Liability and automation: legal issues in autonomous cars

Giuseppe Contissa[a], Francesca Lagioia[a,b], Giovanni Sartor[a,b]

The deployment of highly automated systems, such as autonomous cars, is going through an accelerated expansion: as usual for emerging disruptive technologies a slow start is followed by a more and more rapid surge. One of the most important legal issues concerning these systems is related to liability for accidents. In particular, highly automated systems will make choices and engage in actions – usually with some level of human supervision, or even without any such supervision. In this context, there is the need to analyse how the decision-making process is split between humans and machines, and critically revise the way tasks, roles, and liabilities are allocated. In this contribution we analyse the impact of automation in the allocation of liability within autonomous cars. We first discuss the tasks allocation between human and automation, and the resulting responsibilities. Then, we analyse how the introduction of different levels automation gives rise to a redistribution of tasks between human and automation and, therefore, a reallocation of the liability burden between the user and the manufacturer.

Task-responsibilities and the impact of automation

In order to introduce the analysis of liability issues, we need to refer to the concept of task-responsibility, i.e. the duty pertaining to the correct performance of a certain task or role.

First of all, we need to identify task-responsibilities of the user, since their violation may result in personal liability. In fact, whenever there is a failure in a complex system, we try to connect the failure with the missing or inadequate execution of a task, and so with the (natural or legal) persons who were responsible for that task. As a consequence of the failure to comply with their task-responsibilities, these persons are subject to blame, penalties, and/or the payment of damages.

Secondly, we need to identify task-responsibilities of the automated system, namely, the requirements the system should comply with. These are also relevant, since a failure to meet them may make the system's producers or maintainers liable. With the introduction of higher levels automation, as task-responsibilities are progressively delegated to technology, liability for damages shifts from human operators to the organisations that designed and developed the technology, defined its context and uses, and are responsible for its deployment, integration, and maintenance.

In this context, it is necessary to adopt a systematic approach to match the degree of automation to different responsibilities of users of automated systems at different levels as well as to the responsibilities of other actors involved (managers, producers, and maintainers)(Contissa et al. 2013).

To this end, we consider the Level Of Automation Taxonomy (LOAT), developed by SESAR 16.5.1 (Save and Feuerberg, 2014) used to assess the levels of automation introduced by a new technology and to determine the corresponding impacts on the division of tasks between humans and machines.

The LOAT table provides criteria for assigning a level of automation to a technology with regard to four different cognitive functions: information acquisition (A), information analysis (B), decision-making (C), and action implementation (D). Figure 1 shows a simplified version of the LOAT: all columns start with level 0, corresponding to a fully manual accomplishment of the task, without any technical support. At Level 1 the task is accomplished with "primitive" technical tools, i.e., low-tech non-digital artefacts. From level 2 on upwards, "real" automation is involved, and the role of the machine becomes increasingly significant up to the level where the task is fully automated.

A certain technology may have different levels of automation, according to whether the actors are dealing with the four cognitive functions mentioned above.

The LOAT expresses varying levels of interaction between humans and the technology in question. In the first instance, it is used to better understand technology. It provides an accurate account of human-machine interaction and serves as a tool for refining the concept of automation. By conceptualizing automation on the basis of the human factor, it generates awareness of human-machine interaction.

[a] CIRSFID, Alma Mater – Università di Bologna, Italy, giuseppe.contissa@unibo.it
[b] Law Department, European University Institute, Florence, Italy, Francesca.Lagioia@EUI.eu, Giovanni.Sartor@EUI.eu

