A Decision-Making Architecture for Human-Robot Collaboration: Model Transferability

Mehdi Sobhani[1][a], Jim Smith[2][b], Anthony Pipe[3][c] and Angelika Peer[4][d]

[1]*Department of Engineering Mathematics, University of Bristol, Bristol, U.K.*

[2]*Department of Computer Science and Creative Technologies, University of the West of England, Bristol, U.K.*

[3]*Bristol Robotics Laboratory, University of the West of England, Bristol, U.K.*

[4]*Faculty of Engineering, Free University of Bozen-Bolzano, Bozen-Bolzano, Italy*

Keywords: Decision-Making, Joint Action, Human-Robot Interaction, Internal Simulation.

Abstract: In this paper, we aim to demonstrate the potential for wider-ranging capabilities and ease of transferability of our recently developed decision-making architecture for human-robot collaboration. To this end, a somewhat related but different application-specific example from the generic one used in its development is chosen, a toy car assembling task in which a participant works together with a robot to perform the assembly task. In a "Wizard of Oz" fashion, a comparison is made between the participant's reactions to working with the robot being controlled either by our architecture or by a human "Wizard" who is hidden from view. With regard to the generalisability of the architecture, we also wish to investigate whether specific models trained on the observed human behaviour in a generic assembly task also transfer to this more complex task. Therefore, pre-trained interaction models from a prior generic pick-and-place task are used again in this new application without any re-training. The architecture was implemented on a robotic arm. Participants worked with the robotic arm to perform the task of picking toy car parts one by one and assembling the car while collaborating with the robot. Each participant repeated the task 3 times for each condition, Model or Wizard, in a random order. At the end of each trial participants completed a PeRDITA questionnaire. First, a test to rule out significant differences was performed, which yielded no significant results for any of the subjective and objective measures. As not having a significant difference does not necessarily mean similarity of conditions, to check for similarity, a Bayesian comparison of the conditions was performed next, which indicated a high probability of similarity between the model and Wizard performance. The high similarity to human-like performance observed for this more complex task supports the claim for the transferability of the models trained on a more generic task.

1 INTRODUCTION

Achieving a natural and seamless human-robot interaction (HRI) has been a continuing challenge for roboticists. To tackle this challenge, some tried developing a robot adaptive behaviour to adjust to different situations (Van Zoelen et al., 2020; Mitsunaga et al., 2008; Kumar and Sahin, 2017; Nikolaidis et al., 2017). For mobile robots, attempts have been made to create a natural HRI by improving the robot planners to work in an uncertain situation and produce more socially acceptable trajectories (Ong et al., 2010; Truc et al., 2022). Further, there are works on learning the human affective state to adapt the robot behaviour accordingly (Churamani et al., 2022). However, there has been little work exploring the idea of a co-active human-robot collaboration in which the robot has the required cognitive abilities to attempt understanding of its partner and then adjust its behaviour accordingly.

To address this challenge, a robot needs to be equipped with human-like cognitive abilities. These required abilities include perspective-taking (Trafton et al., 2005), understanding affordances (Moratz and Tenbrink, 2008), forming expectations of the next action (Lohse, 2011), and timing ability (Yamazaki et al., 2008; Chao and Thomaz, 2011).

For all of them, the ability to anticipate the partner's decisions, and any associated future actions, is

[a] https://orcid.org/0000-0001-5138-134X
[b] https://orcid.org/0000-0001-7908-1859
[c] https://orcid.org/0000-0002-8404-294X
[d] https://orcid.org/0000-0002-2896-9011

Sobhani, M., Smith, J., Pipe, A. and Peer, A.
A Decision-Making Architecture for Human-Robot Collaboration: Model Transferability.
DOI: 10.5220/0012210600003543
In *Proceedings of the 20th International Conference on Informatics in Control, Automation and Robotics (ICINCO 2023) - Volume 1*, pages 719-726
ISBN: 978-989-758-670-5; ISSN: 2184-2809

of utmost importance. To this end, in our previous work (Sobhani et al., 2023b), we developed a decision-making architecture for human-robot collaboration (Figure 1). In this architecture, the robot runs an internal simulation to determine its own decisions and to estimate the decisions of the partner given a specific task to be accomplished. The architecture defines a series of policies (like those presented in (Sobhani et al., 2023b)) that are thought to be underlying the agent's decisions and integrates them into an overall decision-making policy model. These policy models are then integrated by the policy integrator to take one's own decisions. In parallel, the models are used to obtain predictions of the partner's decisions through an internal simulation. The outcome of these simulations is then fed into the negotiation layer where potential conflicts are resolved before reaching a final decision at the end of the internal simulation at each time step.

