

Towards A Multidimensional Perspective on Shared Autonomy

Malte Schilling, Stefan Kopp, Sven Wachsmuth, Britta Wrede, and Helge Ritter

Center of Excellence ‘Cognitive Interaction Technology’, Bielefeld University, Germany

Email: mschilli@techfak.uni-bielefeld.de

Thomas Brox, Bernhard Nebel, and Wolfram Burgard

Department of Computer Science at the University of Freiburg, 79110 Freiburg, Germany

Abstract

Shared Autonomy in the traditional sense focuses on the degree of user intervention in the control of artificial systems. We propose to broaden this notion to allow for more interactive scenarios. This requires a shift away from the single system perspective towards the interaction, the participating agents and the cooperation as such. Such a view on the interaction of autonomous agents has to be based on a more fine-grained understanding. Therefore, we extend a differentiation of autonomy into three different levels to interactive tasks as a starting point for a multidimensional perspective on shared autonomy. In particular, we want to point out how this allows for flexible interaction patterns and the negotiation of changing roles in ongoing cooperation.

Introduction

Autonomous systems are capable of organizing the way they behave by themselves, ranging from action execution, over strategic planning to goal selection. To do so, they have to exceed automatic and predefined behaviors and shape their interaction with the dynamic environment by striving for satisficing rather than optimal solutions. Full autonomy still poses a problem for artificial systems that shall deal with difficult and complex tasks. Therefore, in *shared autonomy* a system is supposed to cope autonomously with the task as good as possible (Kanda and Ishiguro 2012). In such settings, usually, the system deals with the low level details of the execution of action, while, on a higher level and whenever the system is not able to deal with the current complexity or unforeseen events, the control of the system is transferred to a user. In this classical view, shared autonomy is understood as a case in between fully autonomous behavior and teleoperation. The term *shared* refers to the aspect that the actions of the system are controlled either by the system or transferred to the user.

From our point of view this leaves out many forms of collaboration. We rather think that the notion of shared autonomy should be understood in a richer sense that allows for cooperation and interaction. Cooperative scenarios provide multifaceted interaction possibilities between participating agents with a large number of degrees of freedom on different levels of abstraction to achieve a goal. Autonomy for a

single system is to organize its behavior given these degrees of freedom, but of course also respecting the constraints as given through the scenario or situation. In shared autonomy this process of shaping behavior has additional degrees of freedom and constraints induced by the cooperation with another autonomous agent and is mediated through communication. Shared autonomy becomes a joint process: the participating agents introduce further constraints for each other and have to respect the constraints introduced by the other agents while at the same time attaining additional degrees of freedom through additional competences and options.

Therefore, shared autonomy requires—in our understanding—always that every participant has to concede some degrees of freedom in order to allow others some autonomy. But at the same time every participant also gains in autonomy in the sense that s/he gets more choices for action and more goals s/he can achieve. Importantly, such scenarios are usually so rich that it is not feasible to reason about or control all possibilities. However, from the perspective of shared autonomy this complexity is actually an asset that can be exploited when multiple agents (technical systems or humans), that are capable of shared autonomy, are introduced and can take over and deal with parts of the complexity. Shared autonomy is thus the setting in which constraints are mediated between agents for achieving their goals.

In this article, we want to briefly discuss the currently used notions of semi-autonomous systems and shared control. Second, we want to review the different levels of autonomy as proposed by Gransche et al. (2014) and we want to extend this perspective to cooperative scenarios involving multiple actors. Further, we want to discuss the adaptivity of shared autonomy and how it emerges as well as evolves throughout interactions. Last, we will provide evaluation criteria and will briefly point out exemplary scenarios as benchmarks.

