

Vicarious Liability: A Solution to a Problem of AI Responsibility

Dr Deepak Kumar Sahoo^{1*}

^{1*}Former Director, Examinations, Biju Patnaik University of Technology, Rourkela, Odisha (India),
e-mail: dks7173@yahoo.co.in; dks7173@gmail.com

***Corresponding author:** Dr Deepak Kumar Sahoo

*Former Director, Examinations, Biju Patnaik University of Technology, Rourkela, Odisha (India),
e-mail: dks7173@yahoo.co.in; dks7173@gmail.com

Abstract

Who is capable when a man-made intelligence machine makes something turn out badly? Or is there a void in the assignation of blame? The responses can either state that there is a single responsibility gap, multiple responsibility gaps, or none at all. The issue can be summarized as follows: On the one hand, it seems right to hold someone accountable for something that an AI machine did wrong; On the other hand, it doesn't appear that anyone deserves to be held accountable for this error. The study concentrates on a specific aspect of the AI responsibility gap in this article: In cases where AI machines have design flaws, it makes sense that someone should bear the legal costs; However, there does not appear to be such a suitable bearer. The study approaches this issue according to the legitimate point of view and propose vicarious responsibility of computer-based intelligence makers as an answer for this issue. Our proposition comes in two variations: The first one is more limited in scope, but it is simple to incorporate into existing legal frameworks; The second one can be used in a wider range of situations, but it requires legal frameworks to be updated. A broader definition of vicarious liability is used in the latter variant. Finally, study draw attention to the important insights that vicarious liability provides for closing the moral AI responsibility gap.

Keywords: AI Responsibility, Regulation, Responsibility Gap, Vicarious liability.

Introduction

Many (not really human) "hands" causally add to making artificial intelligence machines, comprehensively considered (counting web indexes, independent vehicles, caribous, military robots, etc), how they are. In addition, the implemented operations that control a machine's behavior frequently have a high degree of complexity, making it impossible to predict all possible outcomes.¹ Because of these and other factors, it can be challenging to pinpoint the source of a machine's wrongdoing, be it manipulation, racism, bias, data misuse, or even homicide. As a result, study may have wrongdoing without a clearly identified human wrongdoer—a significant issue for AI ethics and law theory and practice. The beginning stage to examinations this issue is the accompanying general inquiry: is it generally the situation that somebody is dependable when a simulated intelligence machine makes something turn out badly, or is there at times a hole in liability? The responses range from asserting a singular responsibility gap, a number of distinct gaps, or none at all. For example, Matthias (2004), who presents the idea of an obligation hole, guarantees that there is an exceptional hole that has numerous viewpoints and that study probably won't have the option to connect. As a result, a problem arises: Either we should stop using the new technology that created the gap, or we should stop thinking in terms of traditional responsibility.

In contrast, Santoni de Sio and Mecacci (2021) assert that there are four distinct responsibility gaps: an active responsibility gap, a culpability gap, a moral accountability gap, and a public accountability gap. However, not everyone agrees that such gaps exist: for example, Kohler et al. al. (Both Tigard (2020) and 2017 assert that there is no "techno-responsibility gap; The former only concentrate on an alleged accountability gap, whereas the latter adopts a pluralistic viewpoint. From a moral and legal point of view, assigning responsibility serves multiple purposes: crediting conduct, accusing miscreants, forcing an obligation to record or deal with any consequences regarding what occurred, rebuffing the blameworthy, forestalling future mischief, repaying casualties, etc.

In this paper, study will zero in for the most part on the lawful viewpoint and on the accompanying comprehension of obligation: If P is a suitable bearer of the legal consequences of a performance or omission of, then a party P is a fitting bearer of responsibility with regard to an action. The reader will be provided with a normative perspective on the issue by reference to fittingness; In other words, study is asking who should bear the normative consequences of the aforementioned

scenarios, and study will provide reasons to support a particular answer. 2 For example, according to our understanding, stating that a party P is a fitting bearer of responsibility with respect to the addition of a defective component in AI machine r means that P should bear the legal consequences when a (possibly different) party Q adds a defective component to r. Likewise, stating that P is a fitting bear Concerning the previously mentioned understanding of responsibility, study would like to make some preliminary remarks. First, we do not get into the subject of a conceptual analysis of the concept of responsibility because study used a working definition of responsibility that is sufficient for the purposes of this article. Second, according to this understanding, study assign responsibility to "actions" (see the parameter above, which technically corresponds to an action-type whose token-based performances are shown). Thirdly, the legal repercussions frequently involve financial compensation. The most straightforward illustration of responsibility as a duty to bear legal consequences is primary legal liability: it tends to be credited, e.g., when a software engineer, programming designer, or an information expert is careless, crazy, or purposefully commits an offense that outcomes in the artificial intelligence machine accomplishing something wrong. If it is possible to determine which employees committed a wrongdoing, it could mean that they bear all legal responsibility or that they do not. Although they are possible, the first category of cases does not cause us the most concern. This is because each case of this kind has a distinct offender. The latter category of cases presents greater difficulty. It does not make sense, for some reason, that the perpetrator of these crimes should bear all of the legal consequences. They might not be able to compensate the victim enough, for example, or it might not be fair to put such a heavy burden on them because they were doing work for their employer rather than working on their own projects. Whether through product liability, vicarious liability, agency law, or insurance, it is frequently the case that the employer (or insurance company) bears the ultimate legal consequences in an employment relationship rather than the employee.

