

Quantifying cooperation between artificial agents using synergistic information

Patricia Wollstadt, Matti Krüger

2023

Preprint:

This is an accepted article published in 2022 IEEE Symposium Series on Computational Intelligence (SSCI). The final authenticated version is available online at: <https://doi.org/10.1109/SSCI51031.2022.10022283> Copyright 2023 IEEE

Quantifying cooperation between artificial agents using synergistic information

1st Patricia Wollstadt

Honda Research Institute Europe GmbH

Offenbach/Main, Germany

patricia.wollstadt@honda-ri.de, 0000-0002-7105-5207

2nd Matti Krüger

Honda Research Institute Europe GmbH

Offenbach/Main, Germany

matti.krueger@honda-ri.de, 0000-0002-2730-1116

Abstract—When designing interactive human-machine systems, it is often assumed that it is desirable for such systems to behave cooperatively towards a human operator in order to improve trust, acceptance, and usability, but also to increase task efficiency. To design cooperative human-machine interaction (HMI) systems, we have to be able to define and quantitatively describe cooperative behavior, for example, to control, optimize, or evaluate the interaction. Despite the increased interest in cooperative HMI in recent years, an approach that provides a suitable definition of cooperation and also a method for its quantification is still missing. In the present work, we therefore develop a novel definition of cooperative behavior in HMI contexts, based on which we propose to quantify cooperation using recent methods from information theory. We define cooperation as joint, coordinated actions that are mutually adapted such as to facilitate the realization of a goal. Thus, cooperation is characterized by a synergistic effect of joint actions towards the goal. Here, we propose to quantify cooperation using the recently introduced partial information decomposition framework from information theory, which proposes measures to quantify the synergistic contributions of two inputs to a target variable. In particular, we propose to apply the synergy measure to two or more input variables describing agents' actions towards a target variable that describes the current goal state. As a first validation, we applied our approach in a grid-world environment, in which two agents solve a cooperative foraging task. We found that synergy was higher for agents implementing cooperative strategies compared to baseline and non-cooperative strategies, and we found higher synergy in trials with high numbers of cooperative actions. We conclude that the synergy is a promising candidate measure for identifying cooperative behavior in goal-oriented interactions.

Index Terms—cooperation, human-machine interaction, partial information decomposition, synergy, information theory.

I. INTRODUCTION

Task demands that surpass the abilities of any individual may in some cases be fulfilled when multiple individuals join forces to solve the task in a cooperative fashion. This concerns interactions between humans, but, increasingly, also interactions between humans and machines. Human-machine interaction (HMI) is not necessarily tailored towards cooperation and instead frequently targets an automation of specific tasks. Designing for automation has been criticized for negative effects on human operators, such as loss of expertise due to a lack of engagement, as well as reduced trust and adaptability [1]. In contrast, systems that behave

cooperatively towards a human operator have been argued to lead to, e.g., higher trust, acceptance, and usability [2]–[5] while circumventing negative automation effects [6], [7]. Thus, more cooperative human-machine interactions are desirable from the perspectives of both the ability to solve a task, but also user satisfaction.

To be able to design cooperative HMI systems, it is necessary to define and quantitatively describe cooperative interactions, for example, to control, optimize, or evaluate system behavior [8], [9]. Even though the interest in cooperative HMI has significantly increased in recent years [5], [10], [11] and other disciplines such as psychology [12], biology [13], [14], or game theory [15]–[18] have a long-standing tradition in researching cooperative behavior and its emergence, the concept of cooperation in HMI systems and its quantification stay elusive [2], [9], [19].

In the present work, we therefore develop a novel definition of cooperative behavior in HMI contexts and present an approach to quantify cooperative behavior based on this definition, using recent methods from information theory, namely partial information decomposition (PID) [20]. In particular, we propose to quantify cooperative behavior as the synergistic information contribution of two agents' actions to a variable describing the current state of the target. As a first demonstration of our approach, we analyze interactions between artificial agents that solve a cooperative task in an environment presented in [21]. In this environment, we show that cooperative behavior is accompanied by more synergistic information between agents' actions and the target.

II. METHODS

A. Cooperation definition

As a first step, we develop a definition of cooperative behavior applicable to the HMI context, derived from existing literature.

1) *Prerequisites*: Our definition makes two assumptions: first, we assume that each of two or more agents—by design or negotiation—agrees or *commits* to working together with the other agents in a coordinated fashion towards a goal [2], [5], [10], [22], [23]. Second, we assume that each agent's goal is reachable via subtasks that are *interdependent* and thus require a coordination of the agents' actions [2], [5], [10], [23]. The first assumption eliminates the need for any form

of negotiation between agents prior to solving the task. The second assumption ensures that the environment and task allow for cooperative behavior as defined in the next section.

Note that agents do not have to pursue the same goal to enable cooperative behavior. Rather, agents can also cooperate on reaching individual goals [22].

2) *Cooperation as mutual facilitation*: We here use a definition of cooperation that follows prior work in [10], [11], [22]. We define cooperation as joint, coordinated activity towards solving interdependent subtasks, where the coordination leads to a *facilitation* of individual agents' actions with respect to the current goal. Central to this definition is the facilitation of actions towards reaching a (sub-)goal in comparison to an agent's individual actions. Some (sub-) goals may even only become realizable through such joint, cooperative activity. Further, cooperation is not symmetrical under our definition. Rather one agent may facilitate another agent's actions without receiving support in return.

This definition of cooperative behavior requires certain abilities in agents to realize cooperative behavior. These abilities are described further in the following.

3) *Effectors for manipulation of the environment and information sharing*: Interacting agents have to be equipped with suitable effectors for acting in the environment and sharing of information, either im- or explicitly. Effectors have to enable suitable motor capabilities that allow to reach the common goal.

