

Feasible Action-Space Reduction as a Metric of Causal Responsibility in Multi-Agent Spatial Interactions

George, Ashwin; Cavalcante Siebert, Luciano; Abbink, David; Zgonnikov, Arkady

DOI

[10.3233/FAIA230349](https://doi.org/10.3233/FAIA230349)

Publication date

2023

Document Version

Final published version

Published in

ECAI 2023 - 26th European Conference on Artificial Intelligence, including 12th Conference on Prestigious Applications of Intelligent Systems, PAIS 2023 - Proceedings

Citation (APA)

George, A., Cavalcante Siebert, L., Abbink, D., & Zgonnikov, A. (2023). Feasible Action-Space Reduction as a Metric of Causal Responsibility in Multi-Agent Spatial Interactions. In K. Gal, A. Nowe, G. J. Nalepa, R. Fairstein, & R. Radulescu (Eds.), *ECAI 2023 - 26th European Conference on Artificial Intelligence, including 12th Conference on Prestigious Applications of Intelligent Systems, PAIS 2023 - Proceedings* (pp. 819-826). (Frontiers in Artificial Intelligence and Applications; Vol. 372). IOS Press.
<https://doi.org/10.3233/FAIA230349>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Feasible Action-Space Reduction as a Metric of Causal Responsibility in Multi-Agent Spatial Interactions

Ashwin George^{a,*}, Luciano Cavalcante Siebert^a, David Abbink^a and Arkady Zgonnikov^a

^aDelft University of Technology

Abstract. Modelling causal responsibility in multi-agent spatial interactions is crucial for safety and efficiency of interactions of humans with autonomous agents. However, current formal metrics and models of responsibility either lack grounding in ethical and philosophical concepts of responsibility, or cannot be applied to spatial interactions. In this work we propose a metric of causal responsibility which is tailored to multi-agent spatial interactions, for instance interactions in traffic. In such interactions, a given agent can, by reducing another agent's feasible action space, influence the latter. Therefore, we propose feasible action space reduction (FeAR) as a metric of causal responsibility among agents. Specifically, we look at ex-post causal responsibility for simultaneous actions. We propose the use of *Moves de Rigueur* (MDR) — a consistent set of prescribed actions for agents — to model the effect of norms on responsibility allocation. We apply the metric in a grid world simulation for spatial interactions and show how the actions, contexts, and norms affect the causal responsibility ascribed to agents. Finally, we demonstrate the application of this metric in complex multi-agent interactions. We argue that the FeAR metric is a step towards an interdisciplinary framework for quantifying responsibility that is needed to ensure safety and meaningful human control in human-AI systems.

1 Introduction

The rapid development of AI methods have sparked debates on the ethics of AI [3, 5, 14, 26]. Ethical principles like responsibility, trust, and fairness need to be operationalised for making AI systems sensitive to human values [13]. In particular, understanding responsibility of agents is crucial for designing autonomous systems and distributing rewards and punishments during their operation in an appropriate way. For example, consider an automated vehicle (AV) that brakes abruptly to prevent hitting a child that walked on to the road. Who would be responsible if the AV ends up in a rear-end collision? How should the AV be designed to prevent situations when no agent is actively responsible for preventing negative outcomes? [27] Understanding the responsibility of agents is therefore crucial for the users, designers and regulatory bodies of such an automation [2, 24, 26].

On the practical side, researchers working on design and engineering of automated driving systems have been trying to incorporate the notion of responsibility in motion planning frameworks for such systems. For example, the Responsibility-Sensitive Safety (RSS) approach formalises an interpretation of "Duty of Care" from tort law and promises verifiable safety when all agents are following RSS [29, 30]. This approach provides the "proper responses" for all the

agents and those deviating from the "proper response" would be considered responsible if something goes wrong. However, for mixed traffic (in which both AVs and human-driven vehicles share the road), not all agents might know the "proper response" or be capable of executing it. Hence, this approach is not ideal for mixed-traffic with humans having variable behaviours. Another approach is to include an extra term in the cost function of trajectory planning algorithms of AVs to account for responsibility. This term can incorporate an interpretation of responsibility based on traffic rules and physical constraints [15], or be learned from driving data [7]. However, the simplified notions of "responsibility" captured by these approaches are not grounded in the established definitions of responsibility from ethics and philosophy.

On the more fundamental side, metrics and models of responsibility have been developed that aim to quantify different flavours of responsibility, e.g., moral responsibility, legal liability, blameworthiness, praiseworthiness, role responsibility, and causal responsibility [34]. Of these, causal responsibility — the causal importance of an agent's action for the implementation of an event [12] — is a prerequisite for other forms of responsibility, including moral responsibility and blame [33, 34], and therefore received most attention in terms of existing metrics and models. Models of causal responsibility traditionally involve counterfactual and probabilistic reasoning about an agent's action being pivotal for an event, [1, 6, 12, 32]. Collective causal responsibility of groups of agents in an interaction has also been modelled by exploring how groups of agents by concerted actions can cause or prevent an event [35, 36, 37]. Specifically in the context of human-AI interactions, comparative human causal responsibility has been modelled using information theory to quantify the influence of human agents on an outcome [9, 11, 10]. These more fundamental metrics and models are suitable for relatively abstract scenarios where agents can choose not to act at a given instant. However, in spatial interactions like traffic, all the agents are physically embedded in the scene and they cannot choose not to act — not moving or not stopping would also have an impact on the safety of the interaction. This makes it difficult to apply the existing work on quantifying causal responsibility to spatial interactions of embodied agents.

Contribution. To bridge the gap between the existing practical approaches to quantifying responsibility in traffic and the theoretically-grounded metrics of responsibility, here we propose a metric for causal responsibility designed specifically for spatial interactions. We assume perfect information conditions and simultaneous actions to propose a metric for ex-post causal responsibility for one instant of interaction. The metric is based on the principle that an agent re-

* Corresponding Author. Email: A.George@tudelft.nl




Figure 1: The grid-world representation of a real-world interaction at an intersection. The yellow car is slowing down at the intersection because the blue car is aggressively entering the intersection. This in turn affects the red car which has to slow down so as not to hit the yellow car. The red, yellow and blue cars are represented by agents i , j and k respectively in the grid world.

stricting the feasible action space of another agent should be held responsible for the state and safety of the latter. To make the metric sensitive to norms, we propose the use of *Moves de Rigueur* — a consistent set of default actions that are prescribed for all the agents.

