POSITION PAPER

Brussels, 8 February 2019

Orgalim comments on the upcoming Impact Assessment of the Machinery Directive

Representing the European technology companies driving the digital transformation of industry and society, Orgalim is certain that artificial intelligence (AI) will be a central pillar of Europe's future competitiveness. Many of the firms we represent are at the leading edge globally in the development of AI applications integrated in consumer and industrial products ('embedded' AI), and embedded AI is already widely used across our industries to deliver efficiency gains along the value chain. More broadly, embedded AI is enabling technology solutions to critical societal challenges – from climate change and the energy transition to the future of mobility, an ageing society and security in the digital age.

Although AI has been deployed safely in manufacturing for decades, the rapid pace of technological evolution in recent years has prompted questions regarding the EU legal framework. Against this background, the European Commission has initiated an evaluation of the Machinery Directive (MD) to assess whether it remains fit for purpose in addressing machinery incorporating new technologies such as AI. The MD is a core piece of legislation for our industries, ensuring a high level of safety while providing legal stability to our companies. Given the importance of a stable and predictable legal framework for attracting investment into key future tech like AI, we believe it will be crucial to pursue a prudent approach to regulation in this arena.

As such, we welcome the publication of the Commission's initial evaluation of the MD, which states that this important piece of legislation remains fit for purpose – primarily thanks to the structure of the Directive and its technology-neutral essential health and safety requirements. Manufacturers can demonstrate compliance with Annex I of the Directive by making use of more than 800 existing harmonised standards, which represent the technological state of the art. Furthermore, our industry continues to actively contribute its experience and technical expertise to the drafting of standards to cover new technologies.

In this same context, Orgalim has also welcomed the Commission Communication on AI of April 2018, which positively recognised the security of the EU safety framework, stating that it: "already addresses the intended and foreseeable (mis)use of products when placed on the market. This had led to the development of a solid body of standards, in the area of AI-enabled devices that are continuously being adapted in line

Orgalim
BluePoint Brussels
Boulevard A Reyers 80
B1030 | Brussels | Belgium

+32 2 206 68 83
secretariat@orgalim.eu
www.orgalim.eu
VAT BE 0414 341 438

with technological progress. The further development and promotion of such safety standards and support in EU and international standardisation organisations will help enable European businesses to benefit from a competitive advantage and increase consumer trust."

As a further step in the evaluation of the MD, the Commission will perform an impact assessment focusing on new technologies, with the aim of analysing whether the Directive can continue in the years to come to guarantee the placing on the market of safe machines and to ensure that users and consumers trust these machines, even when they integrate embedded new technologies.

As the current MD has performed well, increased occupational health and safety and continues to meet its objectives for both ensuring the placement on the market of safe machines (including those that embed new technologies) and establishing a high level of trust in these machines among users and consumers, Orgalim is strongly recommending that this piece of legislation should not be updated in the near future.

THE BOUNDARIES OF AI

AI has climbed up the political agenda in recent years, with discussions often reflecting latent fears that the technology might begin to act autonomously and provoke new, unmanageable dangers. Yet AI has in fact been deployed effectively and safely in manufacturing for a number of decades, supporting European industries in becoming even more competitive on the global market.

It is commonly agreed within the scientific community that the type of AI being used and developed today constitutes what is known as **narrow artificial intelligence**, whereby a machine can only perform an action assigned from the outset by human – whether a designer, computer specialist or manufacturer. If manufacturers wish to place a machine embedding AI functions on the market, they are under an obligation to carry out a risk assessment, followed by conformity assessment procedure(s), taking into account the risks associated with AI agents. Through this process, the manufacturer determines the limits of the machinery and its intended use (which also includes any AI functions).

It is within this framework that all AI applications (industrial or otherwise) operate today and will continue to do so in the future. Buzzwords such as 'general AI' or 'super AI' remain the preserve of science fiction and should not be used in relation to equipment covered under the scope of the MD.

Orgalim suggests the following definition as a workable reference point when examining whether Union harmonisation legislation is fit to ensure user safety and can still be considered trustworthy when applied to products and systems that integrate AI.