Besides, as study demonstrated in the initial passage, it is extremely normal not the case that there is a specific transgressor in settings where simulated intelligence machines are created and utilized. This last category of cases is the most challenging because it does not have a human wrongdoer or the appropriate bearer of legal consequences: it simply gives the idea that the man-made intelligence machine is the "transgressor". As a result, study will concentrate on the cases that are more difficult and propose vicarious liability for AI manufacturers as a solution. Two versions of our proposal will be available: The first one addresses situations in which study do not have the appropriate bearer but do not have the human wrongdoer behind a machine's behavior. There are a number of advantages this strategy has over its rivals. To begin, it can be used in any jurisdiction with the current law. Because the suggestion is merely a previously unconsidered application of vicarious liability, it does not necessitate any revision of the legal frameworks, as opposed to approaches that, for instance, suggest granting the AI legal personhood. Obviously, vicarious obligation was considered in the writing, however not in the manner in which it is proposed, i.e., not as a use of the old-style boss representative relationship to connect (a part of) a hole in artificial intelligence obligation. In addition, study will argue that this proposal has advantages over its rivals, such as product liability, which do not alter the legal framework. In addition to the cases covered by the first version of our proposal, the second version of our proposal is designed to cover situations in which it is impossible to prove that a human being is to blame for a machine's behavior.

For example, it could happen that no AI manufacturer employee committed a tort but the AI machine nevertheless caused harm. The second variant also suggests that manufacturers are the bearers of legal consequences, but instead of employees and manufacturers, the central relationship is one between manufacturers and machines. Naturally, this would necessitate a revision of the existing legal frameworks or a broader definition of vicarious liability, but it also addresses cases in which there is no known wrongdoer. Study would be able to fill in some of the potential responsibility gaps in a novel way, whether the first or second variant was chosen. Most importantly, morality cannot be completely separated from the legal analysis of these issues. As a result, it will occasionally talk about moral AI responsibility and emphasize the significance of vicarious liability in bridging a moral AI responsibility gap. This paper draws inspiration from a number of different places, including: Asaro (2012) and Turner (2019), who examine the pertinence of the vicarious risk of proprietors or clients; what's more, Giliker (2010) and Dim (2018), who give extremely smart examinations of the tenet of vicarious obligation - as it has developed after some time and as it presently exists in different purviews. The remaining parts of the paper are arranged as follows: Segment 2 spotlights on the issue of liability gap(s). From a neutral point of view, it explains how the potential gaps in the attribution of responsibility raise pressing questions about who, if anyone, is responsible for harm caused by an AI machine, in what sense they should be held accountable, and under what circumstances.

Segment 3 gives a brief outline of endeavors to respond to these troublesome inquiries. While study don't give an in that frame of mind of these methodologies, it calls attention to a portion of their assets and shortcomings. The central issue of

(potential) responsibility gaps is the focus of Section 4: the question of who, if any, should bear the legal repercussions of undesirable AI-machine scenarios. Manufacturers emerge as appropriate bearers of responsibility in our analysis of this problem, which has two versions. Vicarious liability is a well-established legal principle that is applied, but we also advocate for its expansion in new directions. Our contribution is outlined in Section 5, and the significance of our findings for AI moral responsibility gaps is highlighted.

The Issue(s) of Responsibility Gaps The development of AI machines right now is illustrative of how technological advancement frequently challenges established norms. The majority of participants in the debate would concur that there at least appear to be responsibility gaps, despite our intention to remain non-committal regarding questions regarding the uniqueness and existence of responsibility gaps related to the use of AI machines. Even though it seems appropriate to hold someone accountable in certain AI-machine scenarios, it is unclear who, if anyone, is the appropriate bearer of that responsibility. Matthias (2004, 176) contends that the affirmation of an obligation hole questions the actual utilization of new innovation, which "definitely lead[s] to a halfway loss of the administrator's command over the gadget", except if it is ready to surrender our customary ideas of obligation (which, as per Matthias, require control; cf. as well as Coeckelbergh 2020a and 2020b, which discuss the significance of control in AI responsibility).

An extreme view is examined by Gunkel (2020), who examines ongoing developments like independent machines, learning calculations and social robots. In fact, according to Gunkel, these advancements are not entirely compatible with an instrumental view of technology, which holds that machines are merely means to an end that is determined by humans. As indicated by Gunkel, self-learning gadgets might have the option to decide their own finishes. In correlation, Dignum (2019, 21) presents a safer view, keeping up with that, essentially at the ongoing phase of mechanical turn of events, via "independence" it typically mean undertaking independence, not objective independence: AI can figure out how to achieve a goal on its own, but it does not set its own goals (for example, an autonomous vehicle may choose how to get you to your destination based on how it is programmed, but it does not choose where to take you). The logic that it wish to seek after here worries the accompanying testing and interlaced questions:

- (Q1) Who (what individual or group) is a definitive conveyor of obligation when the causal specialist (a contributing party) is a man-made intelligence machine?
- (Q2) What kind(s) of obligation can conceivably be attributed when the causal specialist is a simulated intelligence machine?
- (Q3) Under what circumstances is it reasonable to assign a specific type of responsibility to a specific party?