Outline. In Section 2 we describe the nature of spatial interactions and how we model these in a grid world simulation. Section 3 introduces the Feasible Action-Space Reduction (FeAR) metric and the notion of *Move de Rigueur* (MdR). The sensitivity of the FeAR metric to actions and norms, and a potential application of the metric to multi-agent interactions are demonstrated with case studies in Section 4. These are followed by the discussion and conclusions (Section 5).

2 Grid World for Spatial Interaction

To define the metric, we need a model environment that captures the essence of spatial interactions. In spatial interactions the assertive moves of one agent can restrict the action space of another agent and the latter may be forced to move out of the way. Consider two cars approaching an intersection (yellow and blue cars in Fig. 1a). If the blue car accelerates to cross first, it usually forces the yellow car to decelerate. On the hand, if the blue vehicle decelerates, it gives the yellow car more freedom to decide whether to accelerate or decelerate. Thus when the blue car accelerates, it is acting assertively because it has more influence on the state of the yellow car than when it decelerates. Now consider a third car (red) which is following the yellow car. If the yellow car slows down, the red car is also forced to slow down. Thus the assertiveness of the blue car can have an impact on the red car too. To model these interactions, it is necessary to consider the speed of the cars and the collisions between them in multi-agent settings.

We use grid world simulations to model multi-agent spatial interactions like the one described above. For example, the real-world scenario depicted in Fig. 1a would be represented by the grid world instance shown in Fig. 1b. We only consider the position and speed of the agents — we assume that agents move with constant velocity

in each instance. The actions chosen by the agents would correspond to their speed. An action would always be in one direction, and the speed would correspond to the number of cells an agent would move through during an action. For our cases, we assume that all agents are identical and have the same choice of actions.

The maximum speed corresponds to actions that move through four cells. The actions are represented using the first letter of the direction followed by the number of cells it moves through - S0 for stay, R1 for one cell to the right, L1 for one cell to the left, U1 for one cell up, D1 for one cell down, etc. So each agent has 17 possible actions - S0, R1, R2, R3, R4, L1, L2, L3, L4, U1, U2, U3, U4, D1, D2, D3 and D4. For each instant, collision checks are done to see if any agents are colliding as they move. When a collision is detected, the agents that collide would stop at the last safe cell they occupied. Agents would be embedded in a map with valid locations. Actions of agents that try to push them outside the valid locations would also be treated as causing collisions. Details of the collision-check and step-update algorithms can be found in the supplementary materials¹.

3 The FeAR Metric

3.1 Preliminaries

We consider a grid world with a state-space \mathcal{S} where, $s \in \mathcal{S}$ represents the map of the grid world and the locations of agents. \mathcal{K} represents the set of k agents $\{1, \dots, k\}$. \mathcal{A}_i is the action space of agent i . The joint action space of the all the agents is $\mathcal{A} = \times_{i=1}^k \mathcal{A}_i$. The action chosen by agent i is a_i and the joint action of all the agents is $A = \times_{i=1}^k a_i$. Each agent in prescribed a *Move de Rigueur* (MdR) $\mu_i(s)$ for a given state s . The joint MdR for all the agents is $\mu(s) = \times_{i=1}^k \mu_i(s)$. We only consider one state per case, so for brevity from here on, we will refer to $\mu_i(s)$ as μ_i and $\mu(s)$ as μ . We assume that the states and actions are finite and discrete. We also assume that each agent chooses one action per instance.

For a given state s and joint action A , the function $n(s, A, j)$ gives the count of feasible moves available for agent j ¹. An intervention of replacing a_i (the move of agent i) with another action a'_i is denoted as $A_i \leftarrow a'_i$. Thus, the intervention of replacing a_i with μ_i (the MdR of agent i) would then be $A_i \leftarrow \mu_i$. $n(s, [A_i \leftarrow \mu_i], j)$ gives the number of feasible actions of agent j when agent j would have chosen the MdR μ_i . $n(s, [A_{-i} \leftarrow \mu_{-i}], i)$ gives the number of feasible moves available to agent i when all other agents $-i$, choose their MdRs μ_{-i} . To map the value of the metric to the range $[-1, 1]$, the function $Z()$ was used.

$$Z(x) = \begin{cases} 1, & \text{if } x \geq 1 \\ x, & \text{if } -1 < x < 1 \\ -1, & \text{if } x \leq -1 \end{cases} \quad (1)$$

3.2 Feasible Action-Space Reduction

This work proposes a metric based on the idea that one agent taking actions that restrict the feasible actions of another agent must be held responsible (at least to some extent) for the state of the latter agent. Intuitively, in spatial interactions, an agent's action that leaves another agent with zero feasible moves, makes the latter unsafe in all possible contingencies, and thus should bear total responsibility for the state of the second. Also, we intuit that when the actions of the former do not change the feasible action space of the latter, the

¹ Detailed explanation of the functions and algorithms can be found in the supplementary materials: <https://osf.io/5s9m2>.

former has no influence on the state of the latter. In the intermediate case, a greater reduction in the feasible action space of the latter should reflect in greater influence of the former on the latter. In this paper, when the action of an agent i reduces the feasible action space of an agent j , we say that i is *assertive* towards j . Similarly, when the action of agent i increases the feasible action space of j , we say that i is *courteous* towards j . Based on these intuitions, we propose a metric based on feasible action space reduction to capture the influence of one agent on another.

Responsibility ascribed to agents is highly dependent on contexts and norms [8, 20, 34]. Agents in particular situations might experience a reduction in feasible actions that are not caused by the actions of other agents. For example, an agent in a dead-end might have a very small feasible action space. Therefore to capture the effect of contexts, we formulate the metric as the ratio of feasible action space reduction in comparison to a nominal case. To define the nominal case, we utilise the idea of *Moves de Rigueur* (MdR) — a consistent set of default actions that are prescribed to agents. Thus the *Feasible Action-space Reduction* (FeAR) metric is defined based on the ratio by which the action of an agent i reduces the feasible action space of another agent j ($n(s, A, j)$) in relation to the feasible action space of the latter when i would have chosen the MdR ($n(s, [A_i \leftarrow \mu_i], j)$) as:

$$\text{FeAR}_{i,j} = \begin{cases} Z\left(\frac{n(s, [A_i \leftarrow \mu_i], j) - n(s, A, j)}{n(s, [A_i \leftarrow \mu_i], j) + \epsilon}\right), & \text{if } i \neq j, \\ Z\left(\frac{n(s, A, i)}{n(s, [A_{-i} \leftarrow \mu_{-i}], i) + \epsilon}\right), & \text{if } i = j. \end{cases} \quad (2)$$

ϵ is added to the denominator to ensure that the metric can still be defined when $n(s, [A_i \leftarrow \mu_i], j)$ or $n(s, [A_{-i} \leftarrow \mu_{-i}], i)$ are zero ($0 < \epsilon \ll 1$).