DEFINITION OF AI

> **AI (artificial intelligence) refers to computer systems based on algorithms designed by humans that, given a complex task, operate by processing the structured or unstructured data collected in their environment according to a set of instructions, determining the best step(s) to take to perform the given task, via software or hardware actuators. AI computer systems can also adapt their actions by analysing how the environment is affected by their previous actions.**

EXAMPLES OF AI FEATURES

AI is mostly powered by '**machine learning'**: computer systems are configured to learn from data how to carry out a specific task and how to make predictions from it. There are two main machine learning techniques: **supervised learning** and **unsupervised learning**.

Most AI systems used in industrial applications today operate under **supervised learning**: a human inputs a large amount of 'labelled data' (a group of samples with one specific meaning or tag) to train an algorithm to learn the relationship between the input and the desired output, as set by the human designer.

For instance, if the designer wants a robot to pick out a package of a specific colour or dimension from a basket containing all kinds of packages, it can be taught to do this on the basis of labelled data input by the human operator. This is called supervised learning because the human is acting as an instructor in the process. Supervised learning uses the classification technique to classify the data input by category (red/green-true/false) or the regression technique to predict a numeric or continuous value (such as the maximum temperature of a fluid).

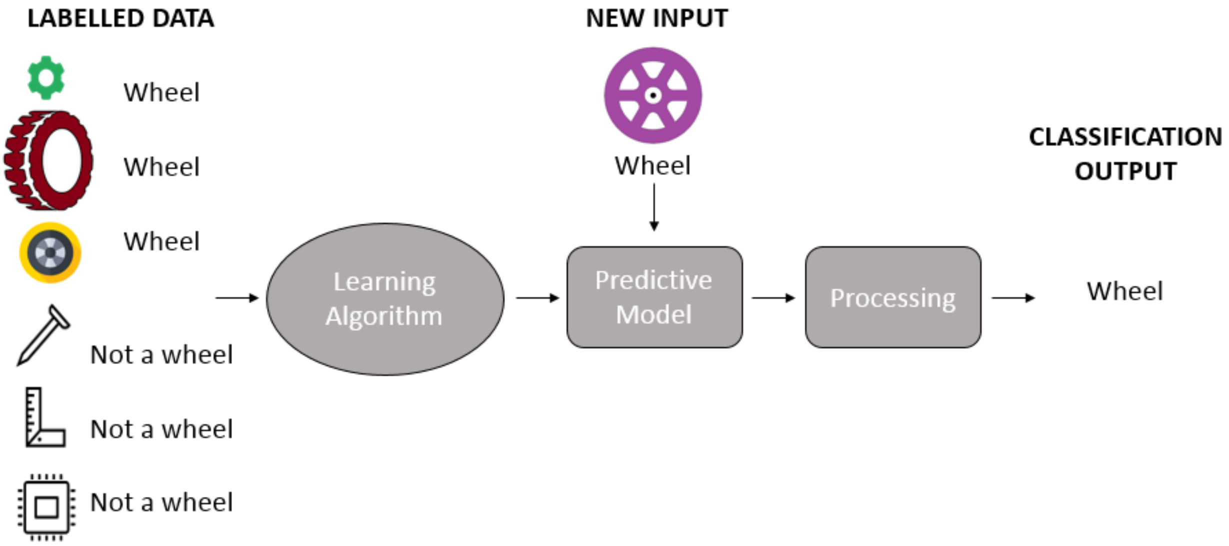

In the example in the diagram above, the objective is to classify wheels and help the machine perform this classification by itself when it receives a new input it has never seen before. Labelling data means that a human/teacher classifies the wheels first on the basis of their shape for instance, then their dimension, then their colour. Once the human has labelled the data, a predictive model is elaborated. When the algorithm 'sees' a wheel it has not encountered before, it will then be able to classify it based on the characteristics (tags/labels) the human assigned to other wheels previously.

Some AI systems, on the other hand, use **unsupervised learning** to find a structure or pattern in large sets of 'unlabelled data' by clustering the data, with no predetermined right or wrong output.

A modern and versatile method to support learning algorithms is the 'deep' (multi-layered) artificial neural network. This approach is inspired by the structure and function of the human brain, which learns to adapt its logical function through observations, actions and feedback on effects.