Study will respond to (Q1) and (Q2) with the intention of bridging potential responsibility gaps by suggesting which parties must bear legal consequences for an AI machine's undesirable behavior. In addition, based on our analysis of the concept of responsibility that study will propose as a response to (Q2), study will draft an answer to (Q3). Before doing that, study will go over a number of approaches to the issue of AI responsibility that have been proposed in the literature (section 3), noting their relevance to the goal of this work and highlighting their advantages and disadvantages. In the context of AI, responsibility naturally raises a plethora of intriguing additional questions. While the study is unable to cover all of these, allow us to briefly discuss a few. To begin, it is difficult to predict whether each of the possible types of AI machines will yield a single response to questions (Q1) through (Q3). Another issue concerns setting: Will our responses to these questions differ based on whether we are in a state of peace, a state of war, a military coup, or even a health crisis? There may not be a one-size-fits-all solution that addresses all AI machine types and (normative) scenarios.

Proposed Approaches to AI Responsibility

The issue of responsibility in the context of AI is the subject of numerous proposals in the literature. Most of the time, these proposals answer questions 1 and 2. Q3, on the other hand, is relatively understudied. Regarding (Q1), a variety of parties have been considered as potential bearers of responsibility, including programmers, engineers, policymakers, lawyers, philosophers, politicians, governmental bodies, hackers, users, operators, those who were using the robot or supervising its use, those who were not using the robot but otherwise came into contact with it, owners, and even AI machines or robots themselves (2017). Concerning (Q2), the majority of the research focuses on who is responsible for bad outcomes (scenarios in which an AI machine "does" something wrong).

This makes sense because whenever something goes wrong, we have to figure out how to handle it. Debates regarding AI responsibility naturally begin with the concept of the user's individual (contributory) responsibility. ³ For instance, Loh and Loh make the observation that "as long as there is something like a driver in autonomous cars, she will have to bear

responsibility for making decisions that are morally relevant that are not covered by traffic rules" (Loh and Loh, 2017, 45). Hence, for this situation the response to (Q1) is 'clients of computer-based intelligence machines', while the solution to (Q2) isn't direct. In point of fact, the extent to which a human user can bear contributory responsibility is both limited and possibly the least problematic. Some proposals offer novel interpretations of key concepts involved in a particular context in order to overcome the limitations imposed by individual responsibility.

For instance, Gurney (2017) proposes giving manufacturers driver responsibilities for autonomous vehicles (they would be treated as the drivers of the vehicles they have manufactured and would be held responsible for damages caused by such vehicles, which is an answer to Question 1). According to him, "one does not have to be in 'actual physical control' of the vehicle to be a driver," and he argues that the definition of a driver (Gurney, 2017, 60) permits this: It is anticipated that the autonomous vehicle's computer program will perform all traditional driving functions, including steering, accelerating, and braking. Subsequently, from a definitional viewpoint, nothing keeps makers from being viewed as drivers of the vehicles. He asserts that many autonomous vehicle manufacturers are willing to accept responsibility for their work, which may indicate that the proposal will be accepted by the responsible party. A benefit of this proposal over product liability, which study will discuss in a moment, is outlined by Gurney (2017, 59). By treating the manufacturer as the driver, what would have been difficult under products liability becomes straightforward under negligence. As a result, a procedure that is comparable to the standard one could be used to investigate potential legal cases. Also, this idea is somewhat fair, since makers and software engineers benefit from creating and selling artificial intelligence. The manufacturer is Gurney's response to (Q1).

However, Gurney (2017, 62) makes the erroneous assumption that the manufacturer and the programmer are the same entity, which simplifies matters. Without this problematic assumption, the account's operation is unclear. It is unclear which of these parties ought to be the driver in the end because the manufacturer is almost never the same party as the programmer (and the programmer is usually a group of people working on an AI machine). This demonstrates that the outcomes of a similar procedure are not entirely attainable. Negligence, which is a type of primary liability, is his response to (Q2) Also, think about the scenario in which a fatal accident was caused by a software error in an autonomous vehicle. A manufacturer or a programmer will be seen as personally liable for this fatal accident if the proposal is accepted. In any case, assuming that were actually thus, they may be criminally dependable. It is true that these parties may accept a fine, provided that the benefits outweigh the costs and the business continues to be profitable for them. However, if the involved party is treated as if they were the driver, they could face prison time for causing a fatal accident. This would probably be unacceptable to the involved party and discourage the development as a whole.