Classical View of Shared Autonomy

Autonomy of artificial systems is a graded quality. This leads naturally to a degree of autonomy which can be traced back to Sheridan’s (Sheridan and Verplank 1978) ten levels of automation (a detailed summary is given in (Hertkorn 2016, chapter 2)). In this view, there are different degrees of autonomy given along a single dimension ranging from di-

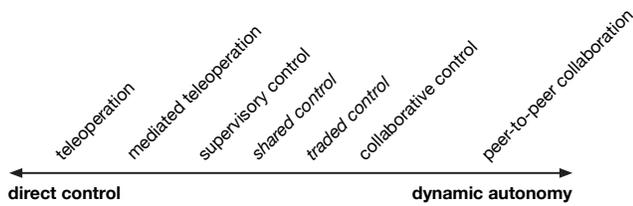


Figure 1: Different degrees of autonomy ranging from teleoperation to autonomy (after (Goodrich and Schultz 2007, Fig. 4.2); added have been additional notions following (Goodrich, Crandall, and Barakova 2013) shown in italics).

rectly operated systems towards fully autonomous systems (shown in Fig. 1 following (Goodrich and Schultz 2007) and (Goodrich, Crandall, and Barakova 2013, Table 5.2)). *Shared autonomy* is the general notion which takes a human-centered perspective. The focus is on the degree of control and supervision of the system through a human (Hertkorn 2016). While the artificial system tries to cope autonomously with the task as good as possible (Kanda and Ishiguro 2012, chapter 6), it usually deals with the low level details of the execution of action. But, on a higher level—and whenever the system is not able to deal with the current complexity or unforeseen events—the control of the system is transferred to a user. *Shared autonomy* is mostly used as the general term. There are multiple different more fine-grained notions of semi autonomous systems which are used more often than the more general term. In the following, we want to give a brief overview of these notions.

Shared/Guided Control puts the focus on the control originating from the user towards the system which has its own control loop and is autonomously reacting to the environment and executes the specified action (Goodrich, Crandall, and Barakova 2013). Still, user and robot work concurrently (in contrast to *traded control* in which they take turns). In *Collaborative Control* user and system share a task and work as a team collaboratively in the same space and at the same time ((Goodrich, Crandall, and Barakova 2013) call this *Mixed-Initiative Control* in which the focus is more on the flexible interaction strategy.). In contrast, in *Supervisory Control* (Endsley and Kaber 1999) the user only monitors the execution of the autonomously working system.

In general, in the perspective of shared autonomy a system is not recognized as a partner in a task, but is more seen as an intelligent tool with the advantage that the system can act autonomously to a certain degree and the user can be freed from many (distant) details of the action. As shared autonomy involves a form of interaction with a user or teleoperation, it requires input and output to the user. Feedback has to be provided to the user in the same way this is required in a teleoperated system. The user should be immersed in the task in order to allow for proper control. Interaction between user and system therefore usually takes place only on a higher level of abstraction (in contrast to teleoperation).

To summarize, shared autonomy is understood as referring to a semi-autonomous system (Vernon 2014) and such a system is, first of all, teleoperated. But, secondly,

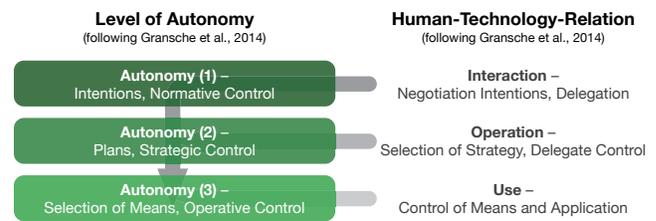


Figure 2: On the left, the three dimensions of autonomy are shown as proposed by Gransche et al. (2014). The filled rounded rectangles visualize the space of possible selections following Bradshaw et al. (2003). On the right, the view of Gransche et al. (2014) on the human relation towards the technical system is summarized.

it is carrying out given actions autonomously (or semi-autonomously). From our point of view this is where the current notion of shared autonomy is lacking. Viewing autonomy as a one dimensional quality is leaving out the fine grained structure of underlying interactions and involved representations. We think it is important to incorporate a rich notion of interaction into the concept of shared autonomy. From our point of view, in an interaction different roles are negotiated and can change over time. The degree of autonomy is therefore also not fixed (comparable to the subnotions *adaptive*, *sliding* (Fong, Thorpe, and Baur 2003) or *adjustable autonomy* (Goodrich et al. 2001)). A multidimensional perspective is required which sees autonomy as a process operating on multiple levels and relying on complementing representations on the different levels (there is already work that relates shared autonomy to representations on different levels, as goals and intentions). Interaction between autonomous systems necessitates to mediate on these different levels of autonomy. We propose that such a detailed view on autonomy can help to refine the concept of shared autonomy for rich interactive scenarios.