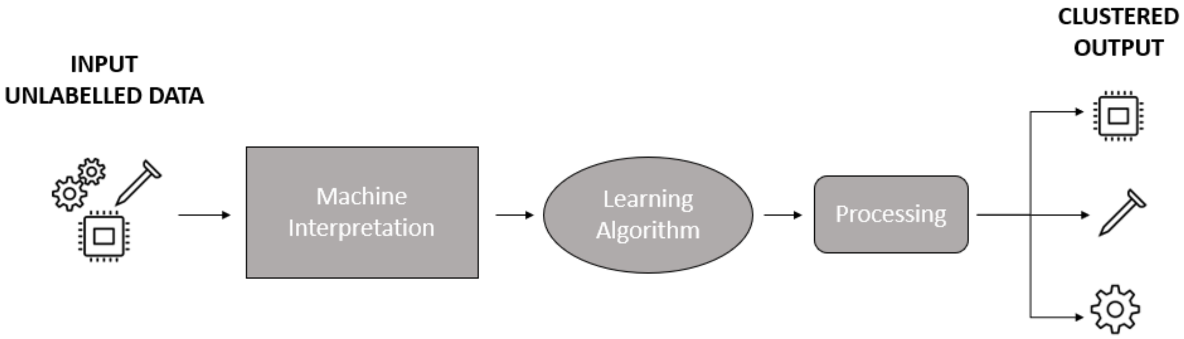

Under unsupervised learning, the data provided is not labelled (the data has not been assigned one or several characteristics); rather, the learning algorithm will try to cluster the information by finding similarities in the data it has received.

AI: FRAMEWORK OF OPERATION

AI and machine learning are subject to **design limitations** that are both mathematical and generic in nature. To train its learning models, the computer system requires:

1) Huge **datasets,** comprising data either labelled by a human (supervised learning) or unlabelled (unsupervised learning). The learning success of the system increases in line with both the amount of data provided and the quality of this data, i.e. the accuracy of labels, amount of noise (additional meaningless information) and ratio of relevant information within the absolute data volume.
2) Technical, legal and business-led **access to this machine-readable data**, which then needs to be pre-processed in order to make it fit to the algorithms used in narrow artificial intelligence.
3) **Algorithms** designed by a human (computer specialist, manufacturer, etc.) that assign a specific task to the computer system.
4) Huge **computational power** to run the algorithms associated with training and querying neural networks (unsupervised learning).
5) **Time** to train the algorithm through a process of trial and error until the AI system has reached the intended level of success in its predictions and analyses.

ADVANCED ROBOTICS: EMBEDDING AI IN MACHINERY

AI systems can be interfaced with either **software or robotic actuators** (a motor function) to perform a specific task intended for professional or consumer use. These actuators can generate either visual, readable or spoken information, make a suggestion to the user (such as a maintenance diagnostic) or initiate an autonomous or semi-autonomous robotic process for a specifically intended use (such as drilling or assembling a machine).

In an industrial environment, AI is used for application such as:

- Optimising product quality in manufacturing operations
- Reducing equipment breakdowns
- Detecting quality issues

- Identifying defects and deviations in equipment
- Improving the transportation of items within a warehouse to identify obstacles and adjust to achieve an optimal route
- Processing more traditional operations such as carving, cleaning, assembling, welding, painting, picking and placing, packing, palletising, etc.
- Detecting cyber-attacks and preventing damage to humans or their environment
- Maintenance of machines
- Increasing occupational health and safety of work equipment machinery

As made clear in the outline of the framework of operation presented above (narrow artificial intelligence), **a machine will never perform any task other than the one it is designed for**. The manufacturer may program the machine learning process to use new patterns during the operational phase to reach an optimal output (e.g. the best temperature of emission in a gas turbine) but this does not change in any way the overall mode of operation of the machine itself. This controlled operation of the machine ensures that the manufacturer's two core objectives are met: placing safe machines on the market and establishing a high level of user and consumer trust in machines embedding AI functions.

The design and operation of the machinery control system remain under the control of the human designer (i.e. the machinery manufacturer or software developer), even when the system uses unsupervised learning and adaptive algorithms: **it can in no way reprogram its own purpose**.

THE CURRENT MACHINERY DIRECTIVE: MACHINES ARE AND WILL REMAIN UNDER THE CONTROL OF THE MANUFACTURER AND THE USER

Taking into account the principles of machine learning as described above, it can be clearly shown that the current MD and the essential requirements it defines can be easily applied to machinery integrating AI functions. AI-embedded machinery functions exclusively within a working perimeter defined by the manufacturer. This working perimeter is either a structured environment or an unstructured environment under strict human supervision, such as the premises of a factory. It is therefore misleading to make any comparison between this kind of machinery operating in a specific working environment and, for example, driverless vehicles in an unstructured environment.