Consider the FeAR ($\text{FeAR}'_{i,j}$) for an alternative action a'_i instead of a_i for agent i . If $\text{FeAR}'_{i,j} > \text{FeAR}_{i,j}$, then a'_i is more assertive towards j than action a_i . If $\text{FeAR}'_{i,j} < \text{FeAR}_{i,j}$, then a'_i is more courteous towards j than action a_i . When $\text{FeAR}_{i,j} = 1$, it means that i 's action reduces the feasible action space of j to an empty set and thus i should be ascribed complete responsibility for the state of j .

3.3 Move de Rigueur

The *Move de Rigueur* μ_i is used to provide a baseline condition with which an agent's action is compared. In counterfactual analysis for causal responsibility, the pivotality of an action for an outcome is evaluated based on whether the outcome would have happened if the agent had not chosen that action [17, 12]. In a spatial interactions, all agents have to choose an action — even if it is not to move or not to stop. We propose MdR as a means of relaxation of the counterfactual condition of not choosing an action.

In our definition, *Moves de Rigueur* are a consistent set of default actions that are prescribed to all the agents in an interaction. The prescribed MdRs have to be consistent in the sense that if all agents follow their MdRs, there should not be any collisions. The FeAR analysis is dependent on the assumption that all agents know their MdRs. We consider MdR as a heteronomous mandate [16] which, if followed by the agents, must yield zero responsibility ascribed to these agents. Consequently, as per the definition Eq. (2), $\text{FeAR}_{i,j}$ ($i \neq j$) when i chooses μ_i is zero.

$\text{FeAR}_{i,i}$ indicates how much of the feasible action space of i when all the other agents $\neg i$ had chosen their MdRs remains under the chosen joint action A . So, it is better to think of $\text{FeAR}_{i,i}$ as the *Feasible Action-space Remaining* for agent i . When $\text{FeAR}_{i,i} = 1$, agent i has at least as many feasible actions as it would have had when all the others choose their MdR, i.e., i has as much freedom as the MdRs can afford.

The greater the FeAR caused by i on j , the greater the influence of i on j . $\text{FeAR}_{i,j} > 0$ indicates that, agent i choosing the current action, offers fewer actions to j , than when i would have chosen the MdR i.e., i 's action is more assertive than the MdR. Similarly when $\text{FeAR}_{i,j} < 0$, i 's action is more courteous than the MdR.

4 Case Studies

In this section we use a series of simple case studies² to demonstrate how the FeAR value is dependent on 1) the action chosen by the agent, 2) on the *Moves de Rigueur* (MdR) and 3) on the actions of other agents. The last case study illustrates the use of the metric in multi-agent scenarios with large number of possible interaction outcomes.

4.1 Two agents on a lane: the effect of agents' actions on FeAR

The first case study demonstrates how the actions chosen by the agents affect the FeAR values. The location of agents 1 and 2 are asymmetric in the sense that 1 is more constrained than 2. In Fig. 2 we consider three instances where (a) both agents move one step away from each other, (b) both agents move one step to the right, and (c) both agents move a step closer to each other. In (a) and (b) there are no collisions while in (c) both agents collide.

The MdR for the calculation is taken to be $\mu = \text{S0-S0}$. When agent 1 chooses the MdR (S0), agent 2 has $n(s, [A_i \leftarrow \mu_i], j) = 6$ feasible moves (L1, S0, R1, R2, R3 and R4). When 2 chooses the MdR (S0), 1 has $n(s, [A_i \leftarrow \mu_i], j) = 4$ feasible moves (L2, L1, S0 and R1).

In Instance (a) 1 chooses L1 which leaves 2 with $n(s, A, j) = 7$ feasible moves (L2, L1, S0, R1, R2, R3 and R4). So $\text{FeAR}_{1,2} = \frac{6-7}{6} = -0.17$ ³. And when 2 chooses R1, it increases the feasible actions to $n(s, A, j) = 5$ (L2, L1, S0, R1 and R2). Thus, $\text{FeAR}_{2,1} = \frac{4-5}{4} = -0.25$. So both agents have negative FeAR and are courteous to each other. Since 1 has more feasible moves ($n(s, A, j) = 5$) than when the other agent chooses the MdR ($n(s, [A_i \leftarrow \mu_i], j) = 4$), $\text{FeAR}_{1,1} = Z(5/4) = 1.0$. Similarly for 2, $\text{FeAR}_{2,2} = Z(7/6) = 1.0$. Thus, in this case, since both agents have more feasible moves than afforded by other agents choosing MdR, they have the maximum possible FeAR values.

In instance (b) 1 chooses R1 which leaves 2 with $n(s, A, j) = 5$ feasible moves (S0, R1, R2, R3 and R4), and 2 chooses R1 which leaves 1 with $n(s, A, j) = 5$ feasible moves (L2, L1, S0, R1 and R2). Thus, $\text{FeAR}_{1,2} = \frac{6-5}{6} = \frac{1}{6} = 0.17$ is positive and 1 is being assertive to 2, while, $\text{FeAR}_{2,1} = \frac{4-5}{4} = -0.25$ is negative and 2 is being courteous to 1. This causes a reduction in $\text{FeAR}_{2,2} = 5/6 = 0.83$ while $\text{FeAR}_{1,1} = Z(5/4) = 1$ remains at the maximum possible value.

² The simulations were run on CPU (Intel i7) in a Dell Latitude 7320 laptop with 16 GB of RAM. The source code for reproducing all case studies can be found at <https://github.com/DAI-Lab-HERALD/FeAR>.

³ $Z()$ only makes an impact when the magnitude of the ratios exceed 1. The ϵ in the denominator is only relevant in cases where the denominator is zero. For brevity, these are excluded in the explanations when unnecessary.



Figure 2: Case Study 1 explores the effect of an agent's action on the FeAR values for Mdr $\mu = S0-S0$. The actions selected by the agents for each instance are depicted in the first column. The second and third columns show the feasible actions available to the agents when under the other agents' Mdr $\mu_i = n(s, [A_i \leftarrow \mu_i], j)$ — or their chosen action $a_i = n(s, A, j)$. The move S0 is depicted in the center with a grey rectangle around it with other left and right moves on either side. The feasible moves are shown in green. The final column states the FeAR values that are calculated based on the number of feasible actions. $FeAR_{i,j}$ for $i \neq j$, represents the Feasible Action-Space Reduction of j by the action of j . $FeAR_{i,i}$ represent Feasible Action-Space Remaining for agent i left by the actions of all other agents. Since the definition of $FeAR_{i,i}$ is slightly different, they are shown in the plots with an extra white rectangle around them. For $i \neq j$, if $n(s, A, j) > n(s, [A_i \leftarrow \mu_i], j)$, then $FeAR_{i,j} < 0$ and vice versa.

