An Autonomous Robotic Assistant for Drinking

Sebastian Schröer Ingo Killmann Barbara Frank Martin Völker Lukas Fiederer Tonio Ball Wolfram Burgard

***Abstract*— Stroke and neurodegenerative diseases, among a range of other neurologic disorders, can cause chronic paralysis. Patients suffering from paralysis may remain unable to achieve even basic everyday tasks such as liquid intake. Currently, there is a great interest in developing robotic assistants controlled via brain-machine interfaces (BMIs) to restore the ability to perform such tasks. This paper describes an autonomous robotic assistant for liquid intake. The components of the system include autonomous online detection both of the cup to be grasped and of the mouth of the user. It plans motions of the robot arm under the constraints that the cup stays upright while moving towards the mouth and that the cup stays in direct contact with the user's mouth while the robot tilts it during the drinking phase. To achieve this, our system also includes a technique for online estimation of the location of the user's mouth even under partial occlusions by the cup or robot arm. We tested our system in a real environment and in a shared-control setting using frequency-specific modulations recorded by electroencephalography (EEG) from the brain of the user. Our experiments demonstrate that our BMI-controlled robotic system enables a reliable liquid intake. We believe that our approach can easily be extended to other useful tasks including food intake and object manipulation.**

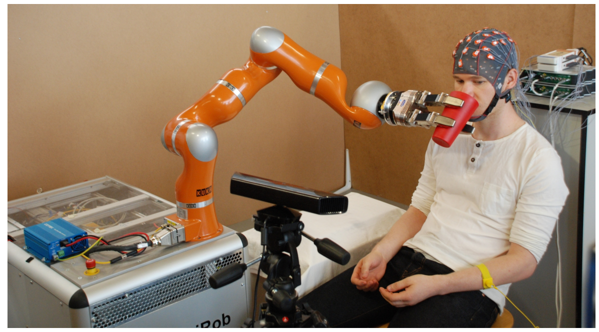

Fig. 1. Our BMI-controlled robot providing a user with a drink. The BMI consists of three components, (i) the EEG recording system, (ii) the RGB-D camera and (iii) the robotic manipulator. The EEG is used to detect go-commands from the user. The RGB-D camera detects the mouth of the user as well as the drinking cup. The robotic manipulator grasps the cup, serves the drink to the user and places the cup back to the table.

I. Introduction

Robotic assistance for disabled users is an application scenario with a great potential. Several previous systems demonstrate that patients may carry out useful tasks without the help of a caregiver using a robotic assistant. Depending on the severity of the user's disability, different modes of interaction and control of the robot can be considered, such as via a joystick, via head- or eye-tracking, or via brain-machine interfaces (BMIs) [7, 4, 3]. While some BMI approaches aim at a high-dimensional motor control of robotic devices, where all aspects of the movement are directly decoded from brain activity, other so-called shared control approaches delegate a part of the control task to the robotic device. In this way, the device might for example autonomously carry out the details of the trajectory control involved to reach a target specified by the user.

However, the requirements to be met by autonomous assistive devices in such scenarios are often demanding. For example, in shared control of liquid intake, the operation of tilting the cup must be well-planned to deliver neither too much nor too little liquid to the user. Furthermore, it is necessary that a liquid-filled cup is brought to the mouth and kept in direct contact with the mouth of the user with high precision. Visual tracking of the mouth for this purpose must also solve the problem that the mouth typically will be occluded during the task execution either by the cup or by the robot arm.

To address these problems, we developed a system integrating all relevant components for enabling reliable task execution in the context of liquid intake using an assistive (mobile) manipulator that performs object detection, motion planning, and motion execution autonomously (see Fig. 1). We used and extended existing technology for object detection and motion planning, in particular to reliably detect the mouth of the user even in case of occlusions. Further we evaluated the compatibility of the robotic assistant with an EEG-based BMI setup using frequency-specific modulation of brain activity translated into control signals derived from dry-electrode EEG recordings [10].

