



Exploring the Difference between Solving and Teaching in Sensorimotor Tasks

Manuel Bied, Mohamed Chetouani

► To cite this version:

Manuel Bied, Mohamed Chetouani. Exploring the Difference between Solving and Teaching in Sensorimotor Tasks. HRI '20: ACM/IEEE International Conference on Human-Robot Interaction, Mar 2020, Cambridge, United Kingdom. pp.139-141, 10.1145/3371382.3378284 . hal-02984868

HAL Id: hal-02984868

<https://hal.science/hal-02984868>

Submitted on 1 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Exploring the Difference between Solving and Teaching in Sensorimotor Tasks

Manuel Bied
Sorbonne Université
Institut des Systèmes Intelligents et de Robotique
CNRS UMR 7222
Paris, France
bied@isir.upmc.fr

Mohamed Chetouani
Sorbonne Université
Institut des Systèmes Intelligents et de Robotique
CNRS UMR 7222
Paris, France
chetouani@isir.upmc.fr

ABSTRACT

Considering a human is not only solving a task, but actively teaching how to solve it to a robot has not been extensively explored and is an important step to improve LfD algorithms. We explored the difference between solving and teaching in a sensorimotor task. In a first experiment participants first solved a continuous maze task and gave demonstrations for a robot afterwards. While teaching the participants could give negative demonstrations (how not to solve the task). In a second experiment we asked new participants to rate how informative they perceive the demonstrations from the first experiment. The results show that significantly more demonstrations from the Teaching-phase are perceived as informative than from the Solving-phase. Furthermore, significantly more negative than positive demonstrations were perceived as informative.

KEYWORDS

LfD; Sensorimotor Communication; HRI; Pedagogy

ACM Reference Format:

Manuel Bied and Mohamed Chetouani. 2020. Exploring the Difference between Solving and Teaching in Sensorimotor Tasks. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20 Companion)*, March 23–26, 2020, Cambridge, United Kingdom. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3371382.3378284>

I INTRODUCTION

Learning from Demonstration (LfD) is a popular approach to transfer new skills to an agent (e.g. robot) in an intuitive manner [1, 2]. LfD research usually assumes an expert giving correct demonstrations, or novice user giving (possibly) flawed demonstrations. The common assumption is that the user just solves the task, giving her best possible solution. However, we could also imagine that the user includes additional information in a demonstration, than simply solving the task alone. Following the *pedagogy hypothesis* stating that “*humans are adapted to transfer knowledge to, and receive knowledge from, conspecifics through teaching*” [5] this assumption seems reasonable. We will call a context like this a pedagogical situation; “*settings in which one agent is choosing information to*




Figure 1: Environment of the continuous maze task. Humans have to solve the continuous maze task (Solving) or teach how to solve it (Teaching) with the possibility to provide negative demonstrations.

transmit to another agent for the purpose of teaching a concept” [14]. The interesting question here is whether people change their demonstrations in these pedagogical situations, and if so, how do they change such demonstrations? The work of [8] investigates the difference between participants in a Doing- and Showing-condition, solving a 2-dimensional grid world task. The Show participants tended to choose paths that disambiguate their goal as compared to Do participants. This difference can be used to speed up the learning of how to solve the task for the agent [9]. In the context of physical interaction, a closely related concept is sensorimotor communication (SMC) [12, 13]. SMC uses the same channel to execute an action and convey additional information resulting in actions that combine pragmatic and communicative goals. Calinon [4] identifies the exploitation of the social interaction as future direction of LfD research, and gives learning from counterexamples as example. Further, the work of [11] identifies the learning from different instruction types as a challenge for LfD research. Similarly to counterexamples, we propose to allow teachers to provide negative demonstrations as one approach to address this challenge. Negative demonstrations are demonstrations that explicitly demonstrate what not to do. Related research (e.g. [3, 6, 10]) integrates possibilities to learn from sub-optimal and flawed demonstrations, but does not offer the possibility to purposefully demonstrate what not to do. In our study we are interested in how humans use sensorimotor communication to teach a sensorimotor task. We are interested in the difference between solving and teaching, furthermore we explore how humans use the possibility of using negative demonstrations to teach the task. We had two hypotheses:

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HRI '20 Companion, March 23–26, 2020, Cambridge, United Kingdom

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7057-8/20/03.