The concept of products liability, also known as "a doctrine that imposes liability on a product manufacturer for harm caused by a defect in its product," is the foundation of other proposals (Gurney, 2017, 54). Subsequently, additionally for this situation one can pay all due respects to (Q1) by highlighting makers and to (Q2) regarding items risk. Since, in the end, their product is to blame, such an attribution seems reasonable. In addition, there are numerous measures that can be taken to minimize product liability, so this option should be acceptable to the suggested bearers (manufacturers outside of the AI sector bear this kind of liability for their products, so it is nothing new). In the context of autonomous vehicles, Wu (2016) says that manufacturers can manage product liability risk through careful planning, a strong commitment to safety, an effective risk management process that starts with a thorough risk analysis, adhering to international standards, obtaining robust insurance coverage, working with other manufacturers, pre-litigation legal strategies, and efficient records and information management practices. Be that as it may, as Cart contends, the choice isn't exactly great as it looks. This is due to the fact that court proceedings would be difficult, time-consuming, and costly (for example, a number of experts would need to be consulted to offer their opinions on the subject and demonstrate whether the defect in question could have been foreseen and avoided), making this option relatively inefficient: According to Gurney (2017, p. 62), "products liability does not provide a cost-effective means of ensuring that the victims of accidents involving autonomous vehicles are adequately compensated."

Study concur that this is a significant drawback from a practical standpoint. Turner (2019) talks about two more challenges. First, it's not clear whether AI can be viewed as a service or a product. According to Chesterman (2021, p. 95), Surprisingly, it's still not clear if software is considered a product. Also, items obligation depends on "the supposition that the item doesn't keep on changing in a flighty way whenever it has left the creation line. This paradigm does not apply to AI" (Turner, 2019, 98). The concept of robot responsibility, or the moral or legal responsibility ascribed to the machine itself, is even at the heart of some proposals to address the issue of AI responsibility. Hakli and M'akel'a (2019) contend against the chance of crediting liability to robots (or to artificial intelligence machines overall). The following is their main

argument (Hakli and Makela, 2019, p. 269): the "cause history of robots inevitably causes a problem for their autonomy; robots are not and won't be fit to be considered ethically dependable in light of the fact that they are planned, constructed, and customized by different specialists to have the 'character' they have".⁴ White and Baum consider lawful robot obligation a practical choice given that "legitimate risk follows straightforwardly from lawful personhood" and legitimate personhood is certainly not an implausible opportunities for robots, since other non-human elements (like companies) are viewed as lawful people too (White and Baum, 2017, 70).

Nonetheless, they know that much relies upon the motivation behind crediting responsibility to robots (assuming the design is revenge, it would look bad to attribute it to robots without cognizance; on the off chance that it is counteraction, it would check out even without cognizance) and on how study answer the inquiry whether the given simulated intelligence machine represents a devastating gamble. According to White and Baum (2017), 71–75, if it does, the current law may not be sufficient to cover such situations. Study presume that machine responsibility is a contentious choice. The proposal of a mandatory insurance policy or tax could make it easier to impose strict liability on the AI machine owners. Accordingly, this proposal addresses (Q1) by pointing to owners (and insurers) and (Q2) by at least assigning them the obligation to compensate for damages (see Hevelke and Nida-R umelin's 2015 discussion).

Despite being appealing, this option does not address many issues. Regarding (Q3), if the option is fair, we must explain why only owners are liable in this manner when many other parties participated in the same risky practice, at least in the sense that their contribution made the practice possible (manufacturers, programmers, distributors, users, etc.). It should also be made clear whether the machine's version matters (should one be strictly liable to a lesser extent if they pay more for a safer machine?). Due to the fact that owners of AI machines cannot anticipate all possible negative outcomes, it is also unclear whether the option would be acceptable to them. Even if it only partially covers the gap(s) in responsibility, mandatory insurance might still be a viable option.

The Proposal

In this section, Study will present our plan the closing a particular (part of the AI responsibility gap), which entails determining the appropriate holder of legal consequences in cases involving AI machines with design flaws. At the point when study discuss a "plan imperfection", here, it is implied that the manner in which the machine has been customized (programming deformity) or collected (equipment imperfection) considers the chance of hurting some party, regardless of whether this can't be anticipated. Our proposal will, for the sake of illustration, concentrate on software defects. The term "vicarious responsibility," which is used less frequently in legal writings but more frequently in moral philosophy, will be used in the proposal. Our proposal will be offered in two variations, as stated in the opening section. The first aims to cover instances in which an AI manufacturer employee (the employer) commits a crime. Because it merely suggests that there is an option within the existing legal frameworks that has not been fully explored in the literature, this variant is conservative in relation to the existing legal frameworks.