Effectors may further enable the explicit communication of relevant information about current states and intentions that can be perceived and understood by other agents. Modes of communication may, amongst others, include visual display or vocal commands (e.g., [24]). Sharing of relevant information is required to establish mutual transparency and predictability, which in turn enables a planning of actions that establishes mutual support and facilitation with other agents [10], [24]. However, information sharing does not have to be explicit and may happen implicitly, for example, through actions or interaction with the environment [24]. The latter phenomenon is termed *stigmergy* [25] and describes an indirect communication between agents through the environment, where typically no agent is specifically targeted as the receiver.

4) *Sensors and interpretation of their input*: Next, agents have to possess sensory capabilities to receive relevant information about the current state of the world and other agent's states. Relevant information is assumed to be either shared explicitly or implicitly, as described in the previous paragraph [25]–[28].

5) *Internal representation of current and predicted states*: To plan actions that adapt to other agent's actions with the goal of executing joint, cooperative activities, an agent must be able to form internal representations of current and future states of the environment and other agents' actions within the environment [5], [6]. Internal representations are formed from perceived information about the current environment state including other agents' actions. Here, information sharing may be explicit by sharing a current task or goal, or more implicit

by observing an other agent's current action [29]. Predictions of future actions and their outcomes may be obtained from explicit communication of other agents' intentions or plans, or predictions may be generated in a data-driven fashion by building models based on earlier observed behavior (e.g., [17]). Meaningful internal representations allow for an anticipatory planning of actions that adapts to and coordinates with other agents' actions [29].

6) *Action planning and control with the ability for mutual adaption*: Agents must be able to plan actions such as to adapt to other agents' (future) states [17]. Such a planning is based on an agent's internal representations of the current and predicted future environment states, including other agents' predicted actions. The ability to plan assumes a sufficient degree of autonomy and motor capability to adapt actions to the current or predicted environment state. The planning of actions should in particular be able to handle interdependencies in current subtasks between its own and other agents' (predicted) actions. As per our definition, cooperative behavior thereby means a planning that strives for a facilitation of own and other agent's actions towards the goal [10], [22].

B. Synergy as a measure of cooperative actions

At the core of the presented definition of cooperation is the facilitation of agents' actions towards a goal by other agents [5], [22], [23]. In particular, central to long-standing cooperation definitions in philosophy, social sciences, and psychology, is the notion that cooperation means a successful management of interferences that leads to a "facilitation" of actions [22], or to "producing something that is a genuine joint project" [10]. We propose that defining cooperative behavior in such a way, corresponds to a *synergistic* effect of the joint actions towards the goal. Here, a synergistic effect means that the joint action achieves a contribution towards the goal that can not be achieved from the actions of one agent alone. We propose to measure such an effect using the recently introduced information-theoretic framework of partial information decomposition (PID) [20], [30]–[32].

Information-theoretic methods have been used previously to investigate and describe collective behavior of artificial [33]–[35] and non-artificial agents [36]–[38], but have also been used to investigate the interaction of agents with their environment [26], [39]–[43]. Information-theoretic methods have desirable properties for the investigation of agent interactions. Methods are model-free and thus make only minimal assumptions on the data [44]–[46]. As a result, methods are applicable to data generated in a wide range of scenarios and require little prior knowledge about the data and their distribution. Successful applications of information-theoretic methods can be found in neuroscience [47], [48], finance [49], artificial life [39], [41], [43], or engineering [50].

We here use the novel framework of PID, which describes how two or more input variables, A_0 and A_1 , contribute to the outcome of a target variable, T . Here A_0 , A_1 , and T denote random variables with alphabets, $a_0 \in \mathcal{A}_{A_0}$, $a_1 \in \mathcal{A}_{A_1}$, and $t \in \mathcal{A}_T$ and probability distributions $p(A_0 = a_0)$,

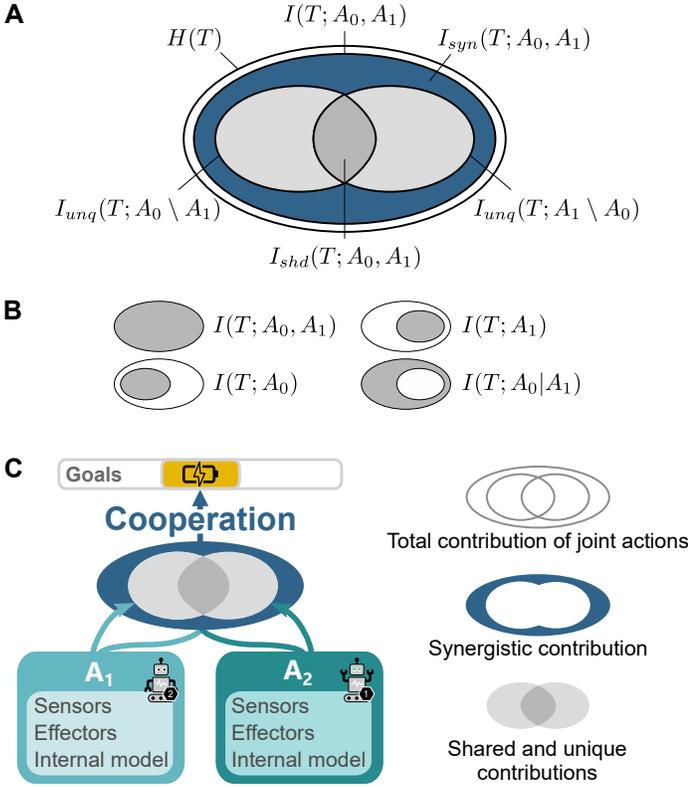


Fig. 1. A) Partial information diagram illustrating the relationship between the target variable's entropy, $H(T)$, and the joint mutual information, $I(T; A_0, A_1)$, with partial information quantities. B) Mutual information terms and their relationship to the joint mutual information. C) Schematic illustration of agent interactions and relationship to partial information atoms.