<https://doi.org/10.1145/3371382.3378284>

- (1) Humans modify their behavior when teaching how to solve a sensorimotor task in comparison to solving it.
- (2) Humans perceive the use of negative demonstrations as informative.

II EXPERIMENTS

We designed and conducted an IRB approved study to investigate the difference between humans solving and teaching a task using the sensorimotor channel.

The study consisted of two experiments:

- (1) We asked participants first to solve a sensorimotor task, and afterwards to teach how to solve it to a robot.
- (2) We asked new participants to rate the demonstrations from the first experiment.

The first experiment was conducted with 42 participants. We asked the participants to teach how to solve a continuous maze task to a robot. In order to add the sensorimotor dimension the task was solved on tablet with a digital pen. The task consists of going from a start zone to a goal zone. The environment has impassable terrain in black, and unsafe terrain in yellow. The impassable terrain can not be crossed. The unsafe terrain could be crossed, but is as the name states; unsafe. An instance of the maze task can be seen in Figure 1. The experiment includes a *Solving*- and a *Teaching*-phase, in order to show that there is a difference between solving and teaching (similar to [8]). In the Solving-phase, the participants were just asked to solve the task correctly. In the Teaching-phase, the participants were asked to give positive and negative demonstrations. Positive demonstrations are correct solutions to the task. Negative demonstrations go through the unsafe zone, but were also required to start at the goal zone and end in the end zone. The Solving- and Teaching-phase was repeated for 15 different instances of the task. The participants did not directly interact with a robot, but were introduced to it with a picture and a description. Additionally, they were told that the robot would learn from their demonstrations later. In the second experiment we asked 72 new participants to rate the collected demonstrations, discarding the demonstrations that did not fulfill the requirements of a correct positive, respectively negative demonstration. The participants were told how the maze task and demonstration process from the first experiment worked. Afterwards, they were asked to rate the statement "*The demonstrator deviates from the simplest way of solving to convey to you other information about the task*"¹ on a 5-point-Likert scale from *strongly disagree* to *strongly agree*. Each demonstration was rated by at least 6 participants.

III RESULTS

Ratings of 1 and 2 were counted as *Non-informative*, ratings of 3 as neutral and ratings of 4 and 5 as *Informative*. Each demonstration was classified according to a majority voting between all participants that gave a rating on a particular demonstration, demonstrations that did not have a majority for neither *Non-informative* nor *Informative* were counted as *Not-clear*. The absolute numbers of the classified demonstrations are reported in Table 1. In the Solving-phase there were no negative demonstrations possible (N/A), the

¹The study was conducted in French, the quoted phrase corresponds to the equivalent English translation.

Table 1: Absolute values of the results of the majority votes how informative an example is.

		Non-informative	Not-clear	Informative
positive demonstrations	Solving	499	48	44
	Teaching	851	132	206
negative demonstrations	Solving	N/A	N/A	N/A
	Teaching	198	174	513

Table 2: Relative numbers for Solving (in %).

	Non-informative	Not-clear	Informative
positive demonstrations	84.43	8.12	7.45

Table 3: Relative numbers for Teaching (in %).

	Non-informative	Not-clear	Informative
all demonstrations	50.58	14.75	34.67
positive demonstrations	71.57	11.1	17.33
negative demonstrations	22.37	19.66	57.97

relative numbers for the Solving-phase can be seen in Table 2 and for the Teaching-phase in Table 3. In the set of positive demonstrations, the relative portion of informative demonstrations in the Teaching-phase (17.33%) is higher than in the Solving-phase (7.43%). This difference between the Solving- and Teaching-phase is significant considering the columns for Non-informative, Not clear and Informative, $\chi^2(2, N = 1780) = 39.52$, $p < .01$, as well when only considering the columns for Non-informative and Informative, $\chi^2(1, N = 1600) = 34.42$, $p < .01$. In the Teaching-phase, the relative portion of informative demonstrations is significantly higher for the negative demonstrations (57.97%) than for the positive demonstrations (17.33%). This difference between positive and negative demonstrations is significant considering the rows for Non-informative, Not-clear and Informative, $\chi^2(2, N = 2074) = 509.73$, $p < .01$, as well when only using the row for Non-informative and Informative, $\chi^2(1, N = 1768) = 486.39$, $p < .01$.

