fall into that category are:

- Laws & regulations: common (in the sense of international) and country-specific driving rules (checked with a driving license).
- Information learned by drivers through their driving experience, which could improve the driver's ability to anticipate hazards…and also reduce driving attention and discipline during routine, daily trips.
- All the softer rules, social contracts and human judgments, for instance the ones that regulate heavy congested situations, leaving way to another vehicle even though driving rules don't say so.

D. *What about all the data that are produced and processed in vehicles?*

In a similar way than the fact that there are a lot of pieces information behind the human intelligence, there are a lot of data in vehicles to ensure its expected (and safe) behavior. Using those data for other purposes outside of the vehicle – and complying to laws and regulations like GDPR – is at the core of our motivation for our work.

E. *Mapping Road Traffic Safety to our simple information flow model*

Starting with our simple information flow model, what can we say about the safety questions that arise for each information flow? The Figure 2 gives a few examples:

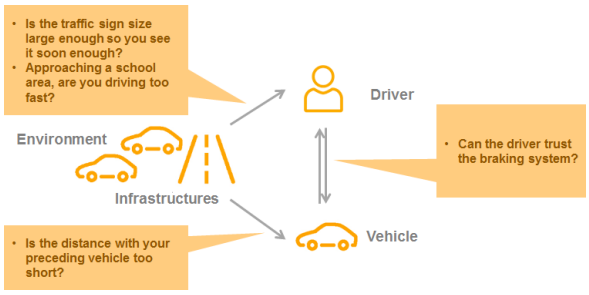

Figure 2: Examples of safety-related questions mapped to information flows

As explained previously, Road Traffic Safety includes a very wide scope of questions in multiple domains (driver, infrastructures, vehicles). What we see also is that a very high diversity on information types is involved, including ones that are - for now - only accessible to humans. That concerns especially soft rules and social contracts. By the way, we can consider the safety itself as a social contract, mixing acceptable safety for the population, and eventual negotiations between stakeholders.

VI. PROPOSED APPROACH & NEXT STEPS

A. *Data outside of Systems*

It can be deduced from assumption (1) that in order to improve safety, a certain amount of information currently managed by the human driver must be handled by the vehicle. However, this transfer is accompanied by a considerable increase in the complexity of the resulting system, as we can guess from "V. A simple information flows model for Road Traffic Safety". Indeed, whether to manage the anticipation of events that will be encountered on the route (see beyond the horizon), or access knowledge about the weather and its consequences on traffic or driving conditions, or driving rules, including temporary or implicit rules of conduct, the car is necessarily connected. This extension of the information to be managed creates complexity, as we have stated in assumption (2).

However, this new complexity is hardly compatible for us with the analysis and design of systems, including system of systems, as it is traditionally done: by defining a system at the highest level, and by breaking it down by refinement to products at lower levels, including equipments inside a vehicle. Moreover: this approach would very quickly meet another limitation: the differences of life cycles between the components of this giga-system, which means that some are operated even before others are defined.

To face up this complexity, we propose to reason directly on the *information*, and to apply a safety-oriented thinking on this flow of information. This inverts the common approach of

system engineering, proposing to build confidence based on *information first*, and *system second*, instead of the contrary.

This results in the data being considered as being outside of the systems.

This point of view is a way to answer to the limit previously highlighted at the end of "IV-State of the art in data safety". By the way, it enables the *decoupling* between systems that produce, consume or manage data, necessary to address the differences of lifecycle stated above.

This approach enables to consider end-to-end *chains of trust*, from the sensing/production of data to the consumption of *knowledge*. It relies on introducing transverse data *checkpoints* in the chain that enable systems to have their own different lifecycles and characteristics, while establishing confidence for the whole chain.

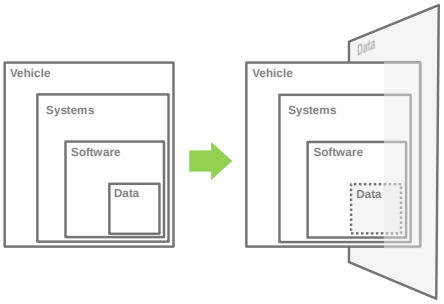

Figure 3 - Data not in systems

This approach is not replacing the traditional way of dealing with data inside systems (and software). It comes in addition to the system-driven approaches, characterized by the following property: the trust we can have in data at system run-time comes from the trust we have in the system at development time.