$p(A_1 = a_1)$, $p(T = t)$. PID provides a decomposition of the total information the two input variables, A_0 and A_1 , provide about a target, T , as measured by the joint mutual information, $I(T; A_0, A_1)$. PID decomposes the mutual information into three types of non-negative contributions (Fig. 1 A): information can either be *uniquely* provided by either one of the inputs, $I_{unq}(T; A_0 \setminus A_1)$ and $I_{unq}(T; A_1 \setminus A_0)$, it can be redundantly *shared* by both inputs, $I_{shd}(T; A_0, A_1)$, or it can be provided *synergistically* by both inputs, $I_{syn}(T; A_0, A_1)$. The synergistic contribution thereby describes a contribution that is exclusively provided when considering both inputs together and that can not be obtained by either input alone.

The original PID framework as proposed by Williams and Beer [20] decomposes the joint mutual information, $I(T; A_0, A_1)$, into the four contributions or “atoms” described above,

$$I(T; A_0, A_1) = I_{unq}(T; A_0 \setminus A_1) + I_{unq}(T; A_1 \setminus A_0) + I_{shd}(T; A_0, A_1) + I_{syn}(T; A_0, A_1). \quad (1)$$

Additionally, the mutual information between the individual input variables and the target can be decomposed into the unique and shared information by

$$\begin{aligned} I(T; A_0) &= I_{unq}(T; A_0 \setminus A_1) + I_{shd}(T; A_0, A_1), \\ I(T; A_1) &= I_{unq}(T; A_1 \setminus A_0) + I_{shd}(T; A_0, A_1). \end{aligned} \quad (2)$$

The proposed PID atoms can not be obtained from classical information-theoretic terms (see also Fig. 1 B). Williams and Beer therefore introduce a set of additional axioms to define the PID and propose an initial measure of the shared information, from which the unique and synergistic information can be derived. However, the proposed axioms do not uniquely determine a measure of each of the atoms and subsequent work has proposed additional assumptions and atoms, as well as alternative measures of the PID terms (see, e.g., [51] for a discussion). We here use a popular measure proposed by Bertschinger et al. [30], who derive a definition of the unique information from the argument that truly unique information should be exploitable in a decision problem. Based on this assumption, the authors argue that the unique information should depend only on the marginal probability distributions, $p(T, A_0)$ and $p(T, A_1)$, but not the full joint distribution, $p(T, A_0, A_1)$. Thus, the unique information should not change for different joint distributions, $q(T, A_0, A_1)$, from a space Δ_P with the same marginals,

$$\begin{aligned} \Delta_P = \{q \in \Delta : q(T = t, A_0 = a_0) &= p(T = t, A_0 = a_0) \\ &\text{and } q(T = t, A_1 = a_1) = p(T = t, A_1 = a_1)\}, \end{aligned} \quad (3)$$

where Δ is the space of all probability distributions over the support of T , A_0 , and A_1 . From their assumption, Bertschinger et al. derive the unique information as

$$I_{unq}(T; A_0 \setminus A_1) = \min_{q \in \Delta_P} I_q(T; A_0|A_1), \quad (4)$$

where $I_q(T; A_0|A_1)$ is a conditional mutual information computed with respect to q . This problem can be solved by means of convex optimization as for example implemented by [52]. Note that this definition of the PID holds for the case of two input variables only. For details refer to [20], [30], [32].

In the following, we will use the PID framework to investigate the behavior of two artificial agents. In particular, we estimate the PID according to [30], between input variables representing the observable behavior of two agents, with respect to a target variable describing the current state of or distance to the goal (Fig. 1 C). We investigate whether the synergy is able to quantify the degree of cooperative behavior in a set of experiments described next.

C. Experiments

As a first evaluation of our cooperation definition and the ability of PID measures to quantify the degree of cooperative behavior, we apply our approach to a model system in which two agents solve a cooperative task, namely the level-based foraging (LBF) environment [21]¹. The LBF environment is

¹Code available from <https://github.com/semtable/lb-foraging>.

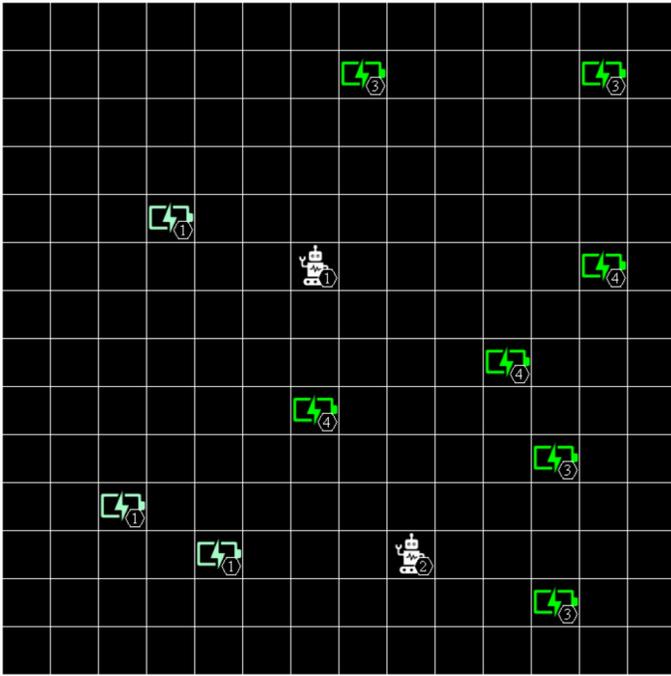


Fig. 2. Example of modified level-based foraging environment of size 14×14 . The environment was originally proposed in [21]. Battery numbers indicate total required agent levels. Agent numbers indicate individual agent levels.

a 2D grid-world, in which multiple mobile agents and static food items are placed. The agents’ goal is to collect as many food items as possible, available actions are steps to the north, south, east, or west, or collecting an adjacent food item. Two agents may not occupy the same tile. To collect a food item, an agent must be located on a tile that is next to a tile with a food item. Agents and items are assigned levels and an item can only be collected if an agent’s level or the sum of the agents’ levels simultaneously collecting the item is equal or larger than the item’s level (Fig. 2). Whenever an agent reaches a tile adjacent to a food item, the agent will attempt to collect the item in the next time step. To collect a food item jointly, both agents have to be located on a tile next to the food item.

We set up the environment to a size of 25 by 25 with two agents with a random assigned level, $l_A \in \{1, 2\}$. In the environment, 10 food items were placed at random. We adapted the LBF environment to allow food items’ levels to be drawn such that a pre-defined percentage of items could only be collected jointly by both agents. We set the number of food items that could only be collected jointly, c , to 0, 20, 40, 60, 80, and 100%, respectively. Whenever half of all placed items were collected, new items were placed on the grid, where the placement was restricted to an area distant from the agents’ current location. This constraint was introduced because respawning in close proximity to an agent could create artifacts that biased measures of cooperation because also non-cooperative behavior would increase another agent’s chance to maximize item intake. (Arguably an avoidance of spawning in close proximity can also be seen as more natural

when interpreted as a substitute for a larger environment.) We introduced the respawning of food items to allow for arbitrarily long running times of trials, such as to obtain a sufficient number of samples for the robust estimation of information-theoretic measures. For the experiments, the length of one run was set to 1000 time steps. We found that the chosen environment size and trial length generated rich dynamics in agent behavior for all heuristics.

We chose two agents such as to analyze one interaction only and not unnecessarily increase the complexity of the experiments presented. However, the LBF environment may be extended to more than two interacting agents. For settings with more than two agents, the synergy as a measure of cooperation may be estimated between pairs of agents. Alternatively, the PID can be estimated for more than two input variables and a target (e.g., [31]). Yet, a fully multivariate decomposition of inputs from $n > 2$ agents that also considers interactions between all possible subsets of inputs as well as interactions between interactions, produces a number of PID atoms that grows super-exponentially (see [32] for details). Hence, the latter approach may be considered infeasible both computationally but also in terms of explanatory value.

We investigated different heuristics for selecting an agent’s next step. As a baseline condition, *BL*, we selected the agent’s next step at random. We further defined an egoistic heuristic, *Ego*, where an agent targeted the closest food and tried to pick it up, irrespective of the item’s level. To prevent agents from getting “stuck” trying to repeatedly collect a food item with a too high level, agents were given rudimentary memory about the success of recent actions, allowing them to disregard a food item after a specified number of failures.

Next, we used a modified version of a more cooperative strategy, *Social*, proposed by [21] in the original version of the LBF environment (termed *H4* in the original publication). The strategy avoids failure in picking up an item by selecting a target not just based on its proximity to oneself, but also consider its proximity to other agents. With knowledge about other agent’s locations, an agent can select targets that are closest to the center of all agents. Additionally, the strategy picks targets such that the food items level is equal or smaller than the summed level of all agents. This heuristic increases the chance for multiple agents to approach the same target and reach it at a similar time, while ensuring that the agents’ summed level is sufficient to pick up the food item. We extend this strategy by the rudimentary memory described above to avoid situations in which two agents blocked each other from collecting the food, which we observed for the original implementation.

However, with the *Social* heuristic, an agent still operates under the assumption that another agent follows a similar strategy. The mutual knowledge of other agents’ intentions is not explicitly established in this strategy. Hence, to remove the need for guessing other agents’ intentions, we added a further strategy, *Coop*, in which agents have access to the goals of other agents. Hence, the strategy models a simple form of information sharing between agents, to enable transparency

about objectives in accordance with our cooperation definition described above. This information enables a more adaptive goal selection such as through recognition of whether or not another agent requires any support in consuming its selected food item.

Last, we implemented a “competitive” strategy, *Comp*, in which agents behaved in an adversarial fashion. While goal transparency enables supportive behavior strategies, it can also be exploited to maximize an adversarial agent’s relative reward. Rather than avoiding a food item that another agent targets and which this other agent may consume on its own, the *Comp* heuristic “steals” this item purposefully whenever the own distance to the item is smaller than the other agent’s distance to it. The other agent’s current target is made transparent to the adversarial heuristic thus simulating a communication of the current goal. We introduced an adversarial heuristic to investigate whether an adversarial interaction would also lead to measurable synergistic effects, and whether such effects were distinguishable from synergistic effects arising from cooperative interactions.

A summary of implemented heuristics is shown in Table I.

TABLE I
OVERVIEW OF INCLUDED AGENT HEURISTICS

Heuristic	Description
<i>BL</i>	Baseline: select next move at random
<i>Ego</i>	Egoistic agent: go to the closest visible food irrespective of its level and try to collect it
<i>Social</i>	Social agent: go to visible food that is closest to the center of all other agents and has a compatible level
<i>Coop</i>	Cooperative agent: share own goal with other agents to allow for mutual support
<i>Comp</i>	Competitive agent: collected food item targeted by another agent if own distance to target is smaller

D. Data collection and analysis

We performed the experiment for all heuristics, *BL*, *Ego*, *Social*, *Coop*, and *Comp*, and all ratios of food items requiring cooperative actions, c . For each combination of heuristic and values for c , we ran 20 repetitions of the experiment, where each trial was ran with a different random initialization of the environment. Each trial lasted 1000 steps such as to obtain a sufficient and equal number of samples for estimation of information-theoretic measures over experiments and trials. An equal number of samples is crucial, because most information-theoretic estimators suffer from bias problems that depend on sample size. By keeping the sample size constant, estimates become comparable, for example, via permutation testing (see below).

For each trial, we recorded the agents’ selected actions over time, as well as the collected food items. We then estimated PID measures using the agents’ actions, $A_0, A_1 \in \{0, 1, 2, 3, 4, 5\}$, as input and the food collected by the first agent as target, T . Agents’ actions encoded a step to the North, East, South, West, or collection of an item. The target was the current sum of food items collected by the first agent, discretized into five bins. Note that while the synergy is a

symmetric measure, we here investigate the synergistic contribution with respect to the food collected by *one* agent only. It is conceivable that in some scenarios, agents’ actions have a synergistic effect towards one agent’s target but not the other agent’s target. In other words, while the synergy is a symmetric measure, we can investigate asymmetric interactions through suitable variable choice (e.g., calculating the synergy with respect to individual goals only).

We estimated the PID using the definition of the measures proposed in [30]. An approach to estimate the PID measure was proposed by [52] and implemented in the IDTxl python toolbox [53], [54], which is used here. The joint probability distributions involved in calculating the measures are obtained from the empirical data collected in the experiments. We tested for differences in estimated measures between heuristics using a permutation test with 500 permutations and using a critical alpha level of $\alpha_{crit} = 0.05$.

III. RESULTS

We first verified that our environment settings led to cooperative actions. We defined a cooperative action as the joint collection of a food item by both agents. We found that with an increasing ratio of cooperative food items, c , the number of cooperative actions increased as well (Fig. 3). We conclude that the settings led to a scenario in which agents jointly collected items. Note that in the non-cooperative condition, $c = 0.0$, the number of cooperative actions was lower also for those heuristics that explicitly target the joint collection of items. This is because any agent attempts to collect a food item whenever possible. In the $c = 0.0$ condition the collection of an item is always possible for individual agents, hence also cooperative agents can collect items individually whenever a single agent reaches a collectable item first. Conversely, also the egoistic and competitive agents sometimes performed cooperative actions whenever two agents arrived at a food item at the same time step.

When estimating the synergy between agents’ actions and the target, we found that the synergy was close to 0 in the setting requiring no cooperation, $c = 0.0$, (Fig. 4) and increased for higher degrees of required cooperation for both cooperative heuristics, *Social* and *Coop*. For all other heuristics, *BL*, *Ego*, and *Comp*, the synergy was low for all levels of c . For all levels of c , both cooperative strategies lead to significantly higher synergy between agents’ actions and the target than the *BL* heuristic. Synergy was also higher than for the *Ego* heuristic, except for $c = 0.0$.

When comparing both cooperative strategies, *Coop* and *Social*, we did not find a difference in estimated synergy, indicating that the difference in the two strategies, a mutual transparency of current targets, did not affect the behavior significantly. Here, more complex environments and tasks may show an effect on the resulting degree of cooperation and the synergy between actions.

Furthermore, for our cooperation strategy, *Coop*, we found higher synergy values per trial as a function of the total number of cooperative actions (Fig. 5). Here, a cooperative action

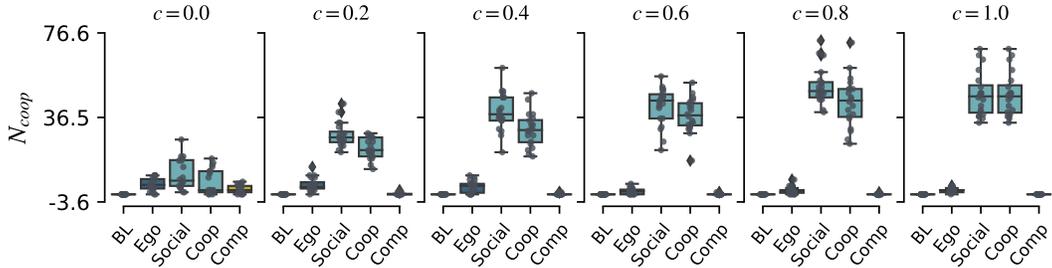


Fig. 3. Total number of cooperative actions, N_{coop} , for different heuristics and cooperation levels, c . Box plots indicate quartiles with whiskers indicating 1.5 times of the inter-quartile range. The horizontal bar indicates the median.

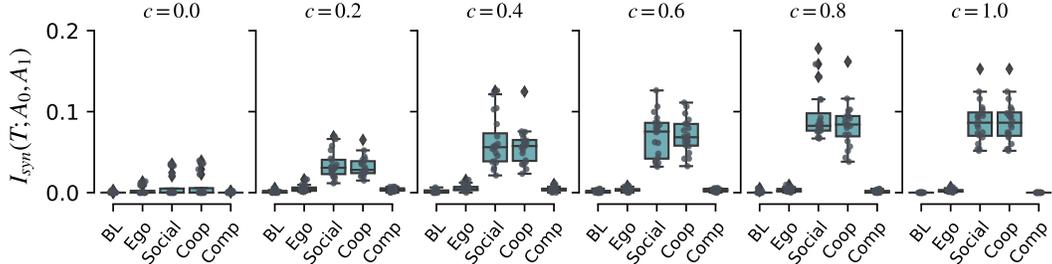


Fig. 4. Estimated synergy, $I_{syn}(T; A_0, A_1)$, for different heuristics and cooperation levels, c . Box plots indicate quartiles with whiskers indicating 1.5 times of the inter-quartile range. The horizontal bar indicates the median.

was defined as a food item that was collected jointly by both agents.

Complementary to the synergy, the unique information contribution of the first agent’s action was highest for the non-cooperative environment settings, $c = 0.0$ to $c = 0.4$ (Fig. 6). Note that here all heuristics, except for the baseline, BL , showed comparable degrees of unique information, indicating that all heuristics succeeded in the individual collection of food items. The low values of the baseline BL can be explained by the low total number of collected items, leading to a lower entropy of the target variable and hence to a lesser amount of information to be explained or contributed to.

In sum, cooperative actions towards collecting food items were accompanied by higher values of synergistic information contribution between the agents’ actions and the target variable. For conditions that required less cooperation, the unique information contribution by the first agent was higher for all heuristics. We conclude that cooperative behavior is reflected by a high synergistic contribution of agents’ actions towards the target.

IV. CONCLUSION

We demonstrated a first application of our novel framework for quantitatively describing cooperative actions from observable behavior using recent methods from information theory. The experiments provide first evidence that information theory and in particular the partial information decomposition (PID) framework are promising tools to describe interactions between agents, in particular cooperation towards a goal.

We use the recently introduced PID framework that provides methods for disentangling contributions of multiple input variables towards an output variable. We here apply the synergy measure to successfully describe cooperative actions as a synergistic contribution of the joint actions of two agents towards a target. The synergy, in particular, captures the joint contribution that is exclusively obtained when considering both inputs together and that can not be obtained from one input alone. Hence, we believe that the synergy is a promising candidate measure to capture the essence of cooperative behavior—a facilitation of actions with respect to reaching a targeted goal.

We provide first evidence that synergy may be used to differentiate cooperative behavior from other types of interactions with respect to a task. However, additional experiments in different task settings are required to further investigate if synergy can reliably differentiate, for example, between cooperative and adversarial behavior. It is conceivable that adversarial behavior in certain scenarios leads to a decrease in the target variable, while the adversarial agent adapts its behavior to that of a cooperative agent, which may lead to a high synergy, while one of the agents does not behave cooperatively. Also, point-wise or local PID measures may be applied in future experiments to provide a time-resolved analysis of interactions between agents [31], [55].

We here present an approach that adopts a definition of cooperation that builds on prior work in philosophy, social sciences, and psychology, which defines cooperation as the management of interferences between individual actions with the goal of (mutual) support and facilitation of individual

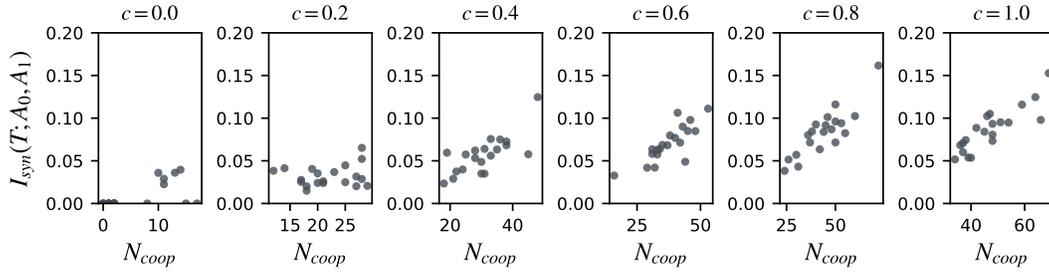


Fig. 5. Scatter plots of total number of cooperative actions versus estimated synergy, $I_{syn}(T; A_0, A_1)$, per trial for cooperation heuristic, *Coop*.

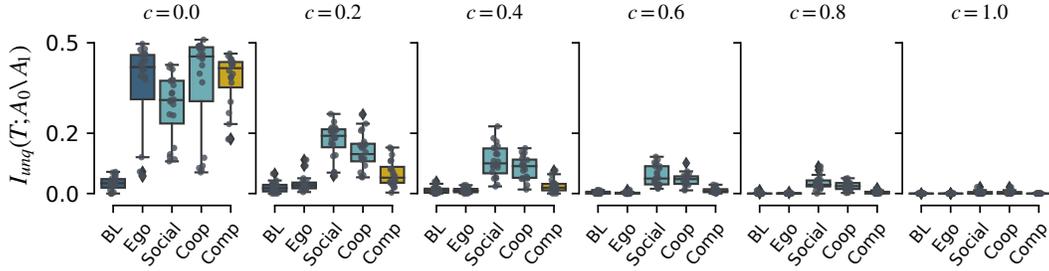


Fig. 6. Estimated unique information contribution of agent A_0 's actions, $I_{unq}(T; A_0 \setminus A_1)$, for different heuristics and cooperation levels, c . Box plots indicate quartiles with whiskers indicating 1.5 times of the inter-quartile range. The horizontal bar indicates the median.

actions [5], [10], [22], [23]. We want to highlight that we here chose a definition that is also applicable to interactions involving at least one human partner such as to evaluate the proposed approach in HMI in future work. Furthermore, the synergy measure is proposed to quantify cooperation under this definition—however, alternative definitions of cooperative behavior exist, under which different quantities may be more suitable as a general measure of cooperation or its value. Examples are, cooperation as adaption in human-robot interaction (e.g., [56]), intention estimates that enable shared operations (e.g., [57]), and game-theoretic definitions of cooperative games as those which satisfy some or all requirements of non-negativity, monotonicity, superadditivity, and convexity [16]. Importantly, superadditivity refers to the synergy generated by cooperative behavior, e.g., quantified as surplus payoff by the Harsanyi dividend [15]. The applicability of the approach proposed here to game-theoretic scenarios and (cooperative) two-player games may be evaluated in the future.

The approach presented here is not tailored to a specific scenario, and does not require a-priori knowledge of possible strategies or behaviors of involved agents. Further, the applied information-theoretic methods and estimators make only mild assumptions about the data collected from observable interactions. We thus believe that our approach and the proposed measure is applicable in a wide range of scenarios and recorded behaviors. We here present a first application to a model system, while evaluation in further model systems, but also more complex scenarios, ideally involving human agents, is subject to future work.

REFERENCES

- [1] J.-M. Hoc, "From human-machine interaction to human-machine cooperation," *Ergonomics*, vol. 43, no. 7, pp. 833–843, 2000.
- [2] K. Bengler, M. Zimmermann, D. Bortot, M. Kienle, and D. Damböck, "Interaction Principles for Cooperative Human-Machine Systems," *Information Technology*, vol. 54, no. 4, pp. 157–164, 2012.
- [3] M. Hassenzahl and H. Klapperich, "Convenient, clean, and efficient? The experiential costs of everyday automation," in *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*, 2014, pp. 21–30.
- [4] B. Sendhoff and H. Wersing, "Cooperative intelligence—a humane perspective," in *2020 IEEE International Conference on Human-Machine Systems (ICHMS)*, 2020, pp. 1–6.
- [5] J. Bütepage and D. Kragic, "Human-robot collaboration: from psychology to social robotics," *arXiv preprint arXiv:1705.10146 [cs.RO]*, 2017.
- [6] M. Krüger, C. B. Wiebel, and H. Wersing, "From tools towards cooperative assistants," *Proceedings of the 5th International Conference on Human Agent Interaction (HAI '17)*, pp. 287–294, 2017.
- [7] M. Krüger, "An enactive approach to perceptual augmentation in mobility," Ph.D. dissertation, Ludwig-Maximilians-Universität München, 2021.
- [8] E. Meisner, S. Šabanović, V. Isler, L. R. Caporael, and J. Trinkle, "Shadowplay: A generative model for nonverbal human-robot interaction," in *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2009, pp. 117–124.
- [9] N. Jarrasse, V. Sanguineti, and E. Burdet, "Slaves no longer: review on role assignment for human-robot joint motor action," *Adaptive Behavior*, vol. 22, no. 1, pp. 70–82, 2013.
- [10] G. Klein, P. J. Feltovich, J. M. Bradshaw, and D. D. Woods, "Common ground and coordination in joint activity," *Organizational Simulation*, vol. 53, pp. 139–184, 2005.
- [11] R. Gervasi, L. Mastrogiacomo, and F. Franceschini, "A conceptual framework to evaluate human-robot collaboration," *The International Journal of Advanced Manufacturing Technology*, vol. 108, pp. 841–865, 2020.
- [12] A. P. Melis and D. Semmann, "How is human cooperation different?" *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 365, no. 1553, pp. 2663–2674, 2010.

- [13] J. L. Sachs, U. G. Mueller, T. P. Wilcox, and J. J. Bull, "The evolution of cooperation," *The Quarterly Review of Biology*, vol. 79, no. 2, pp. 135–160, 2004.
- [14] M. A. Nowak, "Evolving cooperation," *Journal of Theoretical Biology*, vol. 299, pp. 1–8, 2012.
- [15] J. C. Harsanyi, "A simplified bargaining model for the n-person cooperative game," in *Papers in Game Theory*. Springer, 1982, pp. 44–70.
- [16] G. Owen, *Game theory*. Emerald Group Publishing, 2013.
- [17] S. Nikolaidis, J. Forlizzi, D. Hsu, J. Shah, and S. Srinivasa, "Mathematical models of adaptation in human-robot collaboration," *arXiv preprint arXiv:1707.02586 [cs.RO]*, 2017.
- [18] X. Liang and Z. Yan, "A survey on game theoretical methods in human-machine networks," *Future Generation Computer Systems*, vol. 92, pp. 674–693, 2019.
- [19] A. Sciutti, A. Bisio, F. Nori, G. Metta, L. Fadiga, T. Pozzo, and G. Sandini, "Measuring human-robot interaction through motor resonance," *International Journal of Social Robotics*, vol. 4, no. 3, pp. 223–234, 2012.
- [20] P. L. Williams and R. D. Beer, "Nonnegative decomposition of multivariate information," *arXiv Preprint arXiv:1004.2515 [cs.IT]*, 2010.
- [21] S. V. Albrecht and S. Ramamoorthy, "A game-theoretic model and best-response learning method for ad hoc coordination in multiagent systems," *12th International Conference on Autonomous Agents and Multiagent Systems 2013, AAMAS 2013*, vol. 2, pp. 1155–1156, 2013.
- [22] J.-M. Hoc, "Towards a cognitive approach to human-machine cooperation in dynamic situations," *International Journal of Human Computer Studies*, vol. 54, no. 4, pp. 509–540, 2001.
- [23] M. E. Bratman, "Shared cooperative activity," *The Philosophical Review*, vol. 101, no. 2, pp. 327–341, 1992.
- [24] A. Bauer, D. Wollherr, and M. Buss, "Human-robot collaboration: a survey," *International Journal of Humanoid Robotics*, vol. 5, no. 01, pp. 47–66, 2008.
- [25] A. S. Klyubin, D. Polani, and C. L. Nehaniv, "Tracking information flow through the environment: Simple cases of stigmerg," in *In: Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, J. Pollack, Ed. MIT Press, 2004.
- [26] M. Harder, D. Polani, and C. L. Nehaniv, "Two agents acting as one," in *Proceedings of Artificial Life XII*, 2010, pp. 599–606.
- [27] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, "Local measures of information storage in complex distributed computation," *Information Sciences*, vol. 208, pp. 39–54, 2012.
- [28] X. R. Wang, J. M. Miller, J. T. Lizier, M. Prokopenko, and L. F. Rossi, "Quantifying and Tracing Information Cascades in Swarms," *PLoS ONE*, vol. 7, no. 7, p. e40084, 2012.
- [29] N. Sebanz and H. Bekkering, "Joint action: bodies and minds moving together," *Trends in Cognitive Sciences*, vol. 10, no. 2, pp. 70–76, 2006.
- [30] N. Bertschinger, J. Rauh, E. Olbrich, J. Jost, and N. Ay, "Quantifying unique information," *Entropy*, vol. 16, no. 4, pp. 2161–2183, 2014.
- [31] A. Makkeh, A. J. Gutknecht, and M. Wibral, "Introducing a differentiable measure of pointwise shared information," *Physical Review E*, vol. 103, no. 3, p. 032149, 2021.
- [32] A. J. Gutknecht, M. Wibral, and A. Makkeh, "Bits and pieces: Understanding information decomposition from part-whole relationships and formal logic," *Proceedings of the Royal Society A*, vol. 477, no. 2251, p. 20210110, 2021.
- [33] X. R. Wang, J. M. Miller, J. T. Lizier, M. Prokopenko, and L. F. Rossi, "Measuring Information Dynamics in Swarms," in *Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems (ECAL 2011)*, T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, and M. Dorigo, Eds. Paris, France: MIT Press, 2011, pp. 838–845.
- [34] O. M. Cliff, J. T. Lizier, X. R. Wang, P. Wang, O. Obst, and M. Prokopenko, "Quantifying Long-Range Interactions and Coherent Structure in Multi-Agent Dynamics," *Artificial Life*, vol. 23, no. 1, pp. 34–57, 2017.
- [35] P. Cofta, D. Ledziński, S. Śmigiel, and M. Gackowska, "Cross-entropy as a metric for the robustness of drone swarms," *Entropy*, vol. 22, no. 6, 2020.
- [36] Y. Sun, L. F. Rossi, C.-C. Shen, X. R. Wang, J. T. Lizier, and M. Prokopenko, "Information Transfer in Swarms with Leaders," *Collective Intelligence*, pp. 11–14, 2014.
- [37] N. Orange and N. Abaid, "A transfer entropy analysis of leader-follower interactions in flying bats," *The European Physical Journal Special Topics*, vol. 224, no. 17–18, pp. 3279–3293, 2015.
- [38] E. Crosato, L. Jiang, V. Lecheval, J. T. Lizier, X. R. Wang, P. Tichit, G. Theraulaz, and M. Prokopenko, "Informative and misinformative interactions in a school of fish," *Swarm Intelligence*, vol. 12, no. 4, pp. 283–305, 2018.
- [39] D. Tarapore, M. Lungarella, and G. Gómez, "Fingerprinting agent-environment interaction via information theory," *Proceedings of the 8th International Conference on Intelligent Autonomous Systems (IAS-8)*, pp. 512–520, 2004.
- [40] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns, "Methods for quantifying the informational structure of sensory and motor data," *Neuroinformatics*, vol. 3, no. 3, pp. 243–262, 2005.
- [41] N. Bertschinger, E. Olbrich, N. Ay, and J. Jost, "Autonomy: an information theoretic perspective," *Biosystems*, vol. 91, no. 2, pp. 331–345, 2008.
- [42] Y. Yamada, S. Nishikawa, K. Shida, R. Niiyama, and Y. Kuniyoshi, "Neural-body coupling for emergent locomotion: A musculoskeletal quadruped robot with spinobulbar model," *IEEE International Conference on Intelligent Robots and Systems*, pp. 1499–1506, 2011.
- [43] R. D. Beer and P. L. Williams, "Information processing and dynamics in minimally cognitive agents," *Cognitive Science*, vol. 39, no. 1, pp. 1–38, 2015.
- [44] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [45] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2005, vol. 100.
- [46] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, New Jersey: John Wiley & Sons, Inc., 2006.
- [47] M. Wibral, C. Finn, P. Wollstadt, J. T. Lizier, and V. Priesemann, "Quantifying information modification in developing neural networks via partial information decomposition," *Entropy*, vol. 19, no. 9, p. 494, 2017.
- [48] P. Wollstadt, K. K. Sellers, L. Rudelt, V. Priesemann, A. Hutt, F. Fröhlich, and M. Wibral, "Breakdown of local information processing may underlie isoflurane anesthesia effects," *PLoS Computational Biology*, vol. 13, no. 6, p. e1005511, 2017.
- [49] R. Quax, G. Chliamovitch, A. Dupuis, J.-L. Falcone, B. Chopard, A. G. Hoekstra, and P. M. A. Sloot, "Information processing features can detect behavioral regimes of dynamical systems," *Complexity*, vol. 6047846, 2018.
- [50] P. Wollstadt and S. Schmitt, "Interaction-aware sensitivity analysis for aerodynamic optimization results using information theory," in *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2021, pp. 1–8.
- [51] M. Wibral, J. T. Lizier, and V. Priesemann, "Bits from Biology for Computational Intelligence," *Frontiers in Robotics and AI*, vol. 2, pp. 1–25, 2014.
- [52] A. Makkeh, D. O. Theis, and R. Vicente, "BROJA-2PID: A robust estimator for bivariate partial information decomposition," *Entropy*, vol. 20, no. 4, p. 271, 2018.
- [53] P. Wollstadt, J. T. Lizier, R. Vicente, C. Finn, M. Martínez-Zarzuola, P. Mediano, L. Novelli, and M. Wibral, "IDTxI: The Information Dynamics Toolkit xl: a Python package for the efficient analysis of multivariate information dynamics in networks," *Journal of Open Source Software*, vol. 4, no. 34, p. 1081, 2019.
- [54] L. Novelli, P. Wollstadt, P. Mediano, M. Wibral, and J. T. Lizier, "Large-scale directed network inference with multivariate transfer entropy and hierarchical statistical testing," *Network Neuroscience*, vol. 3, no. 3, pp. 827–847, 07 2019.
- [55] C. Finn and J. T. Lizier, "Probability mass exclusions and the directed components of pointwise mutual information," *Entropy*, vol. 20, no. 11, p. 826, 2018.
- [56] M. Gienger, D. Ruiken, T. Bates, M. Regaieg, M. Meibner, J. Kober, P. Seiwald, and A. C. Hildebrandt, "Human-Robot Cooperative Object Manipulation with Contact Changes," *IEEE International Conference on Intelligent Robots and Systems*, pp. 1354–1360, 2018.
- [57] A. Belardinelli, A. R. Kondapally, D. Ruiken, D. Tanneberg, and T. Watabe, "Intention estimation from gaze and motion features for human-robot shared-control object manipulation," *ArXiv Preprint arXiv:2208.08688 [cs.RO]*, 2022.