In this paper, we aim to demonstrate the transferability of these already trained models to the more application-specific task of assembling a toy car. All the trained models were used for the new task without any retraining. Here, the order of the assembling car parts was mapped to the colour order in the colour policy model and the pre-trained distance policy did not need any modification, since it is an underlying part of any task involving some aspects of collaborative pick-and-place. The policy integrator layer and the negotiation layer also did not need any re-training and were used as trained in the previous work.

Comparing our architecture to developments in the state of the art, we can conclude that it differs from most work in that it models behaviour observed in joint action rather than focusing on developing cognitive abilities for individual robotic systems.

For example, several cognitive architectures in the literature, like ACT-R (Anderson et al., 2004), Soar (Laird, 2012) or R-CAST, which are based on Recognition Primed Decision (RPD) models (Fan et al., 2005), are developed for individual agents. These architectures are either based on declarative memory retrieval using instance-based models or rule-based (ACT-R), or they are probabilistic modelling approaches like decision trees (Soar). In our proposed architecture the decision policies were modelled and integrated for individuals using DNFs, and resulting conflicts in joint action were resolved or prevented by the implementation of the negotiation layer.

On the contrary, Bicho et al. (Bicho et al., 2011) presented one of the few existing works that proposes a decision-making system for joint action. They also used DNFs, however, in their work decision policies were hard-coded rather than taken from human experimental data. The workspace was divided into two sides, each covered by only one of the agents allowing to predict actions to be performed by a co-actor. They assume that objects in the area closer to each actor (human or robot) will only be picked by the nearest actor. In contrast, in our work we consider the workspace to be shared equally, since we observed in our human-human collaboration experiments that people did not necessarily act based on the assumption of a divided workspace and, indeed, reached into their partner's area for picking objects. Furthermore, their decision-making system was only tested in joint action scenarios that involve serial actions with collaborators taking turns and performing complementary actions. This reduces potential conflicts significantly, while our proposed architecture has been developed based on both serial and parallel actions with a negotiation layer to resolve conflicts. This means, if there are no physical constraints or limitations imposed by the shared plan, the actor can perform an independent action in parallel to his/her/its partner. A comprehensive comparison to previous work is presented in (Sobhani et al., 2023b).

Designing an experiment with a more realistic task, we aim to evaluate the performance of our recently developed architecture for a more complex task to show the transferability of the models trained based on a rather generic pick-and-place task. For this, we implemented our architecture on a robotic arm and asked human operators to collaborate with the robot for assembling parts of a toy car. We used a Wizard of Oz test in which a human was making decisions and the robot implemented those decisions during the collaboration as a baseline and compared it to the performance of our architecture when it was used by the robot to make decisions.

In addition to objective measures, we evaluated participants' answers to the PeRDITA questionnaire (Devin et al., 2018) taken at the end of each trial. Results reported in Section 4 and discussed in Section 5 indicate that our introduced architecture achieves similar performance as the baseline without requiring retraining of the decision policy models that have been already trained based on a generic pick-and-place task.

2 TRAINED ARCHITECTURE

Our novel decision-making architecture (Figure 1) enables a robot to run an internal simulation to determine its own decisions and to estimate the decisions of its partner given a specific shared task to be accomplished. The architecture defines a series of policies