II. Related Work

Recently, several systems to provide assistance to disabled users have been described. Chen et al. [3] developed an assistive assistive mobile manipulation robot in close collaboration with a patient suffering from tetraplegia and different scenarios for assistance were investigated. These included collecting items from cupboards, distributing halloween candy to children, or scratching and shaving the user. Control over these actions was shared between the user and the robot. The user selected high-level control signals via head tracking and clicking a button, for instance, by selecting the target region from an RGB-D video stream. The robot then analyzed the video data, planed the appropriate motion to approach the target and adjusted the contact force

All authors are with the University of Freiburg, 79110 Freiburg, Germany.
This work has been partly supported by the German Research Foundation under grant number EXC 1086, the German Federal Ministry of Education Research grants 16SV5834 NASS, 0316064C BrainCon, and grant BMI-Bot by the Baden-Württemberg Stiftung.

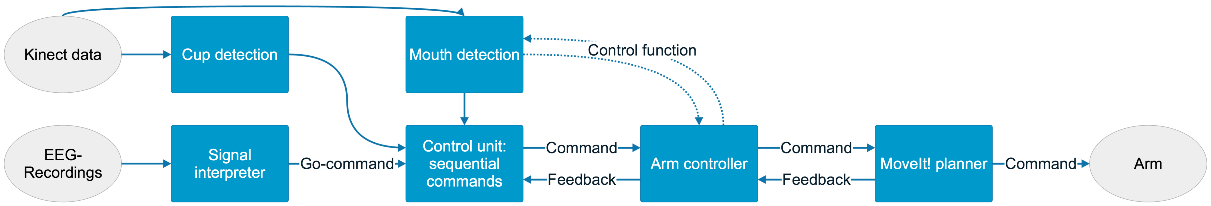

Fig. 2. System Overview

required by the task. Food or liquid intake, however, were not investigated. Devices for feeding elderly and disabled persons have been presented that also rely on interaction with the user via button-click, examples include [20, 18, 17]. They, however, did not carry out tasks autonomously, but entirely relied on the user. Similarly, a vision-based manipulator for liquid intake was described by Higa et al. [6]. They determined the 3-D coordinates of a plastic bottle using a monocular camera and template matching, control required users to manipulate a joystick.

Controlling robotic arms with neural signals was addressed by Hochberg et al. [7]. In this work, patients were implanted with a 96-channel micro-electrode array in the motor cortex, and single neuron activity was extracted to control the prosthetic arm. Collinger et al. [4] present a study that goes into a similar direction. In both cases, training time for the BMI required sessions of several weeks to learn direct motor control with multiple degrees of freedom. In most of the application examples reported in these studies, the user directly controlled the complete trajectory of the arm, which, on the one hand, gives the user control over the exact execution of the task. On the other hand, the performance of these approaches is still limited and below what would be desirable for practical application.

To increase performance, we propose a shared control approach, in which the robot autonomously takes over at least a part of the planning of the trajectory, of detecting the cup and mouth, etc.

Shared autonomy in the context of BMI was previously realized in different application scenarios, for instance, to steer wheelchairs that take care of avoiding obstacles [14, 9] or in tele-presence robots [2]. A concept of high-level BMI control similar to our system was presented by Lampe, et al. [11, 10], without, however, addressing the specific challenges of food or liquid intake.

III. System Overview

Our system consists of several components. The hardware is a KUKA omniRob platform that is equipped with a DLR light-weight arm with seven degrees of freedom. For grasping objects, we attached a Schunk three-finger hand to the robot arm. To observe the entire scene and to detect the cup, the user's mouth, and potential obstacles we use a Kinect RGB-D sensor. Finally, to communicate with the user we employ a BMI that extracts control signals from EEG signals. Fig. 2 illustrates the different components as well as their interaction.

In the following, we describe the parts of the system that are required for the autonomous execution of the task. This includes the detection of the cup and the user's mouth as well as the planning of the motions of the robot to the target locations. Furthermore, we demonstrate the compatibility of this system with a BMI control based on motor-imagery tasks that allows the user to send control signals to the robot.

IV. Object Detection

The goal of our robot is to autonomously deliver a drink to the mouth of the user. Thus, the robot needs to identify the 3D positions of the cup and the mouth of the user in order to plan appropriate motions. As mentioned above, we employ an RGB-D Kinect sensor, which is placed in a way that it is able to observe the required workspace of the robot. As the Kinect sensor measures the positions of the cup and the mouth of the user in its local coordinate frame, we have to determine the transformation between the robot and the camera coordinate system. We calculate this transformation in a calibration step. To this end, we observe the robot end-effector with the camera from four different positions and orientations. We identify the tool-center point in the camera images using an optical marker. In our setting, four observations are sufficient to determine the full 6d transformation, rotation and translation between the two coordinate frames.