A INFORMATION ACQUISITION	B INFORMATION ANALYSIS	C DECISION AND ACTION SELECTION	D DECISION AND ACTION SELECTION
A0 Manual Information Acquisition	B0 Working-memory based Information Analysis	C0 Human Decision Making	D0 Manual Action and Control
A1 Artefact Supported Information Acquisition	B1 Artefact Supported Information Analysis	C1 Artefact Supported Decision Making	D1 Artefact Supported Action Implementation
A2 Low Level Automation Support of Info Acquisition	B2 Low Level Automation Support of Info Analysis	C2 Automated Decision Support	D2 Step by step Action Support
A3 Med. Level Automation Support of Info Acquisition	B3 Med. Level Automation Support of Info Analysis	C3 Rigid Automated Decision Support	D3 Low Level Support of Action Sequence Execut.
A4 High Level Automation Support of Info Acquisition	B4 High Level Automation Support of Info Analysis	C4 Low Level Automatic Decision Making	D4 High Level Support of Action Sequence Execut.
A5 Full Automation Support of Info Acquisition	B5 Full Automation Support of Info Analysis	C5 High Level Automatic Decision Making	D5 Low Level Automation of Action Sequence Exec
		C6 Full Automatic Decision Making	D6 Medium Level Automat. of Action Seq. Execut.
			D7 High Level Automation of Action Seq. Execut.
			D8 Full Automation of Action Sequence Exec

Figure 1. The Level Of Automation Taxonomy (LOAT)

Source: readapted from Save, L., Feuerberg, B.(2012), pp. 48-50.

Assessing the liability impact of autonomous cars

Road transportation is a domain, where the technological development is introducing high levels of automation. In order to analyse the issues related to the introduction of driving automation, the industry has adopted the SAE international standard J3016, a taxonomy describing the full range of levels of driving automation in on-road motor vehicles (SAE International 2016). The classification system is based on the levels of driver intervention and attentiveness required, resulting in a scale of six levels of automation, ranging from level 0 (no driving automation) to level 5 (full driving automation).

The levels of driving automation are defined by reference to the specific role played by each of the three primary actors (the human driver, the driving automation system, other vehicle systems and components) in performing dynamic driving tasks (DDT). Dynamic driving tasks include all real-time operational and tactical functions required to operate a vehicle in on-road traffic (e.g. lateral/longitudinal motion, monitoring the driving environment, etc.), excluding the strategic functions such as trip scheduling and selection of destinations and waypoints.

Functions 3 and 4 are collectively referred to as "object and event detection and response" (OEDR).

In levels from 0 to 2, the driver performs part of all of the DDT, whereas form level 3 to 5 the Automated Driving System (ADS)[1] performs the entire DDT, while engaged.

In this section, we map the driving automation levels described in the SAE standard to the four cognitive functions of the LOAT. We identify, for each driving automation level and for each cognitive function involved, the responsibilities among user and manufacturer, and the resulting legal liabilities. In analysing the different levels, we focus on the driving tasks affected by automation. We do not consider the cases in which the automated driving system is disengaged, since in these cases all dynamic driving tasks

[1] Automated Driving System (ADS) is "the hardware and software that are collectively capable of performing the entire DDT on a sustained basis, regardless of whether it is limited to a specific operational design domain (ODD)." (SAE J3016, pag. 3)

are performed by users. Moreover, we do not consider liabilities not directly related to driving tasks, for example liabilities related to the maintenance of the vehicle and of the ADS systems, which can be apportioned across different actors (the driver, the owner, the manufacturer, etc.).

Level 0

In SAE J3016 standard, level 0 (no driving automation) the driver performs the entire DDT, while the driving automation system (if any), does not perform any part of the DDT on a sustained basis, although the systems in the vehicle may provide warnings or support (e.g. anti-lock brake systems, conventional cruise control, or electronic stability control).

The user is entirely responsible for the following tasks: the acquisition and filtering of information; its analysis; the generation of decision options and the selection of the appropriate ones; the execution and control of actions.

According to the LOAT, the tasks involved in driving at level-0 automation are correspondingly classified as A1 (Artefact-Supported information Acquisition); B1 (Artefact-Supported Information Analysis); C1 (Artefact-Supported Decision Making); and D1 (Artefact-Supported Action Implementation).

In case of accident caused by a failure in executing one of these tasks, the manufacturer is liable only when he provided a defective or non-standard compliant tool, that had a role in the causation of the accident (for example, defective brakes preventing the car to avoid the collision). In all the other cases, user's liability is to be considered.