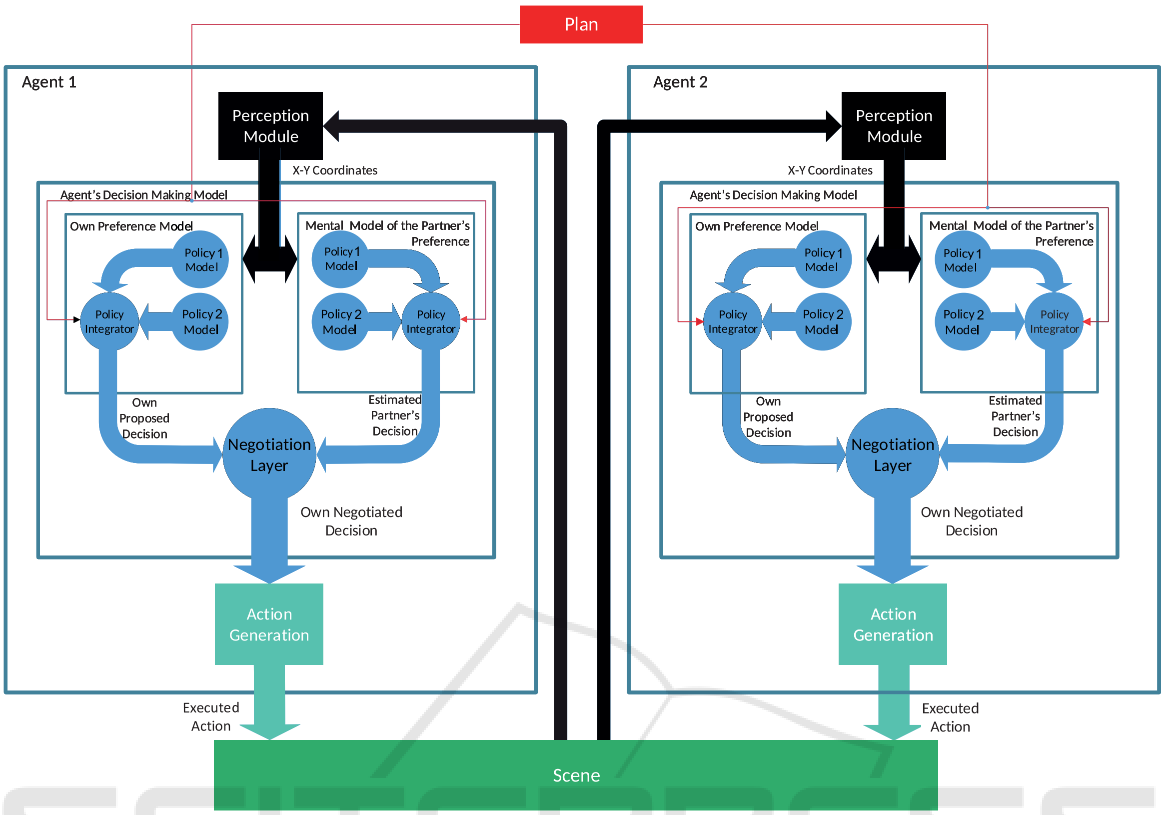

Figure 1: The decision-making architecture for a co-active joint action in human-robot collaboration presented in our previous work (Sobhani et al., 2023b).

that are thought to be underlying the agent's decisions based on the shared plan and integrates them into an overall decision-making policy model. The outcome of these internal simulations of own and partner's decision is then fed into the negotiation layer where potential conflicts are resolved before reaching a final decision at each time step. The negotiation layer works as an implicit communication as, after the actions of both agents are updated in real-time, these updated actions again trigger a new outcome of the internal simulations of both agents. Several modelling techniques were considered to develop the decision policy models, the policy integrator layer, and the negotiation layer (Sobhani et al., 2023b), and Dynamic Neural Field (DNF) was the final choice (Amari, 1977; Schöner, 2008). In developing the architecture, a generic pick-and-place task based on a shared plan with a colour order of objects was used. This required models of Colour and Distance policy along with the policy integrator and the negotiation layer to be trained based on human-human interaction data. The structure of the used models is explained next.

2.1 Structure of Distance Policy

For modelling *Distance policy*, the table-top setup of the experiment is mapped into a 2D DNF. The DNF equations and its parametrers are presented in (Sobhani et al., 2023b).

The projected position of the centre point of each car part and the wrist position of the participants' wrists is mapped on the x-y plane. The x and y axes are then used as features so that each x-y coordinate of the objects and wrist is considered the position of an input stimulus to the neural field. Each stimulus is modelled with a 2D Gaussian and the interaction of these stimuli changes the field activation level in different locations as the input stimuli change due to the agents' motions. The parameters to be learned for this setting are mainly interaction kernel parameters. Having properly trained the model, interaction kernel parameters can change the neural field behaviour such that the response to stimuli will result in an activation of the field at the point of interest, respectively the location of the most likely chosen part.