A. Cup Detection

In our current system, we consider drinks that are provided in colorful plastic cups, which is taught to the system beforehand and which can efficiently be detected using a color-based classifier. To achieve this, we first transform the RGB-image of the Kinect camera to the HSV format. Second, we dilate and erode it based on the color shade of the used cup to obtain a segmentation of the image. Afterwards, we compute the center of the cup from the pixels belonging to the cup. Finally, we compute the 3D coordinate of the center point of the cup using the known offset between depth and RGB pixels. In the future, we plan to let the user select the corresponding cup from a set of cups that are present in the workspace. This can, for instance, be realized by providing another level of user interaction via the BMI.

B. Mouth Detection

The detection of the 3D position of the mouth of the user is achieved using an extension of the Viola-Jones face-detection algorithm [21, 22]. In our system, we transform the RGB image to a grayscale image and detect the face of the user using a haarcascade classifier provided by *OpenCV* [1]. In order to determine the position of the mouth within the detected face, we rely on the golden ratio of human faces. In practical experiments we found out that this method yields better results than directly tracking the mouth of the user using a specific haarcascade classifier. According to the golden ratio, we assume the position of the mouth to be at 7/9 of the height and at 50% of the width of a human face. However, when the user moves during the execution of the drinking task, the position of the mouth might change as well. This requires reliable online tracking of the mouth in order to be able to react to changes and to adapt the path of the robot accordingly. When the robot approaches the user, however, its end-effector might occlude the mouth which makes tracking more challenging. To deal with such occlusions of the mouth and to accurately determine the 3D position of the user's mouth even under occlusions over longer periods of time, we perform the following approach. We consider three specific points in the user's face, which can typically be observed even when the robot's end-effector approaches the user: the first one, $\mathbf{fh}$, is on the forehead, while the second and the third, $\mathbf{lc}$ and $\mathbf{rc}$, are on the left and right cheek, respectively. A plane P is described by these three points with s and $t \in \mathbb{R}$:

$$P = \mathbf{fh} + s(\mathbf{fh} - \mathbf{rc}) + t(\mathbf{fh} - \mathbf{lc}) \tag{1}$$

As faces of different users of our system might exhibit different anatomical features, the plane constructed in this way will not necessarily coincide with the mouth of the user, but needs to be shifted or rotated accordingly (see Fig. 4). This can be described by a user-specific parameter u:

$$P(u) = u\mathbf{fh} + s(\mathbf{fh} - \mathbf{rc}) + t(\mathbf{fh} - \mathbf{lc}), \tag{2}$$

which minimizes the difference between the actually measured z-value of the mouth z_{meas} (obtained from one calibration image in which the mouth is not occluded) and the expected z-value computed from the plane equation $z_{\text{exp}}(u)$:

$$\min_u D(u) = z_{\text{meas}} - z_{\text{exp}}(u) \tag{3}$$

Since the $x-$ and y-coordinates of the mouth (x_m, y_m) result from the golden ratio of the face, we obtain the z-value $z_{\text{exp}}(u)$ of the mouth from the intersection of the line from the camera center to the mouth and the plane. The line from the camera center to the mouth is given by

$$L(u) = w\mathbf{n}(u), \tag{4}$$

with $w \in \mathbb{R}$ and $\mathbf{n}(u) = (x_m, y_m, z_{\text{exp}}(u))$ and finally

$$z_{\text{exp}}(u) = P(u)L(u). \tag{5}$$

This calibration procedure is easy to implement and only requires one RGB-D image of an un-occluded mouth to determine the user-specific parameter u. Given the plane and the (x, y)-position of the mouth, we are able to accurately determine the 3D position of the user's mouth even under occlusions.

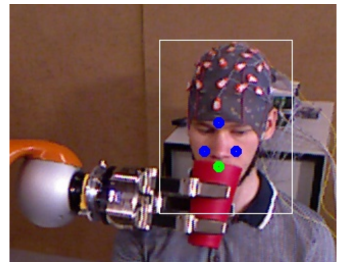

Fig. 3. Typical situation in which the mouth of the user is occluded by the cup. Furthermore, the EEG cap occludes parts of the user's face. Here, only a fraction of around 50% of the face is visible by the RGB-D-camera. With our approach, by considering distinct points in the face (blue points), we are still able to determine the depth value of the mouth (green point) with high accuracy.

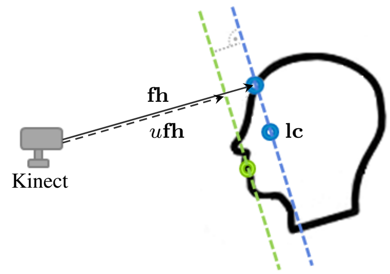

Fig. 4. Identifying the depth of the user's mouth: a plane described by three distinct points $\mathbf{fh}$, $\mathbf{lc}$ and $\mathbf{rc}$ is shifted according to a user specific parameter u to coincide with the mouth. By computing the intersection of the line through the mouth position in the camera frame with the green plane we obtain the depth value for the mouth even under occlusions.

V. Motion Planning

Planning motions for a 7-DOF arm is a complex problem. We use rapidly exploring random trees (RRTs) [8] for this task. The key idea of RRTs is to iteratively probe and explore the configuration space given a starting and goal configuration of the robot. RRTs are designed to handle nonholonomic constraints and are well suited for systems with many degrees of freedom. We use the implementation of *MoveIt!* [19], which provides an integration with ROS and the robot.

Delivering a drink to the user requires the robot to consider different constraints: when carrying the drink, the cup should stay upright; when the user wants to drink, the rim of the cup should be in contact with the user's mouth while the cup is turned. Therefore, we divide the task of providing the drink into different sub-tasks: after grasping the cup, the robot first moves to a position 20 cm in front of the user's mouth in order to avoid collisions with the user. Then it approaches the mouth and turns the cup. To enforce constraints such as an orientation of the end-effector that keeps the cup upright, we sample waypoints along linear trajectories between starting and goal configuration. Waypoints are represented by the positions and orientations of the end-effector in the work space and are used as an input for the RRT planner. The planner connects the waypoints along the trajectory with valid movements and computes the corresponding robot configurations using inverse kinematics. With a distance of

Fig. 5. Mouth detection under occlusions: We artificially introduced black rectangles around the mouth. The three leftmost images illustrate the positions of the detected mouths under occlusions—our approach is able to deal with closed eyes as well as glasses. The image on the right, in contrast, shows a situation in which no face, and hence, no mouth could be detected.

subject	# frames	recall rate (%)	MAE (cm)
P1	430	99.8	0.26±0.21
P2	381	86.6	0.94±0.94
P3	434	97.7	0.50±0.39
P4	448	92.9	0.35±0.28
P5	444	96.4	0.35±0.30
P6	384	98.2	0.70±0.61

TABLE I
EVALUATION OF OUR MOUTH DETECTION APPROACH UNDER OCCLUSIONS FOR DIFFERENT PEOPLE AND DIFFERENT VIEWPOINTS.

2 cm between individual waypoints, we ensure that the end-effector moves along the desired trajectory.

When turning the cup, we have to ensure that the rim of the cup is in contact with the user. Since the rim of the cup has a constant offset from the end-effector, the corresponding trajectory of the end-effector is lying on a circular orbit around this turning point. We obtain waypoints on this orbit by triangulating the described offset of the turning point. Densely sampled waypoints along this trajectory ensure a smooth tilting movement.

VI. Brain-Machine Interface

Our system is designed for the specific needs of paralyzed patients who are not able to control joysticks or other input devices with their muscles, but are still able to drink. As voice commands cannot be used while drinking, we rely on decoding brain activity as a means of communication with the robot. To this end, EEG recordings are used successfully to detect changes in different EEG signal components that occur when a user imagines, for instance, different kinds of movements. These responses can be trained to increase over time and lead to more reliable detection results.

A. User interaction

In a minimal version of our application, in principle it would be sufficient for the user to send one start command. Subsequently, the robot could autonomously execute the entire task, which consists of reaching for and grasping the cup, bringing it to the user's mouth, tilting the cup to let the user drink, and returning the cup to the table. In order to provide the user with a larger degree of autonomy, however, we split the task into several individually initialized subtasks:

1) Find the cup and grasp it
2) Bring the cup to a position in front of the user
3) Move the cup to the mouth of the user
4) Tilt the cup
5) Bring the cup back into an upright position
6) Move away from the users face
7) Return the cup to the table

The execution of each step requires an explicit go-signal from the user. Furthermore, our current system approaches a fixed tilting angle of the cup which allows the user to incorporate a fixed amount of liquid only. This could be easily extended to a more flexible approach by additionally decoding a cancel signal from the user. In such a scenario, a finite state machine could keep track of the state of the robot. Subsequent go commands could be used to tilt the cup in small incremental angles. Finally, the user could decide when to stop the liquid intake by canceling the task.

B. EEG-to-Robot Pipeline

To send a command to the robot, the user imagines finger tapping movements of the right hand. These imaginations lead to a frequency power decrease in brain waves in the alpha and beta frequency band that can be measured with an EEG recording cap. In our experimental setup, we used a dry-electrode EEG cap, for ease of use, quick setup time, and comfort of the user. The EEG data are acquired from 32 g.SAHARA dry active electrodes, sampled at 512 Hz using two amplifiers (GUGER TECHNOLOGIES, Graz, Austria). Electrodes are located according to the 10-20 system [13]. We used BCI2000 [16], a free and thus widely distributed BMI processing framework to interface amplifiers and signal processing. Signal processing, feature extraction and decoding was performed in real-time in MATLAB using an artificial neural network. Details about feature extraction from the EEG signal and decoding methods can be found in [10]. No direct communication method between MATLAB and robotic system was available. Therefore we set up a uni-directional User Datagram Protocol (UDP) network interface to connect MATLAB and the robot. We implemented the communication over the campus' wireless network. Whenever our decoding system detects the imagination of the right hand finger tapping consecutively for 3 s, a 'true' bit is sent to the robot. This threshold of 3 s increases confidence in the decoded command and was determined experimentally to avoid false positives. Of course, this leads to a trade-off between safety and satisfaction of a user who will experience frustration if intended commands are not carried out.

VII. Experimental Evaluation

In several experiments, we evaluated our approach to detect the mouth of the user under occlusions. Furthermore, we demonstrate in a series of experimental runs, how our robotic system successfully and reliably accomplished its task, that is, enabling liquid intake.

A. Mouth Detection

An important component of our system is the ability to identify the mouth of the user in order to precisely move the cup to the desired position. This is required to work even under occlusions, which are inevitable when the robotic arm with the cup approaches the user and the mouth. We evaluated our mouth detection algorithm on image sequences of six different subjects one of which was wearing glasses (P6, see Fig. 3). In total, we collected for each user around two minutes of RGB-D recordings of their faces. In order to reduce similarity between subsequent test images, the frame rate was reduced to 4 Hz. To demonstrate the robustness of the tracking system to movements of the user, the users were asked to move their head arbitrarily in the x, y, and z direction. Furthermore, we asked them to perform moderate rotations of the head. Under optimal circumstances, when the entire face was visible, the mouth could be detected reliably with a detection rate of close to 100 %. We used these detections as ground truth to evaluate how well the system is able to deal with occlusions. We then artificially introduced occlusions to the test images by masking a rectangular region around the detected mouth. We evaluated, how well the mouth could be detected and how accurately the position of the mouth could be determined with respect to the ground truth data. We chose the parameters for the face detection approach such that we did not receive any false positives. Thus, we only report the recall rate, which is the fraction of true positives versus the total number of instances. As our mouth detection algorithm is based on the detection of human faces, it will fail whenever important features of the face are not visible. If, for instance, the user turns his face by about 90 degrees to the left or right, one of the eyes will not be visible and the face cannot be tracked. Therefore, our mouth detection algorithm fails in such cases and leads to negative detection results. This is particularly evident with participant P2 who performed rather large and extreme movements, which challenged the system and lead to a lower recall rate and a larger error in the detected mouth. We evaluated the accuracy of the position tracking in terms of the mean average error (MAE). Table I summarizes the results. Fig. 5 illustrates examples for successful and failed detections. As is illustrated here, the system is able to track the mouth if the user closes his eyes or wears glasses.

B. System Evaluation

We furthermore performed an experiment with the entire system, that is with communication of the user with the robot via EEG, detection of objects and mouth, grasping the cup and bringing the cup to the mouth, letting the user drink, and placing the cup on the table again after the user is finished. Our test user was experienced in BMI control but not engaged in robotics. We executed the experiment in three sessions over two consecutive days. The first two sessions were used to integrate and test the robotic system and the EEG decoding. In the third session, we evaluated the performance of the system in 13 consecutive runs.