Beyond the Single-Dimensional View

We strongly argue that the single-dimensional view on shared autonomy should be extended to better embrace the richness of phenomena that can emerge when autonomous agents engage in interaction. In our view, this requires to open the narrow focus on handling the passing and regaining of control towards a more multidimensional perspective that tries to elucidate the space of interaction patterns that can arise when two or more autonomous agents come together.

An interesting attempt along this line has been put forward by Gransche et al. (2014). They propose a stratification of autonomy into three levels: The first type of autonomy (*freedom of intentions*) allows choosing ones own particular purpose in a current situation. Autonomy of the second type determines a strategy in order to fulfill certain purposes (*freedom of decision*). On the lowest level (their third type of autonomy: *freedom of acting*), autonomy describes the ideally optimal selection of a means to achieve an immediate goal. Conceptually, these levels characterize the freedom of making choices arising at different levels of a hierarchy.

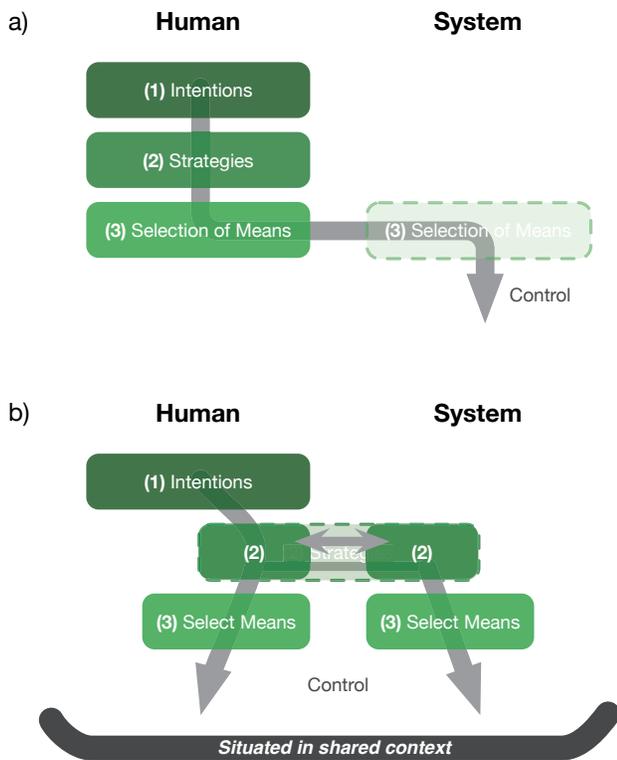


Figure 3: Application of the multidimensional perspective of autonomy on specific cases along the shared autonomy spectrum. Shown is the interaction between user and system (only for cases involving two participants). In a) teleoperation is shown and in b) collaborative control.

The three different levels are shown in Fig. 2. The spaces of possible selections are visualized as rounded rectangles following Bradshaw et al. (2003) who applied such a kind of representation, but only towards the lowest level of potential actions. Importantly, the presented view is simplified as it does not distinguish between capabilities (which actions can be performed) and freedom (which actions are allowed) (Bradshaw et al. 2003). Instead, the simplified figure shall give a rough idea on the possibilities on that level given a current set of constraints (We also exclude the feedback of the system and the shared perception.).