B. *The "Minimum Viable Chain of Trust" that combines Systems & Data*

Following our principle to *decouple* systems using data *checkpoints*, we propose the following *minimum viable chain of trust* that combines systems and data (Figure 4):

- *Sensing Systems* can be embedded in vehicles or in infrastructures. Logging systems, even not attached to physical sensors, are considered as sensing systems in our analysis.
- Sensing systems are the main sources of data, as *Observations*, for which we strongly propose that they should be considered outside of the system that have generated them, with their own quality and safety attributes.
- *Systems of Insights* or *Systems of Intelligence* that turn multiple input data (*Observations*) into new data (*Knowledge*) with additional value. This *knowledge* generation happens according to its own lifecycle, different from the lifecycles of the *observations*.
- The different kinds of *knowledge* produced by the *Systems of Intelligence* are also to be considered apart, like *observations*.
- *Data- & AI-powered systems & services* represent the last link of the chain in the context of connected vehicles. They consume *knowledge*, once again according to their own lifecycle. A typical example is Advanced Driving Assistance Systems (ADAS) that work according to their own "Sense, Plan, Act" operation principle, the *knowledge* being "plugged" to the "Plan" phase, enabling intelligence. Indeed, this knowledge has been elaborated through the collection of *observations*, which are *senses of other vehicles*, and a fusion and compilation of these observation to a knowledge, which is a sort of *collective intelligence*.

C. *Safe Data*

A key objective regarding data safety is to define what *safe data* (and *safe knowledge*) could mean. The quality attributes of the data are naturally essential for this, using for example the *characteristics* defined in [7]. For data close to sensors (*observations*), those attributes are directly linked to the characteristics of the sensing system that has produced the data.

For instance, it seems safe to collect *observations* stating that a vehicle has detected a loss of traction for one or more of its tires (skidding) through its Electronic Stability Control (ESC) system.

The problem is much more difficult when considering the *distilled*, transformed data (*knowledge*). The value of the *knowledge* should be defined, and the core work to do on data safety is precisely on this topic. Among the ideas that could be considered (non-exhaustive list):

- the quantity (diversity, representativity) and quality of source data is of high influence.
- the complexity of the algorithm mixing the source data can influence the final quality.
- a probability of error should be introduced, for instance when the *distillation* uses stochastic models (and distinguish false positive and false negative is key).

Keeping the same previous example, would it still be safe to consolidate all the skidding *observations* from multiple vehicle into a slippery road conditions warning? Would additional information be needed or relevant (like the temperature to help explain that there may be ice on the road)? When would it be safe to predict that the ice may be gone? What if a vehicle does not detect skidding at this location? How to differentiate ice from oil as a cause to slippery road conditions?

Let's assume that we could build a "map" of all those kinds of events, on which criteria could we consider that it's safe for a vehicle to drive autonomously? Obviously, the positive impacts and benefits from additional knowledge must be balanced with the risks of making decisions based on erroneous information. This balance exercise directly helps in the definition of the operation design domain (ODD) defined in [5].

Additionally, in the context of autonomous driving, the location of geographical information is clearly safety-related. Then the choice of how the location is established (including the physical mean), and the way that this location is represented by data will directly enter in the *confidence level* of the geographic information.

For example, a speed limit can be attached to a large road segment, based on the location of the speed limit sign which can be approximative (5-meter accuracy); on the other side, the curvature of a bend must be precisely located, and this location properly described (center of curvature, special points of the curve, etc.).

In other words, we can imagine an approach of the safety design centered on data: the exercise would consist in choosing which data is necessary to enhance the safety of a transport function (for example), considering all the valuable data, fostering the usage of available data, pushing the collection of other new data, evaluating the confidence potentially associated to these data, and the ways of mitigating the potential lack of confidence by another loop in this data investigation, or by alternative means (i.e. usage of another known subsystem).