2.2 Structure of Colour Policy

The *Colour policy* model, is a 1D DNF coupled with a memory trace as presented in (Sobhani et al., 2023b).

The training for this structure is like memorising the colour order by demonstrating the order and showing coloured objects one by one. The memory then forms pre-shapes for the colour order. This model structure is similar to the work by Sandamirskaya and Schöner (Sandamirskaya and Schöner, 2010), but implemented in a way that the neural field stays activated to wait in the order until all objects of the same colour are removed by the participant(s) before moving to the next colour in the order. This is done to simulate tasks with an equal priority of actions in the plan. The parameters of the Colour Policy DNF are chosen to be the same as the ones reported in (Sandamirskaya and Schöner, 2010).

2.3 Structure of Policy Integrator

Having different policies modelled separately, and integrated through the policy integrator layer, makes the architecture adaptable to different tasks. For the task at hand, to have a correct prediction on the chosen part, the colour policy model is coupled with the distance policy model. This provides a measure to decide when there exist multiple objects of the same colour. This means that the colour policy creates a short list of the objects to be picked and the distance policy model predicts which one will be picked up. This is done by having a DNF similar to the *Distance policy* with only shortlisted objects as stimuli being implemented and the final object is chosen from the shortlist according to the distance policy. This process occurs naturally in the DNF of the policy integrator as the amplitude of the input stimuli from the output of the colour policy and distance policy models will intensify the neural field activation for the chosen object.

2.4 Structure of Negotiation Layer

A simulation of the predicted partners' actions runs simultaneously with the 'own' model in the negotiation layer. The aim of this simulation is to adjust the robot's own decision to the predicted partner's decision so as to prevent any conflicts like picking up the same object. This will also adjust the decision based on the plan, so, if the model predicts that the partner would perform the next step, such as when a partner is reaching quicker to an object, the agent should either move on to the next action or wait for the appropriate moment to perform the next action. This is done by inhibiting its own decisions when the model predicts that the partner will perform the same action, or by exciting the decision when it predicts that the partner is waiting or performing another action. To achieve this, the negotiation layer is implemented using a 2D DNF similar to the *Distance policy* and the interaction kernel of two DNFs of the own agent model and the partner model, is adjusted based on the human-human interaction experiment. This means the desired outcome is achieved by learning when each DNF should be inhibited (activation function being locally or globally deactivated) or excited (activation function either locally or globally being further activated). It is noteworthy that unlike works such as the one by Devin and Alami (Devin and Alami, 2016) that utilises a dialogue system, here the negotiation layer works as an implicit communication, as after the actions of both agents are updated in real-time, these updated actions again trigger a new outcome of the internal simulations of both agents.

3 METHOD

An experiment was designed in which the decision-making models developed based on a generic task were to be tested for their adaptability to another more application domain-specific and complex task like assembling parts of a toy car. All the developed models could be used in this experiment without requiring retraining. The order of the assembling part was only mapped to the colour order in the colour policy. The experiment was designed similarly to the previous experiments to have two conditions namely, the "Wizard of Oz" and the "Model" for when the robot used the proposed decision-making architecture. Each participant repeated the task 3 times for each condition without being informed of the conditions. The order of the trials was randomised. Tests were first carried out to rule out significant differences. In case of no significance, these tests would be followed by a Bayesian comparison of the conditions to indicate the probability of similarity of the model and Wizard performance.

3.1 Experimental Setup

The experimental setup was similar to our previous human-robot interaction study (Sobhani et al., 2023a). Participants were asked to work with a Franka Emika Panda robotic arm in a toy car assembling task. Their dominant hand was marked for tracking by Vicon motion capture reflective balls. The robot control was done in real-time using the libfranka library in C++. The decision-making model was run-

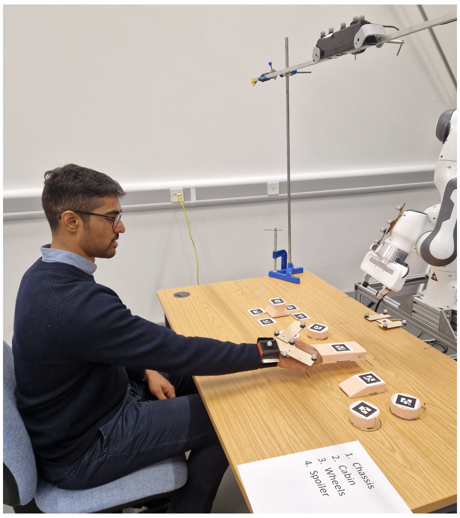

Figure 2: Experiment Setup.