In particular, as we have seen, the possibility of assigning AI manufacturers vicarious liability as employers has neither been considered nor argued for, despite the fact that some people considered vicarious liability to be one of the options. The thought is to treat those engaged with the advancement of computer-based intelligence machines as workers of the computer-based intelligence maker, and to just apply the precept of vicarious obligation by means of business representative relationship. As a result, there won't be any real division of labor in situations like these because there will be an appropriate bearer of this kind of responsibility. The second variant aims to include situations in which this is not the case, such as those in which the AI machine itself appears to be the "wrongdoer." In order to arrive at a plausible analysis of these cases, study will explain how the vicarious liability doctrine can be altered. Since study will utilize the regulation of vicarious obligation in the two variations, our methodology will be like that utilized by Asaro (2012) and Turner (2019), however it will lead us to an alternate class of liability carrier (to the maker as opposed to the client or proprietor). As we will see, vicarious liability is a flexible theory that changes with the times and the world around it. In this segment, study will initially frame the tenet of vicarious responsibility, then, at that point, create and persuade the principal variation of our proposition, lastly create and rouse the subsequent variation. The centuries-old respondent superior doctrine, which holds that the master is responsible for their servant's wrongdoings and assumes control over the servant, is the foundation of vicarious liability. Over time, the doctrine has changed; customarily investing heavily of spot with regards to the business representative relationship, its degree has been reached out to cases past work. First, the rise of highly skilled workers eliminated the need for the employer to have control over the employee, which was a necessary condition for vicarious liability (Magnet, 2015).

The conditions of modern employment are slowly forcing the detachment of vicarious liability from the idea that an employer can only be held liable for the actions of their employees, as it is becoming increasingly difficult to distinguish employees from independent contractors and borrowed workers (Brodie 2006, Giliker 2010, 73–77, 83–93). Lastly, the idea that an employer can only be held liable for actions done for their benefit and with their (implicit or explicit) authorization was divorced from vicarious liability after an increasing need

As previously stated, AI manufacturers typically employ programmers and other parties involved in machine design development (for simplicity, study will concentrate on programmers). As a result, a programmer can be held personally liable for harm caused by negligence or even intentional misconduct that falls within the scope of employment. However, it does not always seem reasonable to anticipate that the programmer will be held liable for all of the harm's legal consequences. It's possible that the programmer won't be able to compensate the victim in full; in that case, the victim should expect compensation from a different party. More importantly, the actions—or lack thereof—of the programmer are within the scope of their employment: in the pertinent cases, developers work for the reasons and the advantage of the artificial intelligence maker (cf. Gray (2018)), who, based on our interpretation, is the employer. That is, the programmer was carrying out tasks that were given to them by their employer, not their own. Typically, this suggests the vicarious liability doctrine rather than primary liability. However, if this is the case, it is unfair to place a heavy burden on the programmers (employees) to deal with legal consequences. Consequently, regardless of whether they were to be viewed as liable for the machine's "direct" at first, the maker would thusly bear liability regarding some unacceptable in which the developer partook.

In fact, a comparative contention can't be formed for clients or proprietors. For instance, in the event that a user abuses the robot, the manufacturer may be held accountable for providing an environment conducive to such abuse (in the sense of posing a risk). However, it would appear that rather than a case of vicarious liability, it is individual responsibility on the part of the owner (for such misuse) and the manufacturer (for creating the possibility of such misuse). Even though none of this is conclusive, contacting manufacturers appears to be the most promising approach. This is our underlying legitimization for replying (Q1) by means of artificial intelligence makers and replying (Q2) through the principle of vicarious obligation. There will be two versions of our proposal, as previously stated. Allow us to begin with the first, which plans to cover situations where a worker of the computer-based intelligence maker (who thus plays the job of an employer) is a miscreant. Study suggest that the existing vicarious liability doctrine can be conservatively applied in these situations. The next step is to assume this kind of liability for mistakes made by the programmer—their employee—by the manufacturer, or employer. It is true that it is notoriously difficult to specify vicarious liability's necessary and sufficient conditions (there are a number of judicial reasoning tests that are of limited assistance). The employee's misconduct must fall within the scope of employment, which is a crucial requirement. If straightforward examples are taken into consideration, it should not be difficult to intuitively comprehend this condition.

For instance, in the event that a developer hassles a lady in his home beyond working hours, this plainly doesn't fall inside the extent of being utilized as a manufacturer. However, he clearly falls under the scope of his employment if, through negligence, he causes a wall to collapse or accidentally injures someone at work. There are many ways to interpret the scope of employment: For instance, an action that results in harm may fall under the purview of employment if the employer explicitly or implicitly commanded or authorized it, if the action is closely related to the duties performed by the employee, or if the risk of harm is unrelated to employment. For example, envision that a software engineer of a kid caring robot executes a standard which permits the robot to hurt a youngster's pets or that such a developer disregards carrying out a standard which permits the robot to identify hazardous items nearby a kid and to keep the kid a long way from these items (comparative guidelines can be either as logical standards or through learning capabilities).