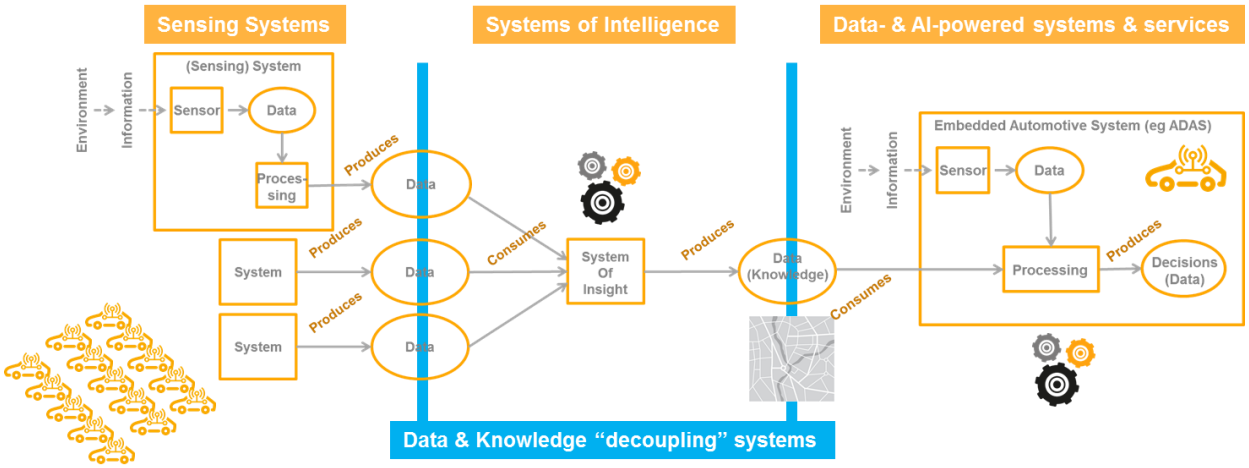

Figure 4: The "Minimum Viable Chain of Trust" that combines Systems & Data

D. *Adapt standards for systems that produce, consume, store or transfer safe data*

For simplification purposes, our minimum viable chain of trust deliberately omits telecommunication systems that transfer data between systems, and storage/access systems that enable lifecycles to be decoupled. We believe that the traditional quality criteria such as integrity, availability, confidentiality, lineage and response times provide the right foundation so that those systems deal with *safe data*, but further analyses are yet to be conducted.

Our minimum viable chain of trust also describes only the case when the last link of the chain is a in-vehicle system (that can rely on off-board *knowledge*). For that case and more generally for all in-vehicle systems, the safety analysis context is the one defined by [2] and [5].

A key question is then "What does it mean for an automotive system [2] or intended function [5] to consume *safe data* or *knowledge*?"

To answer that question, it's crucial to come back to the objectives of the last link of the chain, the actual use case. Using the same weather information for optimizing the energy consumption on one hand and for "calibrating" an emergency braking system on the other hand defines two distinct safety contexts. We believe that we have to introduce some *levels of safety* for data in order to address this point: depending on its usage by the critical system or function, a data item must have a minimum level of safety corresponding to the criticity.

E. *Define Safety for Systems of Intelligence*

Systems of Intelligence produce *safe data* (or *knowledge*). They rely on principles and approaches (cloud computing, distributed systems with specific fault-tolerance solutions, AI, crowdsourcing…) for which there is no or a little experience regarding safety analyses.

We believe that a new area should be opened for the safety of those systems. Two aspects at least should be studied – and balanced: how to make those systems reliable and robust enough, and how to make them produce *safe data*.

VII. Conclusion: Our vision for Data Safety

We have seen that the changing mobility landscape invites to reconsider some aspects of global road traffic safety. Among them, visit and challenge the way data interacts with safety seems to be promising.

On the one hand, to enhance safety, we need more information to give more autonomy to systems, but without losing the huge adaptation abilities of humans. This leads us to ask if safety activities should go from a system-centric view to a data-centric view. Doing so would open the opportunity to build a safety based on decoupled systems, allowing a high level of complexity and dynamic configuration, based on the idea that more and better data means more confidence. The obvious condition to that benefit is that the initial data (the *observations*) are to be produced with a *known level* of confidence.

On the other hand, there is a lot of work to make this shift happen! As any novelty in the quite conservative world of safety, we must show that the new introduced risks are positively balanced by the overall safety enhancement. At the heart of the problem, we must elaborate the rules, the guidance to build *safe data* (*knowledge*) based on source data (*observations*). We must build new criteria and measures, study how they could be linked to existing standards (for system and software). Finally, we must see which types and which level of rules and standardization would best fit *systems of intelligence*.

This work cannot be achieved by a single enterprise, it requires a whole community, and more than that: a new whole community! A community bridging the gap between (at least) two different worlds: the data-digital-web-IT world and the system-and-safety world.

References

[1] ISO/IEC 2382:2015 Information technology – Vocabulary

[2] ISO 26262:2018 Road vehicles -- Functional safety

[3] RTCA DO178C software considerations in airborne systems and equipment certification

[4] SAE ARP4754A Certification considerations for highly-integrated or complex aircraft systems

[5] Road vehicles - Safety of the intended functionality (SOTIF) ISO/PAS 21448:2019

[6] ISO/IEC 25012:2008 Software engineering -- Software product Quality Requirements and Evaluation (SQuaRE) -- Data quality model

[7] ISO/IEC 25024:2015 Systems and software engineering -- Systems and software Quality Requirements and Evaluation (SQuaRE) -- Measurement of data quality

[8] ISO 19157:2013 Geographic information -- Data quality

[9] ISO 20077-1:2017 Road Vehicles -- Extended vehicle (ExVe) methodology

[10] RTCA DO200A Standards for Processing Aeronautical Data

[11] SCSC-127D Data Safety Guidance (Version 3.1 – February 2019) by the SCSC Data Safety Initiative Working Group [DSIWG]. ISBN-13: 9781793375766 https://scsc.uk/scsc-127D