ning in parallel in MATLAB and the decision was communicated to the robot controller through TCP/IP socket communication. This means that the robot controller was not continuously receiving human decision information but at specific times, hence, the robot could not change its behaviour when one motion was being implemented. Participants repeated the same task in two conditions. The Experiment setup is shown in Figure 2. In the human decision-maker condition, the wizard is hidden behind a partition wall and sees the scene from the MS Kinect camera mounted over the table to also track the parts marked by AR markers.

3.2 Wizard Protocol

To ensure consistent behaviour of the human decision-maker, the following protocol was given to the wizard to follow: The wizard makes the first decision as soon as the robot is in the "ready position". This is done by entering a number between 1 to 7, 1 for the car chassis, 2, for the car cabin, 3 to 6 for the wheel depending on the location of the wheels on the table and 7 for the spoiler, from the robot perspective, or entering 8 for just waiting. Later the wizard needs to confirm this decision when the robot goes to the "ready-to-grasp" position. As soon as the robot reaches the "ready-to-grasp" position, the wizard must make the final decision within one second or as soon as the human participant makes his/her move (whichever is faster). At this stage, the wizard has 3 options, either to confirm the decision and continue picking the same part, change to another part, or just wait. If the decision is waiting, the wizard needs to make another decision again within one second until a part is chosen. This will continue until the task is complete.

3.3 Task

The task was to assemble parts of a toy car including the Chassis, Cabin, Wheels and Spoiler. The order of the assembly was given to the participants to follow as 1. Chassis, 2. Cabin, 3. Wheels, and 4. Spoiler. While participants were following this order, they were free to choose the wheels in any order when it was the right time to do so. This assembly task is exemplary for a variety of tasks a robot might be used for.

3.4 Participants

In total, 16 new participants (13 male) took part in the experiment. Participants were students and staff members of the university with a mean age of 29.8 (STD= 6.14) ranging between 22 and 43 years old with an average height of 175.6 cm (STD=7.43) ranging between 164 and 188 cm. Hereby, 14 participants were right-handed, one was ambidextrous but used his left hand and one was left-handed, and all reported normal or corrected-to-normal eyesight (8 wearing glasses). Participants took part in the experiment voluntarily and all gave informed consent. Ethical approval for this experiment was obtained from the ethics committee of the University of the West of England (reference number: UREC16-17.03.10).

3.5 Subjective and Objective Measures

In terms of objective measures, the task completion time, the robot task share, and the number of conflicts were recorded. The robot task share is calculated as the number of parts the robot picks divided by the total number of parts required (7) for assembling the car. A conflict is considered to have occurred when the robot tries to pick the part that the participant has just picked or is picking up.

As there is no verbal communication between the robot and participants, only four dimensions of the PeRDITA questionnaire were used: "Collaboration", "Interaction", "Robot Perception", and "Acting".

4 RESULTS

First, results of the PeRDITA questionnaire (subjective measures) are presented. The results obtained for the human decision-maker are compared to the model in Figure 3. The mean is calculated by averaging over 3 trials of each condition for each participant of the

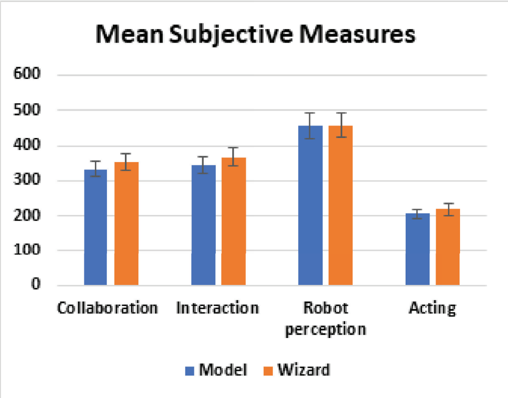

Figure 3: Bar graph for mean value of subjective measures from the PeRDITA questionnaire. Error bars are $\pm 1 SEM$.

