Grasping with your brain: a brain-computer interface for fast grasp selection

Robert Ying and Jonathan Weisz and Peter K. Allen

Abstract Brain-Computer Interfaces are promising technologies that can improve Human-Robot Interaction, especially for disabled and impaired individuals. Non-invasive BCI's, which are very desirable from a medical and therapeutic perspective, are only able to deliver noisy, low-bandwidth signals, making their use in complex tasks difficult. To this end, we present a shared control online grasp planning framework using an advanced EEG-based interface. Unlike commonly used paradigms, the EEG interface we incorporate allows online generation of a flexible number of options. This online planning framework allows the user to direct the planner towards grasps that reflect their intent for using the grasped object by successively selecting grasps that approach the desired approach direction of the hand. The planner divides the grasping task into phases, and generates images that reflect the choices that the planner can make at each phase. The EEG interface is used to recognize the user's preference among a set of options presented by the planner. The EEG signal classifier is fast and simple to train, and the system as a whole requires almost no learning on the part of the subject. Three subjects were able to successfully use the system to grasp and pick up a number of objects in a cluttered scene.

1 Introduction

People with restricted mobility currently require significant infrastructural support in order to perform activities of daily living (ADL), including things like manipulating objects, opening doors, and other basic actions that able-bodied people often take for granted. With the current state-of-the-art robotic arms, hands, and perception systems, it is clear that robotic grasping systems could help reduce the depen-

Robert Ying · Jonathan Weisz · Peter K. Allen
Department of Computer Science, Columbia University,
500 W. 120th Street, M.C. 0401 New York, New York 10027
e-mail: robert.ying@columbia.edu,\{jweisz,allen\}@cs.columbia.edu

dency severely disabled individuals have on live-in caretakers, and provide them with the ability to actively interact with their environment.

However, the robotic grasping systems which show the greatest promise in performing ADL tend to be highly complex, involving high degree of freedom manipulators and precise control to achieve their objectives. It is therefore important to present users with a high-level interface to these grasping systems that can operate with relatively little training.

In previous work [22, 23, 21], we have presented a shared-control online grasp planner that collaboratively determines feasible grasps under the active direction of a user through a low-bandwidth interface. We have demonstrated the efficacy of this system using a variety of facial EMG-based devices in moderately cluttered scenes. However, this interface depends on the ability of the user to trigger relevant facial muscles repeatably and reliably.

In this work, we extend this system to an EEG-based system, which has a number of advantages. Firstly, the neurological phenomena used in the system is a subconscious reaction to visual stimuli, and therefore needs very little relevant user expertise to operate. Secondly, the planner can take advantage of visual ambiguity between functionally similar grasps to achieve fast convergence in the shared-control paradigm. The user acts as a filter for the planner, directing it to a desired approach direction and filtering proposed candidates until a reasonable one is found. Three users were able to use this system with minimal training to pick up a variety of objects in a semi-cluttered scene.

2 Prior Work

Brain-Computer Interface (BCI) control over prosthetic and assistive manipulators has been the subject of a great deal of research, through many different strategies and input modalities. Recently there has been a resurgence of interest in this field. One widely cited recent advance was reported by Vogel et al. [19], who showed that a subject with a BrainGate cortically-implanted electrode can use a robotic manipulator to retrieve a drink container by controlling the end-effector location and the opening and closing of the hand. However. this approach requires an invasive device capable of recording a large number of high quality signals.

Noninvasive EEG systems have been demonstrated in a number of simpler tasks. In [15], surface electrode signals related to eye gaze direction are used to control 2D arm position and EEG signals are used to detect eye blinks to control gripper closing. In [11], hand opening/closing and elbow flexion/extension are controlled by EEG signals.