Before the evaluation session, we re-calibrated the EEG decoding with 10 min of training data. In contrast to an experiment in a dedicated EEG laboratory, the environmental circumstances for the BMI were challenging. Firstly, the robotic arm can lead to EEG recording artifacts and hence to a decreased classification accuracy. Secondly, the unaccustomed presence of the robot in front of the user's face can cause anxiety, which may interfere with properly performing the movement imagination task. Over the course of two calibration sessions, however, we observed that our test user developed trust into the system and got accustomed to the presence of the robot. Furthermore, we observed some false positive detections in the EEG measurements while the robotic arm was moving in front of the user's face. This could possibly be explained by the activity of the mirror neuron system [5, 12, 15] triggered by the observation of a moving robot arm, which could lead to similar effects as movement imagination of the right hand. This possibility could be examined in further studies.

During the third session, however, all 13 runs were successful, that is the user was able to send the commands to the robot, the robot delivered the drink to the correct position, the user was able to drink, and finally the robot put the cup back to the table and returned to its initial position. During the first 10 trials, the required time to complete the task was on average 3 minutes per trial. To speed up the task execution, in the last three trials, we increased the speed of the decoding procedure. Per default, 5 consecutive control signals, that is finger tapping movements needed to be detected before a control signal was sent to the robot. In the last three trials, we reduced this to three consecutive detections. Furthermore, we increased the arm velocity from 10 % to 20 % of the maximum arm speed. The robot also completed these three trials successfully with an average runtime of two minutes per trial. Fig. 6 shows some snapshots of one trial, a video of an experimental run is available online[1].

VIII. Conclusions and Future Work

In this paper we presented a prototypical robotic system for paralyzed patients to drink without having to rely on the assistance of caregivers. The corresponding increase in daily independence is important for patients and might greatly improve the quality of their lives. We realized a robotic system that is able to autonomously detect the 3D positions of objects and users and furthermore is able to plan its movements to autonomously achieve the task. The setup of the dry-electrode EEG used to interface the users brain and the robot was fast, easy and provided stable signals over the whole duration of the experiments. Dry-electrode EEG thus represents a viable alternative to the more widespread conventional, wet electrode systems.

Future development will include improving the speed of the robotic system and adding more user safety, for instance, by using force-feedback sensing to detect contact between user and cup. Up to now we do not detect the entire body of the user. Furthermore we avoid collisions of the robotic arm and the user by including waypoints that the arm has to approach and therefore we demand a very stringent room of possible movement trajectories. As a consequence one part of our future work will focus on the detection of the entire body

[1] `http://ais.informatik.uni-freiburg.de/projects/neurobots/`

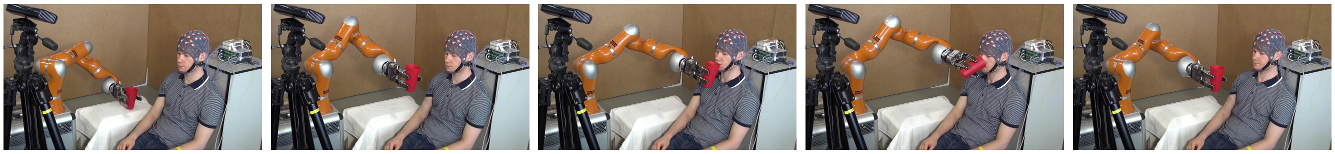

Fig. 6. Successful execution of the drinking task.

of the user and the implementation of a reliable collision avoidance leading in addition to further improvements of the speed of our system due to less stringent constraints on the motion planning algorithm.

As a longterm goal, we envision having the robot, instead of the experimenter, pour the beverage into the cup and extending our approach to self-feeding. The current advances in implantable, wireless recording devices also make it possible to envision transferring our approach to control signals recorded by means of such devices. Such a transfer would solve noise issues and can be expected to substantially improve the information transfer rate from brain to robot.

IX. Acknowledgements

We would like to thank Daniel and Andreas Kuhner for helping us with implementing the autonomous movements of the arm.

References

[1] G. Bradski. *Dr. Dobbs Journal of Software Tools*, 2000.

[2] T. Carlson, L. Tonin, S. Perdikis, R. Leeb, and J. d. R. Millán. A Hybrid BCI for Enhanced Control of a Telepresence Robot. In *Proceedings of the 35th Annual International Conference of the Engineering in Medicine and Biology Society*, 2013.