Levels 1 and 2

Level 1 (Driver Assistance) is defined as "the sustained and ODD[2]-specific execution by a driving automation system of either the lateral or the longitudinal vehicle motion control subtask of the DDT (but not both simultaneously) with the expectation that the driver performs the remainder of the DDT." The driver must supervise the driving system performance by completing the OEDR subtask of the DDT as well as performing the other dimension of vehicle motion control.

Level 2 (Partial driving automation) differs from Level 1 because the driving system is expected to execute both the lateral and the longitudinal vehicle motion control simultaneously. Thus, the difference between Level 1 and 2 is merely quantitative, in the sense that it concerns the extent of the automated motions of the vehicle, under human control.

Regarding information acquisition, the system uses predefined criteria to integrate, filter and highlight information; supports the user in integrating data, filters information items and highlights the most relevant. The user, after being instructed on the information acquisition functions, monitors their performance.

Regarding information analysis, the system supports the user in comparing, combining and analysing information items concerning the status of the system's processes, based on parameters pre-defined by the user, and alerts him if the results of analysis require his/her attention. The user defines the parameters of the process, takes duly into account the system's outcomes and reacts to its alerts.

Regarding decision making, the system proposes decision alternatives and informs the user about its determinations. The user is monitors the determinations of the system.

Regarding action implementation, the system performs automatically a sequence of actions after activation by the user. The user monitors the sequence and interrupts its execution when needed.

According to the LOAT, the tasks involved in driving at Levels 1 and 2 would be classified as A5 (Full Automation Support of Information Acquisition), B4 (High-Level Automation Support of Information Analysis), C4 (Low-Level Automatic Decision Making), and D4 (High-Level Support of Action Sequence Execution) at maximum.

In case of accident caused by a failure in executing one of the functions, the manufacturer is liable only when providing a defective or non-standard compliant tool that had a role in the causation of the accident (for example, a production defect concerning brakes that cannot prevent the car to avoid the collision, or a design defect concerning the user interface being unable to provide correct information, or a warning defect concerning the lack or insufficient information on the functioning of the automation provided to the user). In all other cases, user's liability is considered, since most of the dynamic driving tasks fall under the user's control and responsibility. In particular, the user might be found liable when acting without reasonable care, including when s/he failed in monitoring the performance of the system, taking duly into account the system's outcomes and reacting timely to its alerts or any other risky situation.

Level 3

In Level 3 (Conditional Driving Automation), the driver (while the ADS is not engaged) is expected to (1) verify the operational readiness of the ADS-equipped vehicle; (2) determine when engagement of ADS is appropriate. When the ADS is engaged, the driver shall be ready to in-

[2]The Operational Design Domain (ODD), is the set of the specific conditions under which a given driving automation system or feature thereof is designed to function, including, but not limited to, driving modes.

tervene and to take back the control when requested. The ADS, while engaged, is expected to (1) perform the entire DDT; (2) determine whether ODD limits are about to be exceeded and, if so, issue a timely request to intervene to the driver; (3) determine whether there is a DDT performance-relevant system failure of the ADS and, if so, issues a timely request to intervene to the driver; (4) disengage an appropriate time after issuing a request to intervene, or immediately upon driver request.

Regarding the acquisition of information, the system supports the information acquisition; has predefined criteria for integrating, filtering and highlighting information; supports the user in integrating data, filtering information items and highlighting the most relevant. The user, after being instructed on how to use the system, monitors its performance.

Regarding information analysis, the system performs comparisons and analyses of data available on the status of the process being followed, based on parameters defined at design level. The system triggers visual and/or aural alerts if the analysis produces results requiring attention by the user. The user takes duly into account the system's outcomes and reacts to its alerts.

Regarding the decision and action selection, the system generates decision options, selects the appropriate ones and decides all actions to be performed. The user can safely turn his/her attention away from the dynamic driving task, but must still be prepared to intervene when called upon by the vehicle to do so, or whenever the ODD limits are about to be exceeded.

Regarding action implementation, the system initiates and executes automatically a sequence of actions, while the user monitors all the sequence and interrupts it during its execution, when requested by the system or whenever the ODD limits are about to be exceeded.