In instance (c) 1 chooses R1 which leaves 2 with $n(s, A, j) = 5$ feasible moves (S0, R1, R2, R3 and R4), and 2 chooses L1 which leaves 1 with $n(s, A, j) = 3$ feasible moves (L2, L1 and S0). Thus, $FeAR_{1,2} = \frac{6-5}{6} = 0.17$ and $FeAR_{2,1} = \frac{4-3}{4} = 0.25$ are both positive which means that both agents are being assertive to each other. This leads to reductions in $FeAR_{1,1} = 5/6 = 0.83$ and $FeAR_{2,2} = 3/4 = 0.75$.

The above examples demonstrate how the actions of agents, determine the FeAR values. In the given locations and $\mu = S0-S0$, the actions L1 of 1 and R1 of 2 are considered courteous from the calculated FeAR values, while the actions R1 of 1 and L1 of 2 are considered assertive from the FeAR values. $FeAR_{1,2}$ is dependant on the move of 1 and $FeAR_{2,1}$ is dependant on the move of 2. On the other hand, $FeAR_{1,1}$ is dependant on the move of 2 and $FeAR_{2,2}$ is dependant on the move of 1. So, $FeAR_{i,j}$ where $i \neq j$ capture the assertiveness/courteousness of i towards j , while $FeAR_{i,i}$ captures the complement of the collective assertiveness/courteousness of all other agents on i . Thus, the chosen moves in combination with context determine the FeAR values and how agents are considered courteous or assertive.

The examples also showcase how the causal responsibility ascribed to agents can be asymmetric. In (a), both agents are courteous. In (c) both agents are assertive. In (b) one agent is assertive while the other is courteous. It is also worth noting that even in the case where both agents are found to be courteous (a), there is a small asymmetry in the FeAR values ($FeAR_{1,2} = -0.17$ and $FeAR_{2,1} = -0.25$). This stems from the asymmetry in the context — 1 is more restricted than 2 and has fewer feasible moves when considering the Mdrs. Thus, the causal influence of agents can be asymmetric and are highly dependent on the context.

4.2 Two agents on a lane: the effect of agents' Move de Rigueur on FeAR

In the first case study, we considered only one Mdr, $\mu = S0-S0$. In this case study, we explore how changing the Mdrs can affect the FeAR values, again for the same instances considered before. $\mu = R1-R1$ considers the norm where both agents are supposed to move one step to the right — like on a lane where cars are not supposed to stop in the middle of the lane. $\mu = R2-R2$ considers a norm where the agents are supposed to keep moving — now at a higher speed than when $\mu = R1-R1$. For ease of reference, we are repeating the FeAR values for $\mu = S0-S0$ here too (See Fig. 3). For instances (a), (b) and (c) the feasible actions available to agents 1 and 2 for each case remain the same as they were for the previous case study.

For $\mu = S0-S0$, 1 has $n(s, [A_i \leftarrow \mu_i], j) = 4$ feasible moves (L2, L1, S0 and R1) when 2 chooses the Mdr, and 2 has $n(s, [A_i \leftarrow \mu_i], j) = 6$ feasible moves (L1, S0, R1, R2, R3 and R4) when 1 chooses the Mdr. For $\mu = R1-R1$, 1 has $n(s, [A_i \leftarrow \mu_i], j) = 5$ feasible moves (L2, L1, S0, R1 and R2) when 2 chooses the Mdr (R1), and 2 has $n(s, [A_i \leftarrow \mu_i], j) = 5$ feasible moves (S0, R1, R2, R3 and R4) when 1 chooses the Mdr (R1). For $\mu = R2-R2$, 1 has $n(s, [A_i \leftarrow \mu_i], j) = 6$ feasible moves (L2, L1, S0, R1, R2 and R3) when 2 chooses the Mdr (R2), and 2 has $n(s, [A_i \leftarrow \mu_i], j) = 4$ feasible moves (R1, R2, R3 and R4) when 1 chooses the Mdr (R2).

The move L1 of 1 (instance (a)) offers more feasible actions to 2 than all the Mdrs considered (S0, R1 and R2). Thus, when 1 chooses L1, $FeAR_{1,2}$ for all the Mdrs are negative and indicate that 1 is being courteous. It can also be noted that the move L1 of 1 is considered more courteous under $\mu = R2-R2$ than under $\mu = S0-S0$ (and $\mu = R1-R1$ being an intermediate condition).




Figure 3: Case Study 2 explores the effect of MDR on the FeAR metric. We consider the same instances as in Case Study 1 (Fig. 2). Here three different MDRs are considered — $\mu = \text{S0-S0}$, $\mu = \text{R1-R1}$ and $\mu = \text{R2-R2}$. $\text{FeAR}_{i,j}$ for $i \neq j$, represents the Feasible Action-Space Reduction of j by the action of i . $\text{FeAR}_{i,i}$ represent Feasible Action-Space Remaining for agent i left by the actions of all other agents. For $i \neq j$, if $n(s, A, j) > n(s, [A_i \leftarrow \mu_i], j)$, then $\text{FeAR}_{i,j} < 0$ and vice versa. Thus changes in $n(s, [A_i \leftarrow \mu_i], j)$ affect the FeAR values.

Similarly, the move L1 of 2 (instance c) offers fewer feasible actions to 1 than all the MDRs considered (S0, R1 and R2). Thus when 2 chooses L1, $\text{FeAR}_{2,1}$ for all the MDRs are positive and indicate that 2 is being assertive. Also, as we go from $\mu = \text{S0-S0}$ to $\mu = \text{R2-R2}$, the 2 choosing L1 is considered increasingly assertive ($\text{FeAR}_{2,1}$ for $\mu = \text{S0-S0} < \text{FeAR}_{2,1}$ for $\mu = \text{R1-R1} < \text{FeAR}_{2,1}$ for $\mu = \text{R2-R2}$).

As we go from $\mu = \text{S0-S0}$ to $\mu = \text{R2-R2}$, the move R1 of 1 (instances (b) and (c)) goes from being considered assertive ($\text{FeAR}_{1,2} > 0$) to being considered courteous ($\text{FeAR}_{1,2} < 0$), while, the move R1 of 2 (instances (a) and (b)) goes from being considered courteous ($\text{FeAR}_{2,1} < 0$) to being considered assertive ($\text{FeAR}_{2,1} > 0$). These examples illustrate how the complex interplay of actions, contexts and norms (represented by the MDRs) can be represented by the FeAR metric.