[3] T. Chen, M. Ciocarlie, S. Cousins, P. M. Grice, K. Hawkins, K. Hsiao, C. Kemp, C.-H. King, D. Lazewatsky, A. E. Leeper, H. Nguyen, A. Paepcke, C. Pantofaru, W. Smart, and L. Takayama. Robots for humanity: A case study in assistive mobile manipulation. *IEEE Robotics & Automation Magazine, Special issue on Assistive Robotics*, 20, 2013.

[4] J. L. Collinger, B. Wodlinger, J. E. Downey, W. Wang, E. C. Tyler-Kabara, D. J. Weber, A. J. C. McMorland, M. Velliste, M. L. Boninger, and A. B. Schwartz. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*, 381(9866):557–564, 2013.

[5] V. Gazzola, G. Rizzolatti, B. Wicker, and C. Keysers. The anthropomorphic brain: The mirror neuron system responds to human and robotic actions. *NeuroImage*, 35(4):1674–1684, 2007.

[6] H. Higa, K. Kurisu, and H. Uehara. A vision-based assistive robotic arm for people with severe disabilities. *Transactions on Machine Learning and Artificial Intelligence*, 2(4), 2014.

[7] L. R. Hochberg, D. Bacher, B. Jarosiewicz, N. Y. Masse, J. D. Simeral, J. Vogel, S. Haddadin, J. Liu, S. S. Cash, P. van der Smagt, and J. P. Donoghue. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–275, 2012.

[8] J. J. Kuffner and S. M. LaValle. RRT-connect: An efficient approach to single-query path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2000.

[9] Robert L., S. Perdikis, L. Tonin, A. Biasiucci, M. Tavella, M. Creatura, A. Molina, A. Al-Khodairy, T. Carlson, and J. d. R. Millán. Transferring brain-computer interfaces beyond the laboratory: Successful application control for motor-disabled users. *Artificial Intelligence in Medicine*, 59(2):121–132, 2013.

[10] T. Lampe, L. D .J. Fiederer, M. Voelker, A. Knorr, M. Riedmiller, and T. Ball. A brain-computer interface for high-level remote control of an autonomous, reinforcement-learning-based robotic system for reaching and grasping. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*, 2014.

[11] T. Lampe and M. Riedmiller. Acquiring visual servoing reaching and grasping skills using neural reinforcement learning. In *IEEE International Joint Conference on Neural Networks (IJCNN 2013)*, 2013.

[12] L. M. Oberman, J. P. McCleery, V. S. Ramachandran, and J. A. Jaime. EEG evidence for mirror neuron activity during the observation of human and robot actions: Toward an analysis of the human qualities of interactive robots. *Neurocomputing*, 70(13–15):2194–2203, 2007.

[13] R. Oostenveld and P. Praamstra. The five percent electrode system for high-resolution EEG and ERP measurements. *Clinical neurophysiology*, 112(4):713–719, 2001.

[14] J. Philips, J. d. R. Millán, G. Vanacker, E. Lew, F. Galán, P. W. Ferrez, H. Van Brussel, and M. Nuttin. Adaptive shared control of a brain-actuated simulated wheelchair. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics*, 2007.

[15] G. Rizzolatti and L. Fogassi. The mirror mechanism: recent findings and perspectives. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1644), 2014.

[16] G. Schalk, D.J. McFarland, T. Hinterberger, N. Birbaumer, and J.R. Wolpaw. BCI2000: a general-purpose brain-computer interface (bci) system. *IEEE Transactions on Biomedical Engineering*, 51(6):1034–1043, 2004.

[17] W.-K. Song and J. Kim. Novel assistive robot for self-feeding. *InTech: Robotic Systems - Applications, Control and Programming*, 941(7), 2012.

[18] R. Soyama, S. Ishii, and A. Fukase. The development of meal-assistance robot "myspoon". In *Proceedings of the IEEE International Conference on Rehabilitation Robotics*, 2003.

[19] I. Sucan and S. Chitta. Moveit! [Online] Available: `http://moveit.ros.org`.

[20] M. Topping. An overview of the development of Handy 1, a rehabilitation robot to assist the severely disabled. *Journal of Intelligent and Robotic Systems*, 34:253263, 2002.

[21] P. Viola and M. J. Jones. Rapid object detection using a boosted cascade of simple features. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.

[22] P. Viola and M. J. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57(2):137–154, 2004.