IV CONCLUSIONS

In this study we showed that there is a difference between people solving and teaching a sensorimotor task. However, even when the difference between the Teaching- and Showing-phases is significant, for the positive demonstrations only a relatively small portion (7.43% and 17.33%) are in the Informative category, making it difficult to predict from which phase a single demonstration was taken. **The novelty of this work is that we show that people perceive a significant higher portion of negative than positive demonstrations as informative.** Further, 58% of the negative demonstrations are informative, indicating that our second hypothesis is also verified. In future work, we will consider the use of negative demonstrations in a LfD framework and want to extend into a direction of predictability and legibility [7].

ACKNOWLEDGMENTS

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 765955, further the authors wish to thank the INSEAD-SU Behavioural Lab team for their administrative support.

REFERENCES

- [1] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 5 (2009), 469–483. <https://doi.org/10.1016/j.robot.2008.10.024>
- [2] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. 2008. Survey: Robot Programming by Demonstration. *Handbook of Robotics*, Chapter 59.
- [3] Cynthia Breazeal, Matt Berlin, Andrew Brooks, Jesse Gray, and Andrea L. Thomaz. 2006. Using perspective taking to learn from ambiguous demonstrations. *Robotics and Autonomous Systems* 54, 5 (2006), 385 – 393. <https://doi.org/10.1016/j.robot.2006.02.004> The Social Mechanisms of Robot Programming from Demonstration.
- [4] Sylvain Calinon. 2019. Learning from Demonstration (Programming by Demonstration). In *Encyclopedia of Robotics*, M. H. Ang, O. Khatib, and B. Siciliano (Eds.). Springer. https://doi.org/10.1007/978-3-642-41610-1_27-1
- [5] Gergely Csibra and György Gergely. 2006. Social Learning and Social Cognition: The Case for Pedagogy. In *Progress of Change in Brain and Cognitive Development. Attention and Performance XXI*. 249–274.
- [6] Yuchen Cui and Scott Niekum. 2018. Active Reward Learning from Critiques. In *Proceedings - IEEE International Conference on Robotics and Automation*. 6907–6914. <https://doi.org/10.1109/ICRA.2018.8460854>
- [7] Anca D. Dragan, Kenton C. T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 301–308. <https://doi.org/10.1109/HRI.2013.6483603>
- [8] Mark K Ho, Michael Littman, James MacGlashan, Fiery Cushman, and Joseph L. Austerweil. 2016. Showing versus doing: Teaching by demonstration. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 3027–3035. <http://papers.nips.cc/paper/6413-showing-versus-doing-teaching-by-demonstration.pdf>
- [9] Mark K. Ho, Michael L. Littman, Fiery Cushman, and Joseph L. Austerweil. 2018. Effectively Learning from Pedagogical Demonstrations. In *Annual Conference of the Cognitive Science Society (CogSci)*.
- [10] Carl Mueller, Jeff Venicx, and Bradley Hayes. 2018. Robust Robot Learning from Demonstration and Skill Repair Using Conceptual Constraints. In *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*. 6029–6036. <https://doi.org/10.1109/IROS.2018.8594133>
- [11] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Andrew Bagnell, Pieter Abbeel, and Jan Peters. 2018. An Algorithmic Perspective on Imitation Learning. *CoRR* abs/1811.06711 (2018). arXiv:1811.06711 <http://arxiv.org/abs/1811.06711>
- [12] Giovanni Pezzulo, Francesco Donnarumma, and Haris Dindo. 2013. Human sensorimotor communication: A theory of signaling in online social interactions. *PLoS ONE* (2013). <https://doi.org/10.1371/journal.pone.0079876>
- [13] Giovanni Pezzulo, Francesco Donnarumma, Haris Dindo, Alessandro D’Ausilio, Ivana Konvalinka, and Cristiano Castelfranchi. 2019. The body talks: Sensorimotor communication and its brain and kinematic signatures. *Physics of Life Reviews* 28 (2019), 1–21. <https://doi.org/10.1016/j.plrev.2018.06.014>
- [14] Patrick Shafto, Noah D. Goodman, and Thomas L. Griffiths. 2014. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology* 71 (2014), 55–89.