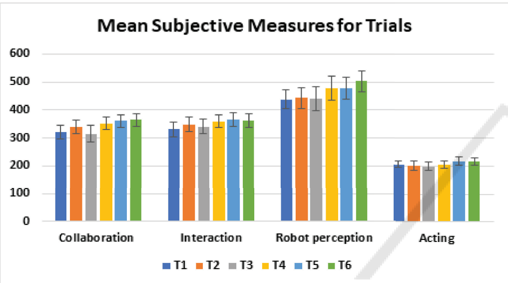

Figure 4: Bar graph for mean value of subjective measures from the PeRDITA questionnaire based on order of the trials. Error bars are $\pm 1 SEM$.

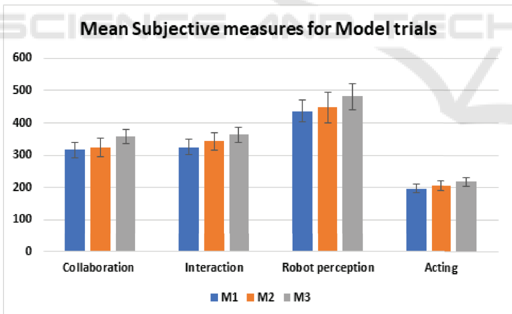

Figure 5: Bar graph of mean subjective measures for 3 trials with model decision-maker. Error bars are $\pm 1 SEM$.

16 participants. Despite a slightly higher score for the human decision-maker across the board, no significant difference was found in the subjective measures. The results are also presented based on the order of the trials in Figure 4. Similarly, the data was sorted based on the order of the trials M1, M2, M3 in each condition, see Figures 5 and 6. An overall improvement in scores can be observed, however, no significant difference was found.

The results for three objective measures namely, Task Completion Time, Robot Task Share, and Con-

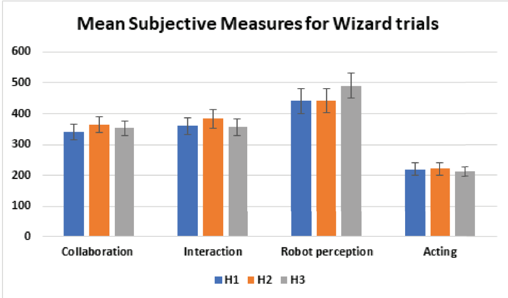

Figure 6: Bar graph for mean subjective measures for 3 trials with human decision-maker. Error bars are $\pm 1 SEM$.

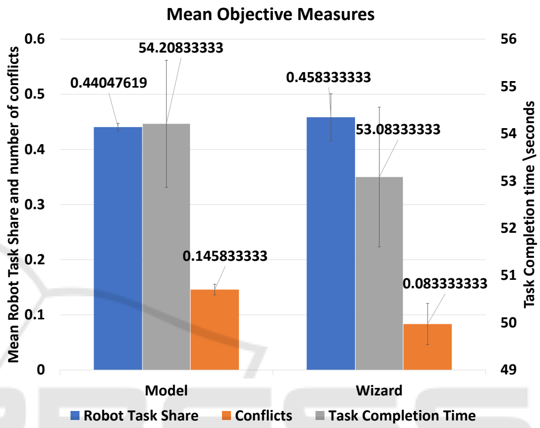

Figure 7: Mean value of the objective measures depicted in bar graphs for Robot Task Share, Conflicts and Task Completion Time. Error bars are $\pm 1 SEM$ (standard error of mean).

flicts are presented in Figures 7. No significant difference was found when comparing the model against the human decision-maker for any of the measures.

Looking at the order of trials, the data shows a significant difference ($F = 2.57, p < 0.0323$) for task completion time between the first and sixth trials (regardless of condition), see Figure 8. When data is sorted based on the order of trials within each condition no significant difference was found for the condition using the model, but for the human decision-maker, a significant difference ($F = 4.04, p < 0.0243$) between the first and third trial in the task completion time was observed.