According to the LOAT, the tasks involved in driving at level 3 would be classified as A5 (Full Automation Support of Information Acquisition); B5 (Full Automation Support of Information Analysis); C6 (Full Automatic Decision Making) and D6 (Medium-Level Automation of Action Sequence Execution).

In case of accident caused by a failure in executing one of the functions, since most of the dynamic driving tasks fall under the system's control, the manufacturer is liable (1) when providing a defective or non-standard compliant tool that had a role in the causation of the accident; (2) whenever the system fails to carry out the assigned task with a level of performance that is (at least) comparable to that reached by a human adopting due care under the same conditions.

User's liability is to be considered when (1) the user does not respond appropriately to a request to intervene; (2) whenever the ODD limits are exceeded, since s/he is expected to monitor the system's performance, take duly into account its outcomes and react to its alerts.

Level 4 and 5

Level 4 (High driving automation) is the level where the driver, while the ADS is not engaged, shall (1) Verify the operational readiness of the vehicle; (2) Determine whether to engage the ADS.

While the ADS is engaged, the driver becomes a passenger (if physically present in the vehicle) or a dispatcher. The passenger/dispatcher is not expected to perform the DDT or DDT fallback and to determine whether and how to achieve a minimal risk condition. Thus, the automated DDT fallback and minimal risk condition achievement capability of the system is the primary difference between level 3 and level 4 ADS features. However, the passenger will become a driver after a request of disengagement.

The ADS, while not engaged, is expected to allow engagement only within its ODD. While the ADS is engaged, it is expected to (1) perform the entire DDT; (2) eventually issue a timely request to intervene; (3) perform the DDT fallback, automatically transiting to a minimal risk condition when (a) a relevant DDT failure occurs; or (b) the user does not respond to a request to intervene; or (c) a user requests that the system achieve a minimal risk condition; (4) disengage if appropriate only after (a) it achieves a minimal risk condition or (b) the driver is performing a DDT; (5) delay eventually user-requested disengagement.

Level 5 (full driving automation) differs from Level 4 only for the unconditional/not-ODD specific performance by the ADS. This means that the system can operate the vehicle under all driver manageable all-road conditions, i.e. there are no geographical, weather or time-based restrictions.

Therefore, with regard to the acquisition of information, the system supports the information acquisition; has predefined criteria for integrating, filtering and highlighting information; integrate data, filter information items and highlight the most relevant for the user.

With regard to the information analysis, the system shall perform comparisons and analyses of data available on the status of the process being followed based on parameters defined at design level. The system triggers visual and/or aural alerts if the analysis produces results requiring attention by the user. The user may take into account the system's outcomes and react to its alerts.

With regard to the decision and action selection, the system shall generate decision options, select the appropriate ones and decide all actions to be performed. The user can safely turn his/her attention away from the dynamic driving task but may intervene when called upon by the vehicle to do so, or, in level 4, whenever the ODD limits are about to be exceeded.

With regard to the action implementation, the system initiates and executes a sequence of actions. The user can only monitor part of it and has limited opportunities to interrupt it, for example, in case of user-requested disengagement under appropriate conditions.

According to the LOAT, the tasks involved in driving at levels 4 and 5 would be classified as A5 (Full Automation Support of Information Acquisition); B5 (Full Automation Support of Information Analysis); C6 (Full Automatic Decision Making) and D7 (High-Level Automation of Action Sequence Execution).

In case of accident caused by a failure in executing one of the functions, since all dynamic driving tasks (in level 4, within the ODD limits) fall under the system's control, the manufacturer is liable (1) when providing a defective or non-standard compliant tool that had a role in the causation of the accident; (2) whenever the system fails to carry out the assigned task with a level of performance that is (at least) comparable to that reached by a human adopting due care under the same conditions. Since at this level of automation the user is not expected to respond to any request to intervene when the ADS is engaged, user's liability is to be considered only when the ODD limits are exceeded, or when, following a user's request, the system permits the ADS disengagement.