Before placing a machine on the market, the manufacturer will make sure that it will remain safe throughout its life cycle. Fulfilling this objective begins at the design phase, following which the product undergoes a compulsory risk assessment, and finally a series of conformity assessment procedures. When implementing the safety concept, the manufacturer will give priority to safety functions over all other functions of the machine. Even if a machine integrates AI functions, the actions it performs will always be within the scope of the intended use as defined by the manufacturer.

What follows is an explanation of the main articles of Annex I of the MD that may be applicable to guarantee the placing on the market of safe machines, whether or not they integrate AI agents. When performing the risk assessment, the manufacturer will consider which of the essential health and safety requirements need to be taken into consideration for the machine to operate safely, on the basis of the intended use. The risk assessment procedure will also consider each situation that may pose a problem for safety, with the MD providing a sound legal framework in this respect.

Annex I Essential health and safety requirements to the design and construction of machinery	Implications for machines embedding AI
General principles	
1. The manufacturer of machinery or his authorised representative must ensure that a risk assessment is carried out in order to determine the health and safety requirements which apply to the machinery.	**AI (artificial intelligence) is never 'purpose free'** or capable of 'free will or decision'. Machines integrating AI are programmed to learn from a pattern and act accordingly, based on the initial programming, in order to perform the intended use of the designer/manufacturer.
The machinery must then be designed and constructed taking into account the results of the risk assessment	**Adaptive algorithms** adapt the patterns and improve the path towards the predefined goal, enabling the machine to better perform the tasks it was initially programmed to complete (the algorithms do <u>not</u> change the *behaviour* or the intended use of the machine).
By the iterative process of risk assessment and risk reduction referred to above, the manufacturer or his authorised representative shall:	AI functions and machine learning techniques cannot change the intended use of the machine that is set by the manufacturer.
- **Determine the limits of the machinery**, which include the intended use	Machinery embedding AI functions is designed to operate with both programming and physical limits.
and	
any reasonably foreseeable misuse thereof,	AI functions cannot autonomously change the intended use set by the manufacturer. **Any change in the *behaviour* of the machine will only occur as designed by the manufacturer.**
- **Identify the hazard** that **can** be generated by the machinery and the associated hazardous situations,	When performing the risk assessment, the manufacturer knows which risks are associated with the use of the machine.

- **Eliminate the hazards** or reduce the risks associated with these hazards by application of protective measures in the order of priority established in section 1.1.2(b).	To this end, the manufacturer will take protective measures.
1.1.2 Principles of safety integration	
(a) Machinery must be designed (...) so **that it is fitted for its function** (...) without putting persons at risk when these operations are carried out under the conditions foreseen but **also taking into account any reasonably foreseeable misuse thereof.** (...) throughout the foreseeable lifetime of the machinery	From the design phase onwards, the manufacturer takes into account the requirements for a specific function (AI functions have to be included in this) and also any foreseeable misuse.
(...) (c) The machinery must be designed and constructed in such a way **as to prevent abnormal use if such use would engender a risk.**	AI functions operated by machine learning techniques cannot change the intended use of the machine as set by the manufacturer from the outset and when performing the risk assessment.
Control systems	
1.2.1. Safety and reliability of control systems	
Control systems must be designed (...) in such a way as to prevent hazardous situations from arising. Above all, they must be constructed in such a way that:	Control systems are increasingly run by AI functions rather than by a human; however, the human still maintains the ultimate control of the machine.
- They can withstand the intended operating stresses and external influences,	The machine control system must be able to resist external influences, regardless of the type of influence. Such influences can be caused by the failure of a safety component such as a light barrier. Or they could be caused by a cyber attack that, for example, changes the parameters of a control system. According to the provisions of Annex I, all of these influences must not lead to dangerous machine situations.