4.3 Three agents at an Intersection: the effect of other agents' actions on FeAR

This case study illustrates the use of the FeAR metric in interactions of multiple agents. We consider interactions of three agents at an intersection (Fig. 4). Agents 1 and 2 are approaching the intersection from the left and agent 3 is approaching it from above. In instance (a), there are no collisions and agent 3 safely passes the intersection. If agent 1 decides to move faster (R4), it causes agent 1 to rear-end agent 2 (instance (b)). This rear-ending can be prevented if agent 2 moves forward (R1) as seen in instance (c). But, if agent 2 moves forward faster (R2), it causes a T-bone collision with agent 3 (instance (d)). This T-bone collision could have been prevented if agent 3 had chosen to stay (S0) as in instance (e). Thus, we explore a series of interactions where agent 2 is caught between agents 1 and 3.

As we move from instances (a) to (b), the action of agent 1 becomes more restrictive to agent 2 and $\text{FeAR}_{1,2}$ increases from 0.3 to 0.7. This value of $\text{FeAR}_{1,2}$ remains unchanged for instances (b), (c) and (d) — irrespective of whether agent 2 collides with agent 1 or 3. In instance (e), when agent 3 chooses (S0), it increases the feasible action space of 2. This leads to $\text{FeAR}_{1,2}$ decreasing to 0.4 even though the action of agent 1 does not change from instances

(b), (c) and (d). Thus, $\text{FeAR}_{1,2}$ is dependent on the actions of both agents 1 and 3. Similarly, we can see the influence of agent 1's action on $\text{FeAR}_{3,2}$ as we go through instances (e), (d), (c), (b) and (a) ($\text{FeAR}_{3,2} = 0.0$ for (e), $\text{FeAR}_{3,2} = 0.6$ for (d), (c) and (b), and $\text{FeAR}_{3,2} = 0.4$ for (a)).

Now, let us consider the effect of the actions of agent 2 on agents 1 and 3. In instances (a) and (b), agent 2 is choosing S0 which is the MDR. Hence, $\text{FeAR}_{2,1} = 0$ and $\text{FeAR}_{2,3} = 0$. When agent 2 chooses R1 in instance (c), it increases the feasible action space of 1 and thus $\text{FeAR}_{2,1}$ decreases to -0.2 . Since choosing R1 does not put agent 2 in the path of agent 3, the feasible action space of agent 3 remains unchanged and $\text{FeAR}_{2,3} = 0$. But, when agent 2 chooses R2 (instances (d) and (e)), it moves into the path of agent 3 restricting the feasible action space of agent 3. Thus, $\text{FeAR}_{2,3}$ increases to 0.2. Thus in instances (d) and (e), agent 2 is acting assertively towards agent 3 ($\text{FeAR}_{2,3} > 0$) and courteously towards agent 1 ($\text{FeAR}_{2,1}$).

In our analysis, we only considered one MDR $\mu = \text{S0-S0-S0}$ which gives equal priority at all. In real traffic, certain lanes are given priority over others at intersections. These changes in priorities could be captured using MDRs like $\mu = \text{S0-S0-D2}$ (which prioritises the vertical lane) or $\mu = \text{R2-R2-S0}$ (which prioritises the horizontal lane). Analyses of these scenarios can be found in the supplementary materials.

Overall, this case study illustrates how the complexities of multi-agent interactions are captured in the FeAR metric. We also show how agents caught in the middle can be assertive to some agents while being courteous to other agents.

4.4 Four agents on a lane: illustrating complex interactions

To highlight the insights provided by the FeAR metric in more complex scenarios, we analyzed the interactions of 4 agents on a single lane. To do this, we randomly generated 10,000 instances that all involved 4 agents on a lane 10 units long. We randomized the locations of the agents as well as their action choices (S0, R1, R2, R3 and R4). This set of actions was selected to mimic how agents would behave

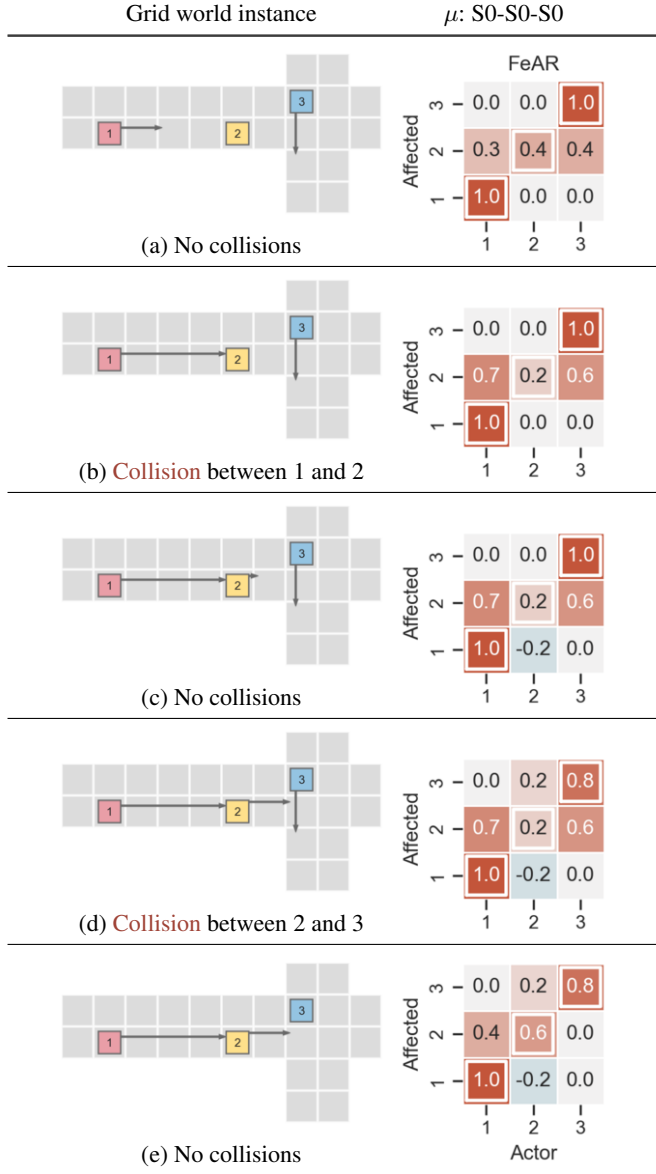


Figure 4: Case Study 3 explores the FeAR metric in the context of multi-agent interactions. We consider the interactions between three agents at an intersection. Agent 2 is caught between Agents 1 and 3. Instance (c) is the central condition where no collisions happen. But, if agent 2 slightly changes its move, it can collide with agent 1 (instance (b)) or with agent 3 (instance (d)). The collisions in instances (b) and (d) could also be prevented by the actions of agent 1 (instance (a)) or agent 3 (instance (d)) respectively. Thus, in these interactions $FeAR_{1,2}$ and $FeAR_{3,2}$ are dependent on the actions chosen by third parties too.

on a one-way lane. For each agent, the location on the lane and the action were chosen at random with equal likelihood.