We propose to extend this multi-level notion of autonomy towards interactive settings and scenarios. Our perspective will be that these levels stratify the shared autonomy space and that a more comprehensive characterization of how the autonomy spaces of the agents merge into this shared space calls for additional dimensions for which we will provide some examples. This perspective allows to further differentiate types of interaction and to differentiate changing contributions from different partners of the interaction. As one example, Fig. 3 a) shows a conceptualization for the case of a fully teleoperated system. Importantly, the different levels of autonomy are shown for the human user and the system as well as the interaction between those. In the case

of a teleoperated system, there is no autonomy on the system side. Instead, the intentions of the user are substantiated into a plan and even the means are selected autonomously (but constrained through the higher levels) by the human. The system is under direct control. Moving towards more semi-autonomous operation—as discussed in the previous section—the hand-over between human and system would be moved to the intermediate level and the user would be freed of the details of the action. The system would be autonomous on the lower level and select its means by itself. This might be moved even further towards a form of interaction in which the system is laying out its own plans.

The single dimensional view on the interaction of autonomy and control becomes problematic for cases like collaborative control. But in a multidimensional view this can be conceptualized as depicted in Fig. 3 b). In that case, human and system both work collaboratively in the same space and at the same time. While the user's intentions should guide the system, the system at least requires some autonomy on the lower level in order to free the user enough, so that s/he has free resources allowing her/him to act autonomously. Additionally, communication is required to coordinate the behavior. Ideally, this is not done on a fully-detailed level, but on an intermediate level which requires a form of common representation accessible to both.

In natural interaction fusing of autonomy spaces brings into view a rich set of important dimensions, most of which cut across the aforementioned levels. These dimensions include, for example:

- **Shared resources**, such as space, e.g. avoiding collisions or creating functional contacts, time, e.g. for catching up in a cooperation, or information, e.g. about goals and plans.
- **Efforts**: how does fusing of autonomy spaces impact on processing load, memory load, or physical forces that need to be exerted?
- **Performance measures**: these are task-specific monitoring dimensions, but they include also generic measures such as throughput or failure rates.
- **Action entropy/predictability**: a limited predictability is intrinsic to autonomy. Thus, to facilitate coordination, sharing autonomy requires agents to include in their policy the control of their action entropy, either through explicit communication or in the form of commitments that are transparent to others.
- **Interaction strength** distinguishes between loosely/tightly coupled agents and can, e.g., be monitored through mutual information or correlation measures.
- **Directionality** sharing may occur symmetric or in a directional leader-follower pattern, with correspondingly richer structures for more agents.
- **Adaptivity** characterizes the capability of the interaction to change in history-dependent ways (and will be taken up again in the next section in the context of shared autonomy and emergence).

This list covers a number of dimensions which we believe to be of major importance, but it is not exhaustive.

It reflects that real world scenarios provide a rich and multitude of interaction possibilities between participating agents which offers a large number of degrees of freedom on different levels of abstraction to achieve a goal. An autonomous system behaves inside a set of constraints given by the current situational context. Extended to shared autonomy the organization of behavior is not only constrained by the current situation, but also by the other agents sharing this situation. And, most importantly, the available degrees of freedom that can be chosen by all the agents is shared. The decisions of each individual agent for those degrees of freedom are influencing all the other agents. The possibilities for each system depend on how the others shape their behavior. As a consequence, in a shared autonomy setting autonomy should not be seen with respect to a single system, but always in the context of the whole situation and the other interacting systems. Shared autonomy requires to leave open choices to the other participating agents—and possibly therefore to constrain oneself. As an advantage, the multi-level notion of autonomy offers here a differentiated view on the interaction between systems. It provides different forms of interaction and the agents take up different roles in these interactions. Importantly, first, these roles might change over time and get more complex with multiple agents taking part in an interaction. Secondly, those exchanges are not always unidirectional. This perspective is from its outset devoid of any subordinative hierarchy between the agent (e.g. robot vs. human). Any such relation is an additional structure and such structure can be added in multiple ways. For instance, a rigid subordinative hierarchy, a context-specific leader-follower pattern (e.g. according to experience), or a potentially complex role dynamics (“social dynamics”, with or without learning).