In such a case, assuming this prompts a damage the software engineer is a miscreant, since excluding or adding the significant order is an activity that can be viewed as firmly associated, or coincidental to the business. Based on the idea that there are a set of deontically relevant actions within the scope of an employment relationship, study offer a first schema for ascriptions of vicarious liability in the following section. The definition of an employment relationship's scope is left open by the schema; However, in this sense, study assume that an action is only deontically relevant if it is connected to a legally protected interest. To put it another way, a deontically relevant action is at the very least mentioned (as a prohibition or obligation)⁶ in some legal document pertaining to the contested employment relationship. This means that actions that are only allowed or required by a particular AI manufacturer but are not mentioned in any legal document governing the production of AI machines in general (like not wearing clothes of a certain color at work) are not considered to be deontically relevant. Study also wants to point out that the first two conditions of the schema below can be changed

depending on the kind of employee behavior that a particular jurisdiction requires for vicarious liability to apply: Vicarious Liability of an AI Manufacturer for the Behavior of a Programmer: A manufacturer m of an AI machine is responsible for the actions of a programmer p if and only if

- i. There is an action α which belongs to a set of deontically relevant actions X ;
- ii. Either p performed α , which is prohibited, or p omitted α , which is obligatory;
- iii. The set X falls within the scope of the relation E of employment between m and p .

To begin, the potential efficacy with which this proposal could address apparent responsibility gaps makes it appealing. All things considered, the vicarious responsibility convention has long made due in regulation unequivocally because of its effectiveness in remunerating casualties. As a result, an easier legal analysis based on vicarious liability can take the place of a difficult case whose legal analysis is based on the products liability doctrine in this instance. Another benefit is that the party that ought to bear lawful outcomes will commonly be a party that can do so. According to Gurney, 2017, 57, manufacturers may accept the proposal because, in the case of autonomous vehicles, they are prepared to take responsibility for their machines. What's more, note that vicarious responsibility typically doesn't reach out to criminal risk (the suit is normally worried about monetary pay).

If it has somehow happened to credit driver hood to makers/software engineers of a self-driving vehicle, it would be indistinct whether risk would stretch out to criminal responsibility. Yet, for what reason should instances of lethally careless driving where the maker driver is to blame outcome just in monetary pay? Because it acknowledges that the relevant action was not the manufacturer's, the current proposal does not encounter these issues. However, as the programmer was acting within the confines of their employment, it seems reasonable to hold the manufacturer accountable as an employer. There are different potential legitimizations that can be utilized to make sense of why this proposition is fair and hypothetically sound. For instance, the close connection can be used to support it: by claiming that the programmer's actions were unrelated to their job (were closely related to it). On the other hand, as far as big business risk, the idea can be legitimate as follows: the maker produced specific dangers that accompanied advantages to their endeavor; The manufacturer should bear the legal repercussions if the risks occur.

There would appear to be a contemporary pattern toward extending the precept of vicarious obligation to represent the requirement for casualty insurance and to mirror the states of present day business and an undeniably innovative world. Study will now concentrate on a more contentious option, despite the fact that the anticipated strategies for extending the doctrine may involve independent contractors or borrowed workers: substituting the employee for the AI machine. This is not nearly as far-fetched as it might appear at first. For example, White and Baum, who, as mentioned earlier, suggested that robots might be held liable in court, make a fascinating comparison between the parent-child relationship and the manufacturer-robot relationship. The gatherings that impact the robot's activities (White and Baum, 2017, 72) can incorporate the robot's fashioner, its producer, and any clients or administrators it might have. These gatherings are similar to a human's folks and managers, however the examination is blemished because of fundamental contrasts among people and robots. A center contrast, they propose, is that "creators have considerably more command over the possible person of robots than of people. This recommends that robot creators ought to hold more risk for robots' activities than human guardians ought to for their kids' activities" (White and Baum, 2017, 72). Note that control is here again connected to liability. Study believe that the manufacturer cannot be viewed as having control over the AI machine, so we must ask whether it makes sense to apply the vicarious liability principle to the manufacturer-AI machine relationship.

If the scope of employment is traded for manufacturing (which is what the machine is designed for), and if we grant the AI machine sufficient agency to feature as the party who performed or omitted a relevant action, study believe it is still a sensible option to pursue. Following Hyman, it is accepted that the typical decrease of activity to human activity is inappropriate. (Hyman, 2015, 43): "A long way from applying only to individuals, the idea of activity applies to each substance ready to cause change." Because AI machines are complex inanimate agents, the philosophical investigation of AI responsibility gaps was ultimately motivated by their capacity to alter the world—particularly to harm. As a result, the idea would be to assign manufacturers (Q1) vicarious responsibility for mistakes made by AI machines, at least as long as the mistakes happen while the machine is doing things for which it was designed (or that are closely related to the original purpose). To respond to (Q3), it must rework the previous scheme as follows: M is the relationship of being a manufacturer, r is a robot (an AI machine), m is its manufacturer, and Manufacturer's Vicarious Liability for an AI Machine: A producer m of a computer-based intelligence machine r is vicariously obligated for the way of behaving of r if and provided that

- (i) There is an action α which belongs to a set of deontically relevant actions X ;
- (ii) Either r performed α , which is prohibited, or r omitted α , which is obligatory;

(iii) The set X falls within the purpose of the manufacturing relation M between m and r .