- a fault in the hardware or the **software** of the control system does not lead to hazardous situations, - Errors in the control system logic do not lead to hazardous situations. (...) Particular attention must be given to the following points: - the machinery must not start unexpectedly, - the parameters of the machinery must not change in an uncontrolled way, where such change may lead to hazardous situations, - the machinery must not be prevented from stopping if the stop command has already been given, - no moving part of the machinery or piece held by the machinery must fall or be ejected, - automatic or manual stopping of the moving parts, whatever they may be, must be unimpeded, - the protective devices must remain fully effective or give a stop command, - the safety-related parts of the control system must apply in a coherent way to the whole of an assembly of machinery and/or partly completed machinery. - **For cable-less control, an automatic stop must be activated when correct control signals are not received, including loss of communication.**	Machine learning techniques are based on input data. It is the human programmer who decides which dataset to use as input for the machine to develop an **algorithm**. The developer ensures that the dataset is not corrupted by external influences. The machine learning phase is ended before the machine is put into service.This is also relevant for cybersecurity. The AI software is pre-programmed to shut down if the control system develops a fault that could lead to a hazardous situation. An AI system is even more effective than a human in detecting errors in the control system logic when this is defined in their programming or as part of machine learning.
1.2.2. Control devices	
From each control position, the operator must be able to ensure that no-one is in the danger zones, **or the control system must be designed and constructed in such a way that starting is prevented while someone is in the danger zone**.	The Directive already stipulates that **systems must be designed and constructed in such a way that starting is prevented while someone is in the danger zone**; this is of course also true for a control system operated by AI functions.

1.2.3. Starting	
It must be possible to start machinery only by voluntary actuation of a control device provided for the purpose.	The human maintains the ultimate control over starting the machine, either by programming the machine to start or starting it directly on site in a factory.
The same requirement applies: - when restarting the machinery after a stoppage, whatever the cause, - when effecting a significant change in the operating conditions.	The operating conditions are external to the machine e.g. heat, resistance, speed. This does not refer to the operating mode, which is part of the programming. Any update in the software is still something that the human operator must confirm.
(…) For machinery functioning **in automatic mode**, the starting of the machinery, restarting after a stoppage, or a change in operating conditions **may be possible** without intervention, provided this does not lead to a hazardous situation.	The parameters of automatic mode are set by a human. All the parameters of all these situations have been programmed by a human.
1.2.4. **Stopping** **1.2.4.1. Normal stop** Machinery must be fitted with a control device whereby the machinery can be brought safely to a complete stop. Each workstation must be fitted with a control device to stop some or all of the functions of the machinery, depending on the existing hazards, so that the machinery is rendered safe. The machinery's stop control must have priority over the start controls. Once the machinery or its hazardous functions have stopped, the energy supply to the actuators concerned must be cut off. **1.2.4.2. Operational stop** Where, for operational reasons, a stop control that does not cut off the energy supply to the actuators is required, the stop condition must be monitored and maintained. **1.2.4.3. Emergency stop**	The provisions of this article reflect that the user has complete control over the machinery, including any change of parameter which may be caused by a cyber attack.

Machinery must be fitted with one or more emergency stop devices to enable actual or impending danger to be averted.	
1.2.5. Selection of control or operating modes	
The control or operating mode selected must override all other control or operating modes, with the exception of the emergency stop. If machinery has been designed and constructed to allow its use in several control or operating modes requiring different protective measures and/or work procedures, it must be fitted with a mode selector which can be locked in each position. Each position of the selector must be clearly identifiable and must correspond to a single operating or control mode. (…)	AI functions cannot change the intended operating modes of a machine, which are defined by the manufacturer. The operating mode can only be selected by an operator within the limits defined by the manufacturer. The operator will still see (in a clear and identifiable manner) which AI programs or operating modes are running at a given time.
1.2.6. Failure of the power supply	
The interruption, the re-establishment after an interruption or the fluctuation in whatever manner of the power supply to the machinery must not lead to dangerous situations. Particular attention must be given to the following points: (…) - the parameters of the machinery must not change in an uncontrolled way when such change can lead to hazardous situations, - the protective devices must remain fully effective	The values of the parameters were defined/selected by the user. If AI functions change some parameters, the user always has the ability to override any modifications. This provision covers all functions of the machine, including AI-driven functions.
1.3.6. Risks related to variations in operating conditions	
Where the machinery performs operations under different conditions of use, it must be designed and constructed in such a way that selection and adjustment of these conditions can be carried out safely and reliably.	The adaptation of AI-driven functions to the operating conditions must not lead to unsafe situations. To this end, when performing the risk assessment the manufacturer will take into consideration any operation of the machine that is run differently and will put in place any adjustment measures.
1.3.9. Risks of uncontrolled movements	
When a part of the machinery has been stopped, any drift away from the stopping	Dangers for the operator or uninvolved bystanders are triggered, for example, by