Tightly connected to the three levels of autonomy are representations of different levels of abstraction. Here, Gransche et al. (2014) do not provide details on the kinds of representation and we think that this is one of the important research questions to complement the different levels of autonomy with rich representations capturing the information on that level of abstraction. While the lower level deals with detailed representation of actions and their execution, an intermediate level will deal with plans, sequences or strategies as combinations of those and on the highest level intentions should be represented. This leads to the question how patterns of shared autonomy might be designed so that they are both beneficial and implementable. In the ideal case, this might again be based on some specifiable optimization criterion. However, the optimization of clearly specifiable criteria might be too difficult to implement in practice, and satisficing or heuristic approaches might be needed instead. This shifts the focus from optimization to a direct implementation of coordination patterns, such as “complementarity of X”. Depending on X, this can cover a wide range of situations: coordination by available resources (e.g. information), skills (e.g. experience), or interests (interference minimization). A different (and more “positive”) pattern would be “synergy for X” where X is only enabled as a result of a suitable cooperation. This is a more complex pattern, instead of interference minimization here the goal is to “create more

than the sum of its parts”.

The proposed levels (Gransche et al. 2014)—proper actions, strategies, and intentions for goals—appear to us as an excellent example of how to “cut” interaction space in a principled fashion. And we believe that there are further important and similarly principled “cuts” that all provide relevant perspectives for a deeper understanding of the fabric of shared autonomy. We propose to broaden the classical concept of shared autonomy to embrace a richer set of phenomena for a deeper understanding of how autonomous agents may interact.

Shared Autonomy and Emergence

Autonomous agents very often are capable of adaptivity and learning. This may allow to gradually establish shared autonomy when it is not available at the outset: we then see shared autonomy emerge as a consequence of adaptive changes. This might enlarge the space of possible selections on the different levels. If agents are cooperating then they might achieve more than simply adding up their single autonomous abilities, exceeding their original possibilities (see Fig. 4).

Imagine two agents passing each other every morning on a narrow lane. The agents cannot sense each other, and they can only pass each other without a bump when they choose opposite sides of the lane on their walk. Initially, they act very uncorrelated (e.g. choosing sides of the lane at random). After a particularly painful bump, one of the agent decides to switch to a deterministic choice and stick to the lane side that is opposite to the bump. Unfortunately, the other agent has the same idea after the event and the situation gets worse. But then, one (or both) agents might include some randomness into their deterministic strategy and, after the first bump-free passage, both agents can stick deterministically to the “discovered pattern”.

In the previously discussed approaches, the successful coordination pattern had to be implemented into the agents by a suitable a-priori design, requiring an analysis of the structure of their interaction. The tiny example illustrates, that in the absence of such an analysis, suitable rules for changing the behavior adaptively can make successful coordination patterns emerge. This could, for instance, also have been achieved by some reinforcement learning approach (with bump-free passage as the obvious reward) or other forms of stochastic search. Already much more elaborate examples have been considered, for example, in biology, e.g. when ants lay down odor traces to share successful navigation patterns. The example already illustrates that there can be rules that converge (on average) much faster than purely statistical unsupervised learning. Discovering such rules is of paramount importance in robotics, where abundant interactions are simply too costly to be practical.

Of course, for straightforward situations like the example, any learning approach seems like “overkill”. But imagine adding further context, for example, one agent unconditionally chooses the lane on the shadow side when the sun is burning, while the other agent doesn’t care about the sun. This would make the successful coordination pattern