Let's use two scenarios to illustrate the condition of the purpose of manufacturing: one in which the relevant action is included in the purpose of manufacturing and another in which it is not. Envision that, because of a plan imperfection, a man-made intelligence lawnmower hurts a kid (or a canine, or a valuable rose, as in Turner 2019, 85) while cutting the grass. In this particular instance, the lawnmower causes harm to a child by cutting the grass, which is one of its manufacturing processes. On the other hand, suppose a user or a hacker reprogrammes an AI kitchen robot to mow the lawn. The robot, which was made to stay inside, goes outside to the garden, tries to cut the grass, and it ends up doing the same damage as the original lawnmower. For this situation, the activity performed doesn't fall inside the reason for assembling. As a result, the manufacturer would not be held vicariously liable in the second scenario. The second proposal inherits a significant amount of the first proposal's discussion: It is appealing due to its potential acceptability, fairness, and efficiency. This is on the grounds that the two recommendations have a typical center: The term "wrongs caused by machines with design defects" can be used to rephrase nearly all, if not all, of the errors made by the machine's designers. This indicates that under both proposals, the outcomes—acts or omissions for which the manufacturer would be vicariously liable—would be comparable in those instances.

The breadth of the two versions of our account is a clear distinction. The latter proposal is more comprehensive because it includes situations in which there is no human wrongdoer. In addition, it has greater ambitions because it aims to cover a subset of philosophically challenging cases. This also indicates that the proposal is more comprehensive than, for example, Gurney's account. This is because Gurney's account seems to only cover situations in which the manufacturer negligently operated the machine. As previously mentioned, the status of AI machines is yet another significant distinction between the two variants: They are not included in the first proposal because the agency of the employee is what matters, and AI machines are regarded as defective products (or services). According to Turner (2019), the second proposal gives AI machines agency—even if we don't consider them to be negligent, intentional, or at fault—which may go along with giving them legal personhood.

Concluding Remarks: The Liable Party Not Necessarily to Blame

This paper has inspected a specific (part of the) simulated intelligence obligation hole comprising in finding a fitting conveyor of lawful outcomes in situations that include a computer-based intelligence machine with a blemished plan. Study have proposed two different ways of managing the issue.

The principal variation of our proposition is somewhat unobtrusive: It merely suggests that the existing legal principle of vicarious liability be applied to AI machine manufacturing businesses. The programmer is viewed as an employee and the manufacturer as an employer. Programming (or all the more extensively, planning) man-made intelligence machines falls inside the extent of business. Study can say, recalling a previous quote from Gurney, that treating the manufacturer as the employer makes what would have been a complicated matter under products liability a simple matter under vicarious liability. Harms that occur as a result of design flaws can then be dealt with as cases of vicarious liability on the part of the manufacturer, insofar as the harm at issue can be traced back to the violation of a prohibition/obligation by programmers (or other categories working on the design of the machine). Inside this image, artificial intelligence machines are not seen as specialists and they don't actually go into the image. Quite, the primary variation of our proposition can be utilized to cover a more extensive scope of situations than it has zeroed in on by and by - plan deserts as well as any assembling surrenders, since different gatherings who add to assembling man-made intelligence machines are regularly representatives of the producer.

The second variant of our proposal is based on Asaro and Turner's idea that AI machines can also enter the picture, possibly in the form of animals or young children, and Hyman's plea that agency should not be reduced to that of humans. An AI machine assumes the role of the agent who contributed to the error rather than the programmer. However, in contrast to Asaro and Turner, study contend that the manufacturer, not the owner or user, should be the other party in this relationship. The justifications of the doctrine of vicarious liability in the context of the relationship between an employer (a manufacturer) and an employee (a programmer) carry over to the context of the manufacturer–AI machine relationship when it comes to the manufacturer as the intended bearer of responsibility. While the focus of this paper was on the legal side of AI liability, study believe that a parallel strategy can be used in morality.

The significant thought from moral way of thinking then, at that point, would be moral vicarious obligation, which is a sort of severe moral obligation. Although understudied in theory, moral vicarious responsibility is prevalent in practice.

Simply think how frequently is one approached to make up for a minor harm done by their kid or pet. And even when the matter could have been pursued legally, it is frequently resolved solely on moral grounds.

The search for appropriate bearers of moral consequences is the focus of the concept of moral vicarious responsibility. It is essential to keep in mind that, like legal vicarious responsibility, moral vicarious responsibility is independent of the blameworthiness and fault of the party liable. Moral vicarious obligation is in this way worried about bearing moral outcomes, not really with accusing individuals. Put in an unexpected way, producers can be the appropriate conveyors of vicarious obligation regarding the machine's way of behaving without being in any sense accountable for what occurred. Relevant moral consequences in this instance may include, as is typically the case in law, a (moral) obligation to provide the harmed party with financial compensation in addition to an apology and an explanation of the circumstances surrounding the incident. Indeed, it is frequently observed that AI manufacturers apologizing for their machines' mistakes. This would imply that such a moral practice already exists, and our account can assist in explaining why the manufacturers' response is morally acceptable.