The majority of previous work using EEG control concentrates on trajectory control. However, it has been shown that users find BCI control easier using even higher-level, goal-oriented paradigms [16]. We have begun to see work that attempts to exploit higher-level abstractions to allow users to perform more complex tasks with robotic arms. In [4], EEG signals were used to select targets for pick and place operations for a small humanoid robot. In [20], the authors used EEG signals to control pick and place operations of a 4-DOF Staubli robot. Bryan et al.[5] presented preliminary work in extending this approach to a grasping pipeline on the PR2 robot. In that work, a 3D perception pipeline is used to find and identify target objects for grasping and EEG signals are used to choose between them.

Recently, some authors including [13, 12] have explored shared control paradigms which integrate computer vision, intra-cortical EEG recording, and online planning to perform reaching, grasping, and manipulation tasks. These works are promising, but rely on the higher fidelity control available from implanted devices. In [6], the planner presented in our work, which is focused on acquiring an appropriate grasp of the object with arbitrarily complex hands, was integrated with a similar system. In this work, we introduce an interface to the user which allows them to control higher level, more abstract goals with lower throughput devices, which could be made complimentary to these other shared controlled paradigms.

3 Methods

3.1 Overview

We present here a prototype of an assistive grasping system which integrates a BCI driven user interface with a perception pipeline, a lightweight mountable manipulator, in this case the 6-DOF Mico arm with a two-finger underactuated gripper [2], and an online grasp planning system to allow a user to grasp an object in moderately cluttered scenes. It decomposes the grasping task into a multi-step pipeline where each step generates a visual representation of the options the user can take. Some options which cannot be visually represented, such as returning to a previous state, are presented as white text on a black background. At each stage, the online planning system derives a set of reasonable possible actions and presents them to the user, reducing the complex task of grasping an object in cluttered scenes to a series of decision points that can be navigated with a low throughput, noisy input such as an EEG headcap. Fig. 1 shows a healthy subject in our validation study using the system to grasp a bottle of laundry detergent in a typical scene.

3.2 Grasp Planning

This system uses the Online Eigengrasp Planner introduced by Ciocarlie et al. in [7]. This planner uses simulated annealing to generate grasp candidates by projecting desired contacts points on to the target object to find grasps likely to result in a force closed grasp. In order to make this task computationally tractable, a reduced subspace of the hand's full configuration is sampled. In the case of the a simple gripper such as that on the Mico, this may not be necessary, but the use of this planner makes the computational cost of using a more complex hand nearly the same as this simpler hand. Candidate grasps in near contact positions are refined to

Fig. 1: The subject guiding the system through the active refinement phase. On the left side is the robotic manipulator and three containers in the grasping scene. On the right is a subject using the system through the B-Alert EEG cap, which is relatively unobtrusive and can be worn for long periods of time. The options for the active refinement stage are presented on the monitor in front of the subject in a grid to allow the subject to pick the one they intend to select during the serial presentation. In this example, at least one of the grasps found in the database for the object was reachable, and is highlighted in blue in the upper left corner of the grid. The user may choose to execute the highlighted grasp, or to re-seed the planner with one of the other nine grasps and then re-enter the active refinement phase with a new highlighted grasp.

completed grasps by kinematic simulation of closing the hand at a predefined set of joint velocities.

The resulting contacts are ranked by the maximum wrench perturbation force they are capable of resisting, as described in [8], and the closeness of the alignment between the hand and the object's surface. If the quality metric is above 0.2 and all of the dot products of the normal direction of the hand and object is above 0.8 for all of the contact points, the grasping pose is tested for reachability using the PRM planner of *MoveIt!*[18]. When the scene is cluttered, the motion planner for the reaching motion is slow and likely to fail. In order to make this problem more computationally tractable, we cache previous solutions as grasps are planned. Whenever a previous solution ends near the new candidate grasp pose, we plan from its end point to the new grasp candidate. Since the nature of our grasp planner produces many nearby solutions, this makes the reachability filter significantly faster and more robust. Grasps are ranked first by reachability, then by the grasp quality, and finally by the maximal surface misalignment.

The neