

Measuring cooperation in Hanabi agents using information theory

Christiane Wiebel, Matti Krüger, Patricia Wollstadt

2023

Preprint:

This is an accepted article published in International Conference on Hybrid Human-Artificial Intelligence (HHAI 2023) . The final authenticated version is available online at: <https://doi.org/10.3233?FAIA230117>

Poster: Quantifying cooperation between rule-based Hanabi agents using information theory

Patricia WOLLSTADT ^{a,1}, Matti KRÜGER ^a, and Christiane B. WIEBEL-HERBOTH ^a

^a*Honda Research Institute Europe, Offenbach am Main, Germany*

Abstract. With increasing abilities of artificial systems and agents, it has been proposed to design those systems to behave cooperatively towards human users. However, it is often unclear how to evaluate whether a cooperation in human-machine interaction takes place or how to even quantitatively describe such a cooperation. We have previously proposed a framework to quantify the degree of cooperation in an interaction between two agents, using novel measures from information theory. We here extend the initial evaluation of this framework by applying the proposed measure to rule-based artificial agents playing Hanabi. We show that the measure correlates with the number of points scored by the agents and that the framework allows to describe uni-directional interactions between stronger and weaker players. The proposed framework may be used to evaluate and guide HMI design towards more cooperative interactions, which is believed to lead to a more pleasant user experience.

Keywords. human-machine interaction, cooperation, information theory, partial information decomposition, Hanabi

Introduction Today, many intelligent systems are developed towards operating in interaction with human users instead of acting autonomously. It has been proposed to enable these systems to act *cooperatively* [17,14,2,3], as cooperation is thought to be beneficial with respect to task success, trust, user satisfaction, and counters negative effects of pure automation [7,17,3,8,6]. However, cooperation in this setting is not dichotomous—rather, cooperation is a continuous phenomenon such that different degrees of cooperative behavior have to be distinguished [5]. Today, it is still difficult to benchmark and evaluate cooperative behavior in human-machine interaction (HMI) [2,16,15,12]. To successfully design cooperative HMI, it is therefore desirable to develop metrics that reliably identify different qualities of interactions and in particular, identify different degrees of cooperation. Ultimately, such a measure may also be used to guide agent behavior, as it has been shown that optimizing e.g., for task success, may not always lead to the most satisfying user experience [18]. We have previously proposed a measure for cooperation and have provided a first, successful evaluation of the measure in a toy system [21,22]. Here, we extend previous results by applying the proposed measure in the popular cooperation benchmark Hanabi and demonstrate its ability to differentiate between degrees of cooperative behavior in different agent strategies.

¹Corresponding Author: Patricia Wollstadt, Honda Research Institute Europe, Carl-Legien-Str. 30, 63073 Offenbach am Main, Germany; E-mail: patricia.wollstadt@honda-ri.de.




Figure 1. A) Illustration of PID framework. B-C) Synergy versus score for agent self-play and the Piers agent playing other strategies. D) Synergy towards individual task success for VDB versus Internal agent.

Methods In prior work [21,22], we have proposed an information-theoretic measure of cooperative behavior using the partial information decomposition (PID) framework [20]. PID allows to decompose the contribution of two or more input variables towards a target variable into *unique*, redundantly *shared* and *synergistic* contributions. Here the synergy describes a contribution that is exclusively provided by both inputs together and that can not be obtained from one input alone (Figure 1A). We proposed the synergy as the operationalization for a common notion of cooperation in HMI, namely the “facilitation” of agents’ actions towards a goal [13,9]. We successfully demonstrated our approach in a toy system. We here extend this evaluation in the popular benchmark for cooperative interaction, Hanabi [1]. We used an existing implementation of the Hanabi learning environment and five rule-based agents presented in [4] (originally proposed in [19]). To estimate the synergy between agents’ actions and the score in each round, we use a measure of the synergy proposed by Ince [10] which is implemented in the dit Python toolbox [11]. We estimated the synergy between the first agent’s action, A_0 , the subsequent action by the second action, A_1 , and whether this second action resulted in an error, no point, or a point. We distinguished between agents’ actions *Reveal* (either color or number), *Discard*, and *Play*. The investigated agents were the *Internal*, *Piers*, *IGGI*, *Outer*, and *VDB* (see [19] and [4] for a detailed description).

Experiments and Results We simulated 1000 episodes of Hanabi in different experimental settings. First, we evaluated the synergy as a marker of cooperative behavior in agent self-play, where we found a strong correlation between the synergy and the average score per episode (Fig. 1B). Second, we estimated the synergy for games in which the *Piers* agent (the most successful agent) played all other agents, and again found a strong correlation with the average score (Fig. 1C). Last we investigated whether the synergy was able to reflect also asymmetric interactions between a well-performing and a poorly performing agent. Here, we estimated the synergy between the *VDB* and *Internal* agent (Fig. 1D), where the former achieved higher scores in all games compared to the latter. We estimated the synergy once between *VDB*’s and *Internal*’s actions with respect to scores assigned after *VDB*’s action and once with respect to scores assigned after *Internal*’s actions. We found that the synergy towards *Internal*’s actions was higher than towards *VDB*’s actions, indicating a stronger one-sided or asymmetric cooperation towards the weaker *Internal* agent.

Conclusion We have recently introduced a novel framework for quantifying cooperative behavior using methods from information theory. Here we extend previous findings by providing evidence for its ability to differentiate between different cooperative strategies in a famous benchmark of cooperative behavior for artificial intelligence [1].

We show that the synergy measure correlates closely with points scored during a game, which is a direct outcome of cooperative game plays. Additionally, we show that the synergy allows insights into uni-directional interactions.

While it could be argued that in this particular case there are little additional insights from estimating the synergy, in other contexts, the relationship between cooperative behavior and task success may be less clear: i.e., cooperation may not necessarily lead to the highest task success or may not be necessary for solving a task. Nevertheless, we may be interested in evaluating whether cooperative behavior takes place for other reasons than optimal performance, e.g., user satisfaction [18]. Here, our proposed measure fills a gap as it allows to make a statement about the quality of the interaction, independent of overall task success, and thus may be an alternative guidance for designing HMI systems towards a more pleasant user-experience. An evaluation of our approach in human-machine interaction will be subject to future work.

References

- [1] Nolan Bard, Jakob N. Foerster, Sarath Chandar, Neil Burch, Marc Lanctot, H. Francis Song, Emilio Parisotto, Vincent Dumoulin, Subhodeep Moitra, Edward Hughes, et al. The Hanabi challenge: A new frontier for AI research. *Artificial Intelligence*, 280:103216, 2020.
- [2] Klaus Bengler, Markus Zimmermann, Dino Bortot, Martin Kienle, and Daniel Damböck. Interaction Principles for Cooperative Human-Machine Systems. *Information Technology*, 54(4):157–164, 2012.
- [3] Judith Büttepage and Danica Kragic. Human-robot collaboration: from psychology to social robotics. *arXiv preprint arXiv:1705.10146*, 2017.
- [4] Rodrigo Canaan, Xianbo Gao, Youjin Chung, Julian Togelius, Andy Nealen, and Stefan Menzel. Evaluating rl agents in hanabi with unseen partners. In *AAAI’20 Reinforcement Learning in Games Workshop*. (Code available from <https://github.com/rocanaan/hanabi-ad-hoc-learning>), 2020.
- [5] Allan Dafoe, Yoram Bachrach, Gillian Hadfield, Eric Horvitz, Kate Larson, and Thore Graepel. Cooperative AI: machines must learn to find common ground. *Nature*, 593(7857):33–36, 2021.
- [6] Joost C. F. de Winter, Sebastiaan M. Petermeijer, and David A. Abbink. Shared control versus traded control in driving: A debate around automation pitfalls. *Ergonomics*, 0(0):1–27, 2022.
- [7] Marc Hassenzahl and Holger 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*, pages 21–30, 2014.
- [8] Jean-Michel Hoc. From human-machine interaction to human-machine cooperation. *Ergonomics*, 43(7):833–843, 2000.
- [9] Jean Michel Hoc. Towards a cognitive approach to human-machine cooperation in dynamic situations. *International Journal of Human Computer Studies*, 54(4):509–540, 2001.
- [10] Robin A. A. Ince. Measuring multivariate redundant information with pointwise common change in surprisal. *Entropy*, 19(7):318, 2017.
- [11] R. G. James, C. J. Ellison, and J. P. Crutchfield. dit: a Python package for discrete information theory. *The Journal of Open Source Software*, 3(25):738, 2018.
- [12] Nathanael Jarrasse, Vittorio Sanguineti, and Etienne Burdet. Slaves no longer: review on role assignment for human-robot joint motor action. *Adaptive Behavior*, 22(1):70–82, 2013.
- [13] Gary Klein, Paul J. Feltovich, Jeffrey M. Bradshaw, and David D. Woods. Common ground and coordination in joint activity. *Organizational Simulation*, 53:139–184, 2005.
- [14] Matti Krüger, Christiane B. Wiebel, and Heiko Wersing. From tools towards cooperative assistants. *Proceedings of the 5th International Conference on Human Agent Interaction - HAI ’17*, pages 287–294, 2017.
- [15] Eric Meisner, Selma Šabanović, Volkan Isler, Linnda R. Caporael, and Jeff Trinkle. Shadowplay: A generative model for nonverbal human-robot interaction. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction, HRI’09*, pages 117–124, 2009.

- [16] Alessandra Sciutti, Ambra Bisio, Francesco Nori, Giorgio Metta, Luciano Fadiga, Thierry Pozzo, and Giulio Sandini. Measuring Human-Robot Interaction Through Motor Resonance. *International Journal of Social Robotics*, 4(3):223–234, 2012.
- [17] Bernhard Sendhoff and Heiko Wersing. Cooperative intelligence-a humane perspective. In *2020 IEEE International Conference on Human-Machine Systems (ICHMS)*, pages 1–6. IEEE, 2020.
- [18] Ho Chit Siu, Jaime Peña, Edenna Chen, Yutai Zhou, Victor Lopez, Kyle Palko, Kimberlee Chang, and Ross Allen. Evaluation of human-AI teams for learned and rule-based agents in hanabi. In *Advances in Neural Information Processing Systems*, volume 34, pages 16183–16195, 2021.
- [19] Joseph Walton-Rivers, Piers R. Williams, Richard Bartle, Diego Perez-Liebana, and Simon M. Lucas. Evaluating and modelling Hanabi-playing agents. In *2017 IEEE Congress on Evolutionary Computation (CEC)*, pages 1382–1389. IEEE, 2017.
- [20] Paul L. Williams and Randall D. Beer. Nonnegative decomposition of multivariate information. *arXiv Preprint arXiv:1004.2515 [cs.IT]*, pages 1–14, 2010.
- [21] Patricia Wollstadt and Matti Krüger. Quantifying cooperation between artificial agents using information theory. In *Hybrid Human Artificial Intelligence (HHAI) 2022*, page 7538, 2022.
- [22] Patricia Wollstadt and Matti Krüger. Quantifying cooperation between artificial agents using synergistic information. In *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1044–1051. IEEE, 2022.