position, for whatever reason other than action on the control devices, must be prevented or must be such that it does not present a hazard.	movements of the machine. If AI-driven movements trigger a hazard or a risk, the manufacturer must take protective measures that correspond to the state of the art. For example, they could program the machine (robot) to stop when at a safe distance from an obstacle (whether human, animal or object).
1.6. MAINTENANCE	
1.6.1. Machinery maintenance	
(...) In the case of automated machinery and, where necessary, other machinery, a connecting device for mounting diagnostic fault-finding equipment must be provided.	The fault-finding equipment may be controlled by AI functions; however, the parameters of this equipment are defined by a human.
1.7. INFORMATION	
1.7.4.2. Contents of the instructions	
Each instruction manual must contain, where applicable, at least the following information: (...) (g) intend use (k) instruction for putting into service	The data used for AI functions will be part of the general information on the intended use and of the detailed information in the instructions for putting into service.

EXAMPLES OF EXISTING APPLICATIONS OF AI IN THE PRODUCTS AND PROCESSES OF THE TECHNOLOGY INDUSTRIES

The companies represented by Orgalim have been using industrial robots for a number of decades already. Whereas in the past these robots comprised automated arms that performed tasks considered too tedious or dangerous for humans, recent technological developments allow for the wider use of robots in the workplace, with or without interaction with workers. The evolution of these technologies has been accompanied by a sharp rise in standardisation activities at ISO level (TC 299). To enable a common understanding, a clear distinction is made between two main types of robots:

Industrial robots: robots that are automatically controlled, reprogrammable (2.4), multipurpose (2.5), manipulator (2.1) programmed in three or more axes[1], which can be fixed in place or mobile for use in industrial automation applications. These robots operate **autonomously or semi-autonomously**.

Service robots: robots that are programmed to perform useful tasks for humans or equipment, but excluding industrial automation applications. These robots normally operate in the consumer and public environment, autonomously or semi-autonomously.

[1] See ISO 8337:2012 definition: https://www.iso.org/obp/ui/ - iso:std:iso:8373:ed-2:v1:en:term:2.1

Examples of industrial robots:

- Purpose-designed robot systems performing complex specific tasks such as assembling, welding, painting, ironing, assembly, pick and place, palletising, etc.
- Robots for product inspection (e.g. detecting product defaults), testing and removing products from the production chain at various stages before they are further processed, assembled and packaged for distribution
- Agricultural machinery equipped with cameras that take pictures of, for example, lettuce to determine the precise amount of fertilizer to apply to each lettuce head
- AI systems interpreting video feeds from drones carrying out visual inspections of infrastructure such as oil pipelines, fires in buildings, windmills etc.
- Automated guided vehicles (AGVs) used in warehouses for shelf movers and fork-lift trucks

Industrial robots may or may not have features for collaborative use. For instance, the 'co-bots' that interact with human co-workers are a sub-group of industrial robots[2].

Service robots are programmed to perform applicative tasks in the professional, consumer and public environment, autonomously or semi-autonomously. Examples include:

- Autonomously guided vehicles such as robotic lawnmowers or postman letter caddies
- Cleaning robots for public places
- Delivery robots in offices or hospitals; fire-fighting robots
- Rehabilitation robots or surgery robots in hospitals
- Robots guiding customers to the item they seek in a shop
- Dirt-detecting and cleaning robots (vacuum cleaners)

Automated systems – AI software

- Electricity grid management systems that prevent a power grid overload or black-out
- Program debugger
- Prevention of cybersecurity attacks

CONCLUSION

The technology industries represented by Orgalim have been placing on the market machines embedding AI functions for many years, without any negative impact on the safety of users. Despite the rapid evolution of these AI-embedded machines – operating on the basis of narrow artificial intelligence – manufacturers can continue to rely on a robust regulatory framework and the development of a large number of standards (a process which our companies are actively involved in). This will continue to guarantee the placing on the market of safe machines while ensuring a high level of trust among both machine users and consumers for many years to come.

[2] See ISO/TS 15066:2016 definition: https://www.iso.org/obp/ui/ - iso:std:iso:ts:15066:ed-1:v1:en

In conclusion, we firmly believe that the current Machinery Directive 2006/42/EC does not present any legal gaps and that no additional requirements are needed at this point in time. We are therefore convinced that a revision of the Machinery Directive 2006/42/EC is not needed.

Responsible adviser: Eleonora Piccinni