References

1. Asaro, P. M. (2012). A body to kick, but still no soul to damn: legal perspectives on robotics. In *Robot Ethics: The Ethical and Social Implications of Robotics*, 169–186.
2. Brodie, D. (2007). Enterprise liability: justifying vicarious liability. *Oxford Journal of Legal Studies*, 27(3), 493–508.
3. Cambridge, MA, London: MIT Press. Brodie, D. (2006). *The Enterprise and the Borrowed Worker*. *Industrial Law Journal*, 35(1), 87–92.
4. Cane, P. (2002). *Responsibility in Law and Morality*. Oxford and Portland: Hart Publishing.
5. Cane, P. (2016). Role Responsibility. *The Journal of Ethics* 20(1-3), 279–298.
6. Chesterman, S. (2021). *We, the Robots? Regulating Artificial Intelligence and the Limits of the Law*. Cambridge: Cambridge University Press.
7. Clarke, R., McKenna, M. and Smith, A. M. (Eds.). (2015). *The Nature of Moral Responsibility: New Essays*. New York: Oxford University Press.
8. Coeckelbergh, M. (2020a). Artificial intelligence, responsibility attribution, and a relational justification of explainability. *Science and Engineering Ethics*, 26(4), 2051–2068.
9. Coeckelbergh, M. (2020b). *AI Ethics*. Cambridge, MA: MIT Press.
10. Dignum, V. (2019). *Responsible Artificial intelligence: how to develop and use AI in a responsible way*. Cham: Springer.
11. Giliker, P. (2010). *Vicarious Liability in Tort: A Comparative Perspective*. New York: Cambridge University Press.
12. Gray, A. (2018). *Vicarious liability: critique and reform*. Oxford: Hart Publishing.
13. Gunkel, D. J. (2020). Mind the gap: responsible robotics and the problem of responsibility. *Ethics and Information Technology* 22, 307—320.
14. Gurney, J. (2017). Applying a reasonable driver standard to accidents caused by autonomous vehicles. In:
15. Hakli, R. and M'akel'a, P. (2019). Moral responsibility of robots and hybrid agents. *The Monist* 102(2), 259–275.
16. Hevelke, A. and Nida-Rümelin, J. (2015). Responsibility for crashes of autonomous vehicles: an ethical analysis. *Science and Engineering Ethics*, 21(3), 619–630.
17. Howard, C. (2018). Fittingness. *Philosophy Compass* 13: e12542.
18. Hyman, J. (2015). *Action, Knowledge, and Will*. Oxford: Oxford University Press.
19. Isaacs, T. (2011). *Moral Responsibility in Collective Contexts*. Oxford University Press, USA.
20. Kohler, S., Roughley, N. and Sauer, H. (2017). Technologically blurred accountability? Technology, responsibility gaps and the robustness of our everyday conceptual scheme. In:
21. Lin, P., Abney, K. and Jenkins, R. (eds.), *Robot Ethics 2.0*, 51–65, New York: Oxford University Press.
22. Loh, W. and Loh, J. (2017). Autonomy and responsibility in hybrid systems. In: Lin, P., Abney, K. and Jenkins, R. (eds.), *Robot Ethics 2.0*, 35–50, New York: Oxford University Press.
23. Magnet, J. (2015). Vicarious liability and the professional employee. *Canadian Cases on the Law of Torts* 6, 208–226.
24. Matthias, A. (2004). The responsibility gap: ascribing responsibility for the actions of learning automata. *Ethics and Information Technology* 6, 175— 183.
25. Murphy, J. (2004). The merits of *Rylands v Fletcher*. *Oxford Journal of Legal Studies* 24 (4), 643–669.

26. Narveson, J. (2002). Collective responsibility. *The Journal of Ethics*, 6(2), 179- 198.
27. Routledge. Lin, P., Abney, K. and Jenkins, R. (Eds.). (2017). *Robot ethics 2.0: From autonomous cars to artificial intelligence*. New York: Oxford University Press.
28. Santoni de Sio, F. and Mecacci, G. (2021). Four responsibility gaps with artificial intelligence: Why they matter and how to address them. Online first in *Philosophy & Technology*.
29. Sullins, J. P. (2011). When is a robot a moral agent. In Anderson, M., and Anderson, S. L. (eds.), *Machine ethics*, 151-161, New York: Cambridge University Press.
30. Tigard, D. W. (2020). There is no techno-responsibility gap. Online first in *Philosophy & Technology*.
31. Turner, J. (2019). *Robot Rules: Regulating Artificial Intelligence*. Cham: Palgrave Macmillan. White, T. N. and Baum, S.D. (2017). Liability for present and future robotics technology. In: Lin, P., Abney, K. and Jenkins R. (eds.), *Robot Ethics 2.0*, 66–79, New York: Oxford University Press.
32. Ulbert, C., Finkenbusch, P., Sondermann, E. and Debiel, T. (eds.), *Moral agency and the politics of responsibility*, 51-68,
33. Wu, S. S. (2016). Product liability issues in the US and associated risk management. In: Maurer, M., Gerdes, J. C., Lenz, B. and Winner, H. (eds.), *Autonomous Driving*, 553–569, Berlin, Heidelberg: Springer.