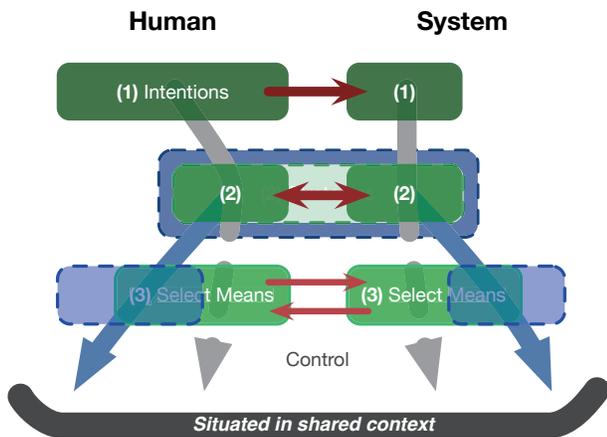


Figure 4: Shared Autonomy for a complex interactive scenario: On the highest level intentions are given to the system. In addition, there is explicit and implicit communication between the human and the system (red arrows) also on the lower levels. The blue space of possibilities signifies the resultant action space in cooperation. On the one hand, this can grow larger, for example as seen on the intermediate level there are more possibilities, even though many of those probably require coordination. On the other hand, this might become shifted, for example as is the case for the lowest level: in a cooperation other means become possible, but also other might not be possible anymore. A simple example might be redecorating a room with large furniture. While a single agent can not move a huge table around, two agents in cooperation might be able to. With respect to the weight, their overall capabilities (indicated by the blue area) might be increased to the sum of what each agent individually is able to carry. But even if one of the agents could lift the weight of the table by himself, only together it might become feasible to maneuver the bulky furniture in a goal-directed way.

context-dependent and increase the adaptivity of the behavior.

Fields such as game theory (“coordination games”), cooperative planning, or reinforcement learning, offer useful tools to craft forms of adaptivity that can lead to emergent shared autonomy. However, in complex tasks contexts and agent value functions (“attention”, “interest”, “mood”), as well as agent roles, can shift, switch or interact in a multitude of ways, depending on the given task and situation, but also on the interaction partner or what is known about him and possible cooperation. Since most real world interactions have to converge within a very limited number of interactions, pure observation based identification or learning needs to be complemented with additional strategies for fast coordination, such as some form of implicit or explicit negotiation. This involves communication and should be continuous to allow tracking of role changes and to make it possible that all participants can contribute according to their capabilities. For example, when tidying up, a child and a robot

might engage in placing objects in different boxes to stow those away, with the child leading the activity and the robot falling into the role of the follower, just selecting the means for the strategy given by the child. Along with the emerging role pattern, the need for communication becomes more and more reduced. But after some time the robot might suggest a different order or different placements. This breaks the established role pattern, brings back communication and a level change: the robot now proposes a strategy (intermediate level of autonomy) which might be adopted or further discussed and refined between both. The overall coordinated behavior emerges out of the interaction of actions as well as the shared autonomy. Moreover, during the process the agents discover each others capabilities which additionally impacts on the fusing of their autonomy spaces.

Real world situations abound with such complexities (most of them much harder to describe than in the example), making adaptivity and emergence important for maintaining or establishing shared autonomy in dynamical contexts and between agents whose policies vary in time. Due to this richness, we expect emergent shared autonomy to become an exciting research direction where methods from different areas have to be combined in interesting and challenging ways to create new levels of adaptivity enabling real world agents with sophisticated autonomy to achieve coordination within a very short time span.

Evaluating Shared Autonomy

An obvious key question is: how can we measure the degree of “successfulness” of shared autonomy?

Clearly, the multidimensional perspective on shared autonomy can also provide guidance to this question: once we pick a number of dimensions along which we characterize the fusing of the agents’ autonomy spaces, measuring shared autonomy becomes transformed into a multicriterial optimization problem and we might apply the toolset that is available in this established framework. For instance, we might consider “engineering type” dimensions such as amount of task achievement, required communication, failure rates, throughput, negotiation time, mutual anticipation success and many more. Already more challenging are measures that require human judgement, such as user acceptance, ease of interaction, user satisfaction or fatigue. Even more challenging is to measure effects such as the emergence of new capabilities, or individual influences, such as the dependence on agent-individual skills and biases.

We also would like to point out that benchmarking shared autonomy is in an interesting way connected with robustness of policies to disturbances: by utilizing their autonomy spaces, both agents create a degree of unpredictability for each other, and this can only work when the unpredictability is “shaped” in such a way that it does not hit any “sensitive spots” of the policies of the other agent(s). Therefore, high robustness benchmarks for policies should be indicative of a good capability of fusing one’s autonomy space with others.

But there is also little doubt that the biggest challenge is to implement any of the “conceptual” measures for interaction scenarios of real world agents, and obtaining meaningful results within the restricted number of interaction turns

feasible under such constraints. This may require to relax the aspiration of quantitative benchmarking to evaluating success in tournaments. Fortunately, the situation of shared autonomy is very suggestive of a generic tournament design that measures whether the “whole” (the fusing of autonomy spaces in a particular implementation) is “more than its parts”. Such “team gain” can be measured by a simple tournament between two agents that act in parallel and isolation, versus a team of agents acting under identical conditions, but being allowed to fuse their autonomy spaces. However, we should expect that there exist task-dimension pairings for which an optimal team solution might require to sacrifice autonomy in favor of performance, while for others the optimal solution is characterized by an optimal, intermediate level of autonomy of each agent. The characterization of task-dimension combinations with regard to these outcomes appears to be another exciting research question pertaining to a deeper understanding of shared autonomy, which should be approached from both an empirical and a theoretical side.

Conclusion

The fine grained view on autonomy (Gransche et al. 2014) helps to extend the current (classical) view on shared autonomy which mainly considers shared autonomy as a differentiation in between manual execution, automation and autonomy itself. Originally, this view focused on the single artificial system as such and the system is seen as a tool that on some (low) level relieves the user from details of the actual execution of a behavior while leaving the high level autonomy to the user. However, this is a massive limitation of the capabilities that a team of a user and an autonomous system can achieve. Our multidimensional perspective on shared autonomy extends this to interactive scenarios. This allows all agents to contribute on different levels which—in our opinion—is an important requirement for cooperation and communication: as soon as the cooperative behavior becomes a little more complex it is required that all participants are able to understand on different levels and also have a certain degree of freedom on that level. These different levels of understanding and decision are complementing the types of autonomy as proposed by (Gransche et al. 2014): an autonomous system should not only be able to realize currently carried out actions (the means), but also should have a broader understanding of the current intention and the currently followed strategy, including the partner’s intentions as well as the consequences of action selection on the partner’s state of autonomy.

As a goal for our future research such systems should not be restricted to jump in and take only over a certain means, but should rather be able to propose alternative strategies and offer their complementing abilities on all levels, thus shaping the joint behavior autonomously. We want to apply this in different interactive robotic scenarios, for example tidying a room which relies on manipulation (Li, Haschke, and Ritter 2015), perception (Oliveira et al. 2016; Eitel et al. 2015) and interaction capabilities (Twardon and Ritter 2015; Kopp et al. 2014; Renner, Pfeiffer, and Wachsmuth 2014). This requires differentiated underlying rich representations (Nomikou et al. 2016; Schilling and Narayanan 2013), in

particular for communication like the use of ‘Pragmatic Frames’ (Rohlfing et al. 2016) which humans use in their everyday interactions and which are emergent interaction patterns where pragmatic meaning emerges within the interaction situation. It further involves adaptation and learning (Kuderer, Gulati, and Burgard 2015; Boedecker et al. 2014) on those different levels of representation from few trials in complex scenarios.

Acknowledgments

This research/work was supported by the Cluster of Excellence Cognitive Interaction Technology ‘CITEC’ (EXC 277) at Bielefeld University, which is funded by the German Research Foundation (DFG) and the BrainLinks-BrainTools Cluster of Excellence (grant number EXC 1086).