It should be noted that vehicles under Level 4 or 5 may be designed to be exclusively operated by ADS for all trips. In this case, they may be designed without user interfaces, such as braking, accelerating, steering, etc. These categories of vehicles do not necessarily involve a driver. Whenever user interfaces are missing, the user will not be able to intervene in any of the Dynamic Driving Tasks, and therefore cannot be subject to the related liabilities.

Conclusions

When automated systems are increasingly introduced into a complex system, the main effect is that liability for damage or harm is gradually transferred from humans to enterprises using the automated technology, that replaced the human operator and /or to the technology developer (programmer, manufacturer) that created the technology.

While the trend of transferring liability from the individual to the enterprise has been observed for quite a long time (Brüggemeier 2006), new AI technologies accelerate this trend, since they deeply impact on the tasks of human operators, not only quantitatively but also qualitatively, replacing human operators in their higher cognitive functions, ranging from the analysis of information to the selection of a decision or an action, to the fully automated implementation of the chosen action.

Of course, not all advanced technological systems possess these cognitive functions to the same extent. For example, many currently employed automated systems are not designed to automatically implement the chosen actions, but only to suggest actions to be executed by the human operator.

In order to evaluate the final liability allocation between different actors, it is necessary to assess each technology's different levels of automation in performing different cognitive functions (acquiring information, analysing information, making decisions, and acting on them).

It should be noted that intermediate levels of automation are sometimes those that create higher levels of legal risk for certain actors. It happens because of a high fragmentation of task-responsibilities between the automated technology and the operator in these levels, possibly leading in some circumstances to uncertainty in the assignment of tasks. In addition, intermediate levels of automation usually imply also a greater complexity in the human-machine interface, since fragmented tasks require more interaction between a technology and its operator.

In legal terms, this may translate to an increased duty of care, resulting in a higher liability risk for (a) the operator; (b) the organisation employing the operator, both for vicarious liability and for organisational liability; and, finally, (c) the producer of the technology, since higher complexity in the human-machine interface would increase the risk of technological failure (Schebesta et al. 2015).

Moreover, from a safety perspective, in intermediate levels characterised by higher levels of automation in acquisition and analysis of information, and lower levels in the action selection and implementation, it is questionable, whether the user has concrete possibilities to intervene in the decision-making process, since s/he will not be in the position to critically revise and evaluate (in due time) all the environmental information required to take a decision. In time-sensitive situations, the action by the user will be mostly instinctive, rather than based on a rational process. These may result in worse outcomes. For example, the choice of braking may be taken and executed faster by automation, rather than analysed and proposed to the user, that will subsequently assess, select and implement the proposed action, or even disregard it. Besides, even in non-time-sensitive situations, it is possible to argue that an autonomous car is able to collect the relevant information and evaluate all the possible outcomes of alternative actions more precisely and faster than any human should, for example, considering the car speed and its distance from an obstacle (e.g. a pedestrian or another car) to evaluate the more appropriate manoeuvre (e.g. swerve or stay on course and break). For these reasons, we believe that in many domains, full automation is preferable to intermediate levels of automation, both from the legal and the safety perspective.

References

Brüggemeier, G. Common principles of tort law: a pre-statement of law. British Institute of International and Comparative Law, 2006.

Contissa, G., Laukyte, M., Sartor, G., Schebesta, H., Masutti, A., Lanzi, P., Marti, P., and Tomasello, P. Liability and automation: Issues and challenges for socio-technical systems. Journal of Aerospace Operations 2, 1-2 (2013), 79–98.

SAE International, Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems (J3016), https://www.sae.org/standards/content/j3016_201609/

Save, L., and Feuerberg, B. Designing human-automation interaction: a new level of automation taxonomy. In Proc. Human Factors of Systems and Technology 2012 (2012).

Schebesta, H., Contissa, G., Sartor, G., Masutti, A., Paola, T., and Taurino, D. Design according to liabilities: Acas x and the treatment of ads-b position data. In Proceedings of the SESAR Innovation Days (2015) EUROCONTROL (2015).