Since most of the tests didn't show any significant differences and since a test for difference does not allow making any statement about similarity, a Bayesian comparison (Benavoli et al., 2017) of the data for the two conditions; the "Wizard of Oz" and the "Model", was performed. The data was normalised for this analysis based on the maximum values of each measure. The comparison results are depicted in Tables 1 and 2. They show a high likelihood

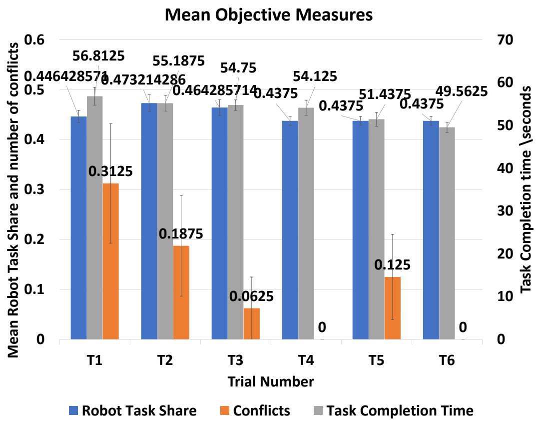

Figure 8: Mean value of the objective measures depicted in bar graphs for Robot Task Share, Conflicts and Task Completion Time. Error bars are $\pm 1 SEM$ (standard error of mean).

Table 1: Bayesian Comparison for the objective measures.

Measure	Model is better	Difference is negligible	Human is better	Rope
Robot Task Share	0.000364	0.9996192	1.680241e-05	0.1
Task Completion Time	1.555164e-05	0.99992125	7.719776e-05	0.1
Conflicts	0.0002701	0.99970587	2.4007e-05	0.5

Table 2: Bayesian Comparison for the subjective measures.

Measure	Model is better	Difference is negligible	Human is better	Rope
Collaboration	9.5738e-05	0.975926	0.02397785	0.1
Interaction	8.44871e-05	0.958927	0.040988432	0.1
Robot Perception	0.00182163	0.9954326	0.002745791	0.1
Acting	0.00205497	0.922336	0.0756091	0.1

of similarity between the two conditions for all subjective and objective measures.

5 DISCUSSION

The presented car assembling HRC experiment was designed to demonstrate the transferability and expandability of the proposed architecture for human-robot collaboration presented in our previous work. To do so, no model was retrained and the order of the assembling task was mapped to the already learnt colour order in the colour policy. The distance policy was used as it was previously trained.

Results indicate that when data was averaged and sorted per conditions, the human decision-maker scored descriptively slightly better than the model in all the objective and subjective measures, but no significant differences were found. Having ruled out eventual significant differences and since tests for differences don't allow any statement on similarities, a further Bayesian comparison of the data was performed, which revealed a high probability of similarity between the two conditions. This high similarity of the model to human-like performance found for the case of the studied more complex task, supports the claim of the transferability of the models trained on a more generic task, which was the main aim of the performed study.

Sorting the data in the order of the appearance of trials, a significant difference was found in the task completion time between the first and the last trial. The significant difference was also observed when only looking at the human decision-maker trials, however, not for the model trials, although the time descriptively improved over trials. The reason for this decrease in the task completion time from the first trial to the last could be due to learning effects as the same task was repeated 6 times during this experiment. This learning effect was prevented in the previous HRC experiments (Sobhani et al., 2023a) as the task was changing randomly for trials and in each trial, participants were making a different alphanumeric character. Having a significant change in the Task Completion Time for the Wizard condition and lack of this significant difference in the model condition could also indicate that participants could adapt better to the robot in the wizard condition. However, further investigation with more participants is required for a concrete conclusion here.

Finally, comparing the tasks the models were trained on with this task, fewer conflicts were recorded. This could be because of having fewer parallel actions compared to previous experiments. In the task used for data collection and training of models, there were 3 pairs of wooden blocks of the same colours giving them the same priority to be picked in parallel, while in the car assembling task, only wheels had the same priority. Nonetheless, having several ways of assembling wheels in this task and the success of the architecture in dealing with this uncertainty without any specific prior training could be a good indication of the models' transferability.

Overall, we can conclude that our results indicate a good potential for transferability of the models trained on a generic task to different more application-specific tasks if they shared some characteristics, such as serial orders. In this context, it is important to note, that the Colour policy was trained for an order of four colours and the car assembling task had also four main steps. However, this does not resemble a limitation of the model as it is also possible to map unequal sizes of orders to the "Colour policy". If required to add new steps to the order, for example more objects, they can be presented to the model to create a new stimulus, which can be added to the order.