References

- Boedecker, J.; Springenberg, J. T.; Wlfling, J.; and Riedmiller, M. 2014. Approximate real-time optimal control based on sparse gaussian process models. In *Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*.
- Bradshaw, J. M.; Feltovich, P. J.; Jung, H.; Kulkarni, S.; Taysom, W.; and Uszok, A. 2003. Dimensions of adjustable autonomy and mixed-initiative interaction. In Nickles, M.; Rovatsos, M.; and Wei, G., eds., *Agents and Computational Autonomy*, volume 2969 of *Lecture Notes in Computer Science*, 17–39. Springer.
- Eitel, A.; Springenberg, J. T.; Spinello, L.; Riedmiller, M.; and Burgard, W. 2015. Multimodal deep learning for robust rgb-d object recognition. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Endsley, M., and Kaber, D. 1999. Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics* 42(3):462–492.
- Fong, T.; Thorpe, C.; and Baur, C. 2003. Collaboration, dialogue, Human-Robot interaction. In Jarvis, R. A., and Zelinsky, A., eds., *Robotics Research*, volume 6 of *Springer Tracts in Advanced Robotics*. Springer Berlin Heidelberg, 255–266.
- Goodrich, M. A., and Schultz, A. C. 2007. Human-Robot interaction: A survey. *Foundations and Trends in Human-Computer Interaction* 1(3):203–275.
- Goodrich, M. A.; Jr., D. R. O.; Crandall, J. W.; and Palmer, T. J. 2001. Experiments in adjustable autonomy. In *Proceedings of the IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents*, 1624–1629.
- Goodrich, M.; Crandall, J.; and Barakova, E. 2013. Teleoperation and beyond for assistive humanoid robots. *Reviews of Human Factors and Ergonomics* 9.1 175–226.
- Gransche, B.; Shala, E.; Hubig, C.; Alpsancar, S.; and Harlach, S. 2014. *Wandel von Autonomie und Kontrolle durch neue Mensch-Technik-Interaktionen: Grundsatzfragen autonomieorientierter Mensch-Technik-Verhältnisse*. Fraunhofer Verlag.

- Hertkorn, K. 2016. *Shared Grasping: a Combination of Telepresence and Grasp Planning*. KIT Scientific Publishing.
- Kanda, T., and Ishiguro, H. 2012. *Human-Robot Interaction in Social Robotics*. Taylor & Francis.
- Kopp, S.; van Welbergen, H.; Yaghoubzadeh, R.; and Buschmeier, H. 2014. An architecture for fluid real-time conversational agents: integrating incremental output generation and input processing. *J. Multimodal User Interfaces* 8(1):97–108.
- Kuderer, M.; Gulati, S.; and Burgard, W. 2015. Learning driving styles for autonomous vehicles from demonstration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*.
- Li, Q.; Haschke, R.; and Ritter, H. 2015. A visuo-tactile control framework for manipulation and exploration of unknown objects. In *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th International Conference on*, 610–615. IEEE.
- Nomikou, I.; Schilling, M.; Heller, V.; and Rohlfing, K. 2016. Language – at all times: Action and interaction as contexts for enriching representations. *Interaction Studies* 17(1).
- Oliveira, G.; Valada, A.; Bollen, C.; Burgard, W.; and Brox, T. 2016. Deep learning for human part discovery in images. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Renner, P.; Pfeiffer, T.; and Wachsmuth, S. 2014. Towards a model for anticipating human gestures in human-robot interactions in shared space. *Cognitive Processing* 15(1 Supplement):59–60.
- Rohlfing, K. J.; Wrede, B.; Vollmer, A.-L.; and Oudeyer, P.-Y. 2016. An alternative to mapping a word onto a concept in language acquisition: pragmatic frames. *Frontiers in Psychology* 7(470).
- Schilling, M., and Narayanan, S. 2013. Communicating with executable action representations. In *AAAI Spring Symposium: Designing Intelligent Robots*, volume SS-13-04 of *AAAI Technical Report*. AAAI.
- Sheridan, T. B., and Verplank, W. L. 1978. Human and computer control of undersea teleoperators (Man-Machine systems laboratory report).
- Twardon, L., and Ritter, H. 2015. Interaction skills for a coat-check robot: Identifying and handling the boundary components of clothes. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 3682–3688. IEEE.
- Vernon, D. 2014. *Artificial Cognitive Systems: A Primer*. MIT Press.