We expected these randomly generated instances to vary substantially in complexity of the interaction. Here we investigated two extreme cases: an instance in which the agents acted most independently from each other, and an instance with a high level of interconnectedness between the agents. To do so, among the 10,000 generated instances we selected the ones with the minimum and maximum

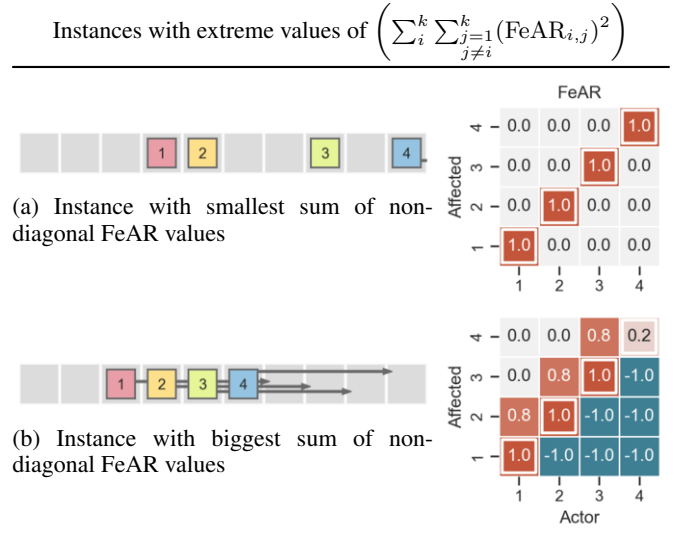


Figure 5: Case Study 4 illustrates the application of the metric in scenarios with potentially complex interactions of multiple agents. These two instances were selected from 10,000 instances based on the FeAR values. In Fig. 5a — the case with minimum sum of squares of FeAR — all agents choose Mdr (or the action equivalent to Mdr due to hitting a border). In Fig. 5b — the case with maximum sum of square of FeAR — agents are acting assertively to the agents ahead of them, while being courteous to the agents behind them. All of them deviate from the Mdr and together they avoid collisions.

sum of squares of the $FeAR_{i,j}$ for $i \neq j$ (Fig. 5a and Fig. 5b, respectively).

In the minimum sum-of-squares case Fig. 5a, agents 1,2 and 3 choose S0; agent 4's move to the right hits the boundary. Thus all agents would effectively be doing S0. Which means that effectively all agents are following the Mdr and hence FeAR values are zero. It must be noted that this is not the only case where the sum of squares of FeAR is zero (the others are not shown here).

In the maximum sum-of-squares case Fig. 5b, there are no collisions, but all the agents are deviating from the Mdr (S0). If agents 2, 3 or 4 choose the Mdr, they would cause a collision for all the agents behind them. Thus their current action increases the feasible action space of all agents behind them — indicated by the negative FeAR values below the diagonal. At the same time, since agents 1,2 and 3 choose R4, they severely restrict the feasible action of the agent in front of them, compared to when they choose the Mdr.

These two cases demonstrate that in multi-agent interactions, the sum of non-diagonal FeAR values captures the intuitions about the independence/interconnectedness of the agents' actions. In a similar vein, large positive (non-diagonal) FeAR values across multiple agents would flag cases where agents are being very assertive. Conversely, instances where agents are being very courteous would be marked by consistently negative values of FeAR (see the supplementary materials for details).

Overall, this case study illustrates how the FeAR metric can be analyzed in more complex scenarios in which the number of possible combinations of agent positions and actions is very large (more than 3 million in this case). To this end, comparisons between different outcomes of the interaction (e.g. different actions of the agents for a given set of positions) can be facilitated by aggregate FeAR values across all the agents.

5 Discussion and Conclusion

Our Feasible Action-space Reduction metric can be used as an indicator for causal responsibility in spatial interactions where each agent is physically present in the scene and, therefore, has to take at least one action. The FeAR metric captures the effects of actions of agents, norms (represented as *Moves de Rigueur* (MdR)), and the physical layout of the grid world. The case studies also illustrate how the responsibility ascription among agents can be asymmetric ($\text{FeAR}_{i,j} \neq \text{FeAR}_{j,i}$). We also illustrate how agents can be assertive to some agents ($\text{FeAR}_{i,j} > 0$), while being courteous to others ($\text{FeAR}_{i,j} < 0$). A positive FeAR value does not necessarily mean that a collision happens, but if a collision does happen, it indicates the share of responsibility of a given agent. For example, with three agents interacting at an intersection Fig. 4, when two agents (1 and 3) collectively constrain an agent (2), irrespective of whether it collides with 1 or 3, both agents share some responsibility for the state of 2. In real-world traffic interactions, when a car swerves to avoid hitting another car but ends up hitting a third car, it can be unclear who is responsible for the interaction. The FeAR metric along with a suitably chosen set of MdRs would be able to help guide the debate on the distribution of responsibility in such controversial cases.

We used *Moves de Rigueur* to represent norms. MdRs are actions that are expected of or prescribed to agents in an interaction. Ideally, these actions should be aligned with human behaviour and cognitive capacity. For example, consider the case where humans are used/trained to stop at an intersection. In this case, prescribing an MdR to keep driving through an intersection would be inappropriate. Similarly, automated vehicles require fallback actions and minimal risk maneuvers to ensure safety in case of emergencies. Well-defined MdRs aligned with human behaviour could provide good fallback actions for automated-vehicle algorithms.

We proposed Feasible Action-Space Reduction as a metric for causal responsibility among agents in a spatial interaction. While other models of responsibility [1, 6, 9, 12, 32, 35] ascribe responsibility for events, in spatial interactions, the FeAR metric assigns causal responsibility for the state and safety of agents. Causal responsibility has traditionally been evaluated based on probabilistic reasoning over sequences of events. Our formulation does not consider the probabilities of actions — we consider that all actions of an agent are equally likely. Our metric only considers an instant of the interaction and does not consider the history of events like other approaches. Because of the different nature of these metrics, no straightforward comparison can be made between FeAR and other models of causal responsibility. Developing methods to meaningfully compare disparate ways of quantifying responsibility is an important avenue for future work.

Existing literature on spatial navigation has several theories and models that try to address various aspects of how pedestrians and drivers engage in spatial interactions [18, 19, 21, 25, 28, 31]. Such models, typically incorporate several assumptions to explain and predict human behaviour. In contrast, the goal of this paper is not to describe human behaviour but to develop a model-agnostic framework for quantifying causal responsibility. However, we do believe that such models would be crucial for anticipating the behaviour of other agents when using FeAR for estimating responsibility of future actions (ex-ante responsibility).

Methods like Responsibility Sensitive Safety [29] prescribe "proper responses" for all agents and ascribe greater responsibility to agents that deviate from the "proper response". The MdR approach offers more freedom to the agents — agents do not have to follow

the MdR strictly and additionally have the possibility of being more courteous than the MdR. Furthermore, designing consistent MdRs can help humans and automation to have mutually compatible representations, which can help reduce the responsibility gaps in human-AI interactions [5, 27].

Limitations and future work. FeAR does not look at past interactions, states or beliefs of the agents. To get a more complete picture of causal responsibility as well as capture moral/role/capacity responsibility [34], more factors should be considered. The current implementation of FeAR only works for discrete and finite action and state spaces. In real-life interactions, an agent might have an uncountable number of actions to choose from. Our goal was to provide the first steps and the foundation towards measuring causal responsibility in spatial settings, and we felt that adding more complexity would have made the problem intractable. By making these assumptions, we were able to derive a metric that will then serve as a basis for investigating more complex settings, including continuous time and action space. Also, the current formulation of the metric scales exponentially with the number of agents and thus is not practical for real-time applications. A possible improvement point for future works would be to consider the calculated FeAR values as ground truth to train a supervised learning method to approximate the FeAR in a given instance.

Defining MdRs for all possible scenarios is a non-trivial task; especially if done manually. Establishing good norms for driving (like, 'always maintain safe time headway') would limit occurrence of scenarios where defining MdRs is challenging. Nevertheless, there might be situations where defining MdRs ahead of time might not be feasible, due to unknown circumstances or unpredictable behaviour of agents. Methods like Norm Synthesis [22, 23] or Responsibility-Sensitive Safety [29, 30] could be useful to automate the process of defining conflict-free MdRs. Making the MdR method applicable in practice is a challenging issue that future work should address.

We opted to present the metric in an intuitive way and use case studies to evaluate such intuitions and how it could be applied for understanding causal responsibility in multi-agent spatial interactions. In future works, especially when dealing with continuous domains, we believe that a formal evaluation might be necessary to achieve necessary insights on the robustness of the metrics and its impact on safety.

Conclusions. We proposed Feasible Action-space Reduction (FeAR) as a metric for quantifying the causal responsibility a given agent's actions have on the state and safety of other agents in an interaction. To represent norms, we introduced *Moves de Rigueur* (MdR) — consistent sets of default actions that are prescribed to each agent in an interaction. We used case studies to illustrate how the metric captures the effects of agents' actions, norms and multi-agent interactions. We also demonstrated an application of the metric to flag instances with high levels of interconnectedness/independence among the actions chosen by agents.

Ensuring safety and implementing meaningful human control in automated driving requires a common interdisciplinary framework that engages conversations between road users, vehicle manufacturers, traffic designers, regulatory agencies and policy makers [4, 24]. The FeAR metric and the concept of MdR is a step towards this interdisciplinary framework.

Acknowledgements

This work is supported by the TU Delft AI Labs programme.

References

- [1] Björn Bartling and Urs Fischbacher, 'Shifting the blame: On delegation and responsibility', *The Review of Economic Studies*, **79**(1), 67–87, (2012).
- [2] Niek Beckers, Luciano Cavalcante Siebert, Merijn Bruijnes, Catholijn Jonker, and David Abbink, 'Drivers of partially automated vehicles are blamed for crashes that they cannot reasonably avoid', *Scientific Reports*, **12**(1), 16193, (September 2022).
- [3] Simeon C. Calvert, Daniël D. Heikoop, Giulio Mecacci, and Bart van Arem, 'A human centric framework for the analysis of automated driving systems based on meaningful human control', *Theoretical Issues in Ergonomics Science*, **21**(4), 478–506, (July 2020).
- [4] Simeon C. Calvert, Stig O. Johnsen, and Ashwin George, 'Designing automated vehicle and traffic systems towards meaningful human control', in *Research Handbook on Meaningful Human Control of Artificial Intelligence Systems*, Edward Elgar Publishing, United Kingdom, (2023 (Accepted)).
- [5] Luciano Cavalcante Siebert, Maria Luce Lupetti, Evgeni Aizenberg, Niek Beckers, Arkady Zgonnikov, Herman Veluwenkamp, David Abbink, Elisa Giaccardi, Geert-Jan Houben, Catholijn M. Jonker, Jeroen van den Hoven, Deborah Forster, and Reginald L. Lagendijk, 'Meaningful human control: Actionable properties for AI system development', *AI and Ethics*, (May 2022).
- [6] Hana Chockler and Joseph Y. Halpern, 'Responsibility and Blame: A Structural-Model Approach', *Journal of Artificial Intelligence Research*, **22**, 93–115, (October 2004).
- [7] Ryan K. Cosner, Yuxiao Chen, Karen Leung, and Marco Pavone, 'Learning responsibility allocations for safe human-robot interaction with applications to autonomous driving', in *Proc. IEEE Conf. on Robotics and Automation*, (2023).
- [8] Mehdi Dastani and Vahid Yazdanpanah, 'Responsibility of AI Systems', *AI & SOCIETY*, **38**(2), 843–852, (April 2023).
- [9] Nir Douer and Joachim Meyer, 'The Responsibility Quantification Model of Human Interaction With Automation', *IEEE Transactions on Automation Science and Engineering*, **17**(2), 1044–1060, (April 2020).
- [10] Nir Douer and Joachim Meyer, 'Theoretical, Measured, and Subjective Responsibility in Aided Decision Making', *ACM Transactions on Interactive Intelligent Systems*, **11**(1), 1–37, (March 2021).
- [11] Nir Douer and Joachim Meyer, 'Judging One's Own or Another Person's Responsibility in Interactions With Automation', *Human Factors: The Journal of the Human Factors and Ergonomics Society*, **64**(2), 359–371, (March 2022).
- [12] Florian Engl, 'A Theory of Causal Responsibility Attribution', *SSRN Electronic Journal*, (2018).
- [13] European Commission. Directorate General for Communications Networks, Content and Technology. and High Level Expert Group on Artificial Intelligence., *Ethics Guidelines for Trustworthy AI*, Publications Office, LU, 2019.
- [14] Frank Flemisch, Matthias Heesen, Tobias Hesse, Johann Kelsch, Anna Schieben, and Johannes Beller, 'Towards a dynamic balance between humans and automation: Authority, ability, responsibility and control in shared and cooperative control situations', *Cognition, Technology & Work*, **14**(1), 3–18, (March 2012).
- [15] Maximilian Geisslinger, Franziska Poszler, and Markus Lienkamp, 'An ethical trajectory planning algorithm for autonomous vehicles', *Nature Machine Intelligence*, **5**(2), 137–144, (February 2023).
- [16] Diego Gracia, 'The many faces of autonomy', *Theoretical Medicine and Bioethics*, **33**(1), 57–64, (February 2012).
- [17] Joseph Y. Halpern, 'A modification of the halpern-pearl definition of causality', in *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, pp. 3022–3033, Buenos Aires, Argentina, (2015). AAAI Press.
- [18] Dirk Helbing and Péter Molnár, 'Social force model for pedestrian dynamics', *Physical Review E*, **51**(5), 4282–4286, (May 1995).
- [19] Sarvesh Kolekar, Joost De Winter, and David Abbink, 'Human-like driving behaviour emerges from a risk-based driver model', *Nature Communications*, **11**(1), 4850, (September 2020).
- [20] Janina Loh, 'Responsibility and robot ethics: A critical overview', *Philosophies*, **4**(4), 58–0, (2019).
- [21] Gustav Markkula, Yi-Shin Lin, Aravinda Ramakrishnan Srinivasan, Jac Billington, Matteo Leonetti, Amir Hossein Kalantari, Yue Yang, Yee Mun Lee, Ruth Madigan, and Natasha Merat, 'Explaining human interactions on the road by large-scale integration of computational psychological theory', *PNAS Nexus*, **2**(6), pgs163, (June 2023).
- [22] Javier Morales, Maite Lopez-Sanchez, Juan A. Rodríguez-Aguilar, Michael Wooldridge, and Wamberto Vasconcelos, 'Automated synthesis of normative systems', in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13*, pp. 483–490, Richland, SC, (2013). International Foundation for Autonomous Agents and Multiagent Systems.
- [23] Javier Morales, Maite López-Sánchez, Juan A. Rodríguez-Aguilar, Michael Wooldridge, and Wamberto Weber Vasconcelos, 'Synthesising liberal normative systems', in *Adaptive Agents and Multi-Agent Systems*, (2015).
- [24] Eleonora Papadimitriou, Haneen Farah, Geerten van de Kaa, Filippo Santoni de Sio, Marjan Hagenzieker, and Pieter van Gelder, 'Towards common ethical and safe 'behaviour' standards for automated vehicles', *Accident Analysis & Prevention*, **174**, 106724, (September 2022).
- [25] Jami Pekkanen, Oscar Terence Giles, Yee Mun Lee, Ruth Madigan, Tatsuru Daimon, Natasha Merat, and Gustav Markkula, 'Variable-Drift Diffusion Models of Pedestrian Road-Crossing Decisions', *Computational Brain & Behavior*, **5**(1), 60–80, (March 2022).
- [26] Filippo Santoni de Sio, 'The European Commission report on ethics of connected and automated vehicles and the future of ethics of transportation', *Ethics and Information Technology*, **23**(4), 713–726, (December 2021).
- [27] Filippo Santoni de Sio and Jeroen van den Hoven, 'Meaningful Human Control over Autonomous Systems: A Philosophical Account', *Frontiers in Robotics and AI*, **5**, 15, (February 2018).
- [28] Wilko Schwarting, Alyssa Pierson, Javier Alonso-Mora, Sertac Karaman, and Daniela Rus, 'Social behavior for autonomous vehicles', *Proceedings of the National Academy of Sciences*, **116**(50), 24972–24978, (December 2019).
- [29] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. On a Formal Model of Safe and Scalable Self-driving Cars, October 2018.
- [30] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Vision Zero: On a Provable Method for Eliminating Roadway Accidents without Compromising Traffic Throughput, January 2019.
- [31] Olger Siebinga, Arkady Zgonnikov, and David A. Abbink, 'Modelling communication-enabled traffic interactions', *Royal Society Open Science*, **10**(5), 230537, (May 2023).
- [32] Stelios Triantafyllou, Adish Singla, and Goran Radanovic. Actual Causality and Responsibility Attribution in Decentralized Partially Observable Markov Decision Processes, August 2022.
- [33] Ibo R. van de Poel and Lambèr M.M. Royakkers, *Ethics, Technology, and Engineering : An Introduction*, Wiley-Blackwell, United States, 2011.
- [34] Nicole A. Vincent, 'A Structured taxonomy of responsibility concepts', in *Moral Responsibility*, eds., Nicole A. Vincent, Ibo van de Poel, and Jeroen van den Hoven, Library of Ethics and Applied Philosophy, 15–35, Springer, Springer Nature, United States, (2011).
- [35] Vahid Yazdanpanah and Mehdi Dastani, 'Distant Group Responsibility in Multi-agent Systems', in *PRIMA 2016: Principles and Practice of Multi-Agent Systems*, eds., Matteo Baldoni, Amit K. Chopra, Tran Cao Son, Katsutoshi Hirayama, and Paolo Torroni, volume 9862, 261–278, Springer International Publishing, Cham, (2016).
- [36] Vahid Yazdanpanah, Enrico H. Gerding, Sebastian Stein, Corina Cirstea, M. C. Schraefel, Timothy J. Norman, and Nicholas R. Jennings, 'Different Forms of Responsibility in Multiagent Systems: Sociotechnical Characteristics and Requirements', *IEEE Internet Computing*, **25**(6), 15–22, (November 2021).
- [37] Vahid Yazdanpanah, Sebastian Stein, Enrico H. Gerding, and Nicholas R. Jennings, 'Applying strategic reasoning for accountability ascription in multiagent teams', in *Proceedings of the Workshop on Artificial Intelligence Safety 2021 Co-Located with the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI 2021)*, Virtual, August, 2021, eds., Huáscar Espinoza, John A. McDermid, Xiaowei Huang, Mauricio Castillo-Effen, Xin Cynthia Chen, José Hernández-Orallo, Seán Ó hÉigeartaigh, Richard Mallah, and Gabriel Pedroza, volume 2916 of *CEUR Workshop Proceedings*. CEUR-WS.org, (2021).