REFERENCES

Amari, S.-i. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological cybernetics*, 27(2):77–87.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, 111(4):1036.

Benavoli, A., Corani, G., Demšar, J., and Zaffalon, M. (2017). Time for a change: a tutorial for comparing multiple classifiers through bayesian analysis. *Journal of Machine Learning Research*, 18(77):1–36.

Bicho, E., Erlhagen, W., Louro, L., and e Silva, E. C. (2011). Neuro-cognitive mechanisms of decision making in joint action: A human–robot interaction study. *Human movement science*, 30(5):846–868.

Chao, C. and Thomaz, A. L. (2011). Timing in multimodal turn-taking interactions: Control and analysis using timed petri nets. *Journal of Human-Robot Interaction*, 1(1):1–16.

Churamani, N., Barros, P., Gunes, H., and Wermter, S. (2022). Affect-driven learning of robot behaviour for collaborative human-robot interactions. *Frontiers in Robotics and AI*, page 20.

Devin, S. and Alami, R. (2016). An implemented theory of mind to improve human-robot shared plans execution. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 319–326.

Devin, S., Vrignaud, C., Belhassein, K., Clodic, A., Carreras, O., and Alami, R. (2018). Evaluating the pertinence of robot decisions in a human-robot joint action context: The perdita questionnaire. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 144–151.

Fan, X., Sun, S., and Yen, J. (2005). On shared situation awareness for supporting human decision-making teams. In *AAAI Spring Symposium: AI Technologies for Homeland Security*, pages 17–24.

Kumar, S. and Sahin, F. (2017). A framework for an adaptive human-robot collaboration approach through perception-based real-time adjustments of robot behavior in industry. In *2017 12th System of Systems Engineering Conference (SoSE)*, pages 1–6.

Laird, J. E. (2012). *The Soar cognitive architecture*. MIT press, Cambridge, Massachusetts.

Lohse, M. (2011). The role of expectations and situations in human-robot interaction. *New Frontiers in Human-Robot Interaction*, pages 35–56.

Mitsunaga, N., Smith, C., Kanda, T., Ishiguro, H., and Hagita, N. (2008). Adapting robot behavior for human–robot interaction. *IEEE Transactions on Robotics*, 24(4):911–916.

Moratz, R. and Tenbrink, T. (2008). Affordance-based human-robot interaction. In *Towards Affordance-Based Robot Control*, pages 63–76. Springer.

Nikolaidis, S., Hsu, D., and Srinivasa, S. (2017). Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634.

Ong, S. C., Png, S. W., Hsu, D., and Lee, W. S. (2010). Planning under uncertainty for robotic tasks with mixed observability. *The International Journal of Robotics Research*, 29(8):1053–1068.

Sandamirskaya, Y. and Schöner, G. (2010). An embodied account of serial order: How instabilities drive sequence generation. *Neural Networks*, 23(10):1164–1179.

Schöner, G. (2008). Dynamical systems approaches to cognition. In Sun, R., editor, *The Cambridge Handbook of Computational Psychology*, pages 101–126. Cambridge University Press.

Sobhani, M., Giuliani, M., Smith, J., Pipe, A., and Peer, A. (2023a). Evaluating a decision-making architecture in human-robot collaboration experiments. *International Journal of Social Robotics*. (under revision).

Sobhani, M., Smith, J., Pipe, A., and Peer, A. (2023b). A novel mirror neuron inspired decision-making architecture for human-robot interaction. *International Journal of Social Robotics*.

Trafton, J. G., Cassimatis, N. L., Bugajska, M. D., Brock, D. P., Mintz, F. E., and Schultz, A. C. (2005). Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 35(4):460–470.

Truc, J., Singamaneni, P.-T., Sidobre, D., Ivaldi, S., and Alami, R. (2022). Khaos: a kinematic human aware optimization-based system for reactive planning of flying-coworker. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 4764–4770.

Van Zoelen, E. M., Barakova, E. I., and Rauterberg, M. (2020). Adaptive leader-follower behavior in human-robot collaboration. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1259–1265. IEEE.

Yamazaki, A., Yamazaki, K., Kuno, Y., Burdelski, M., Kawashima, M., and Kuzuoka, H. (2008). Precision timing in human-robot interaction: coordination of head movement and utterance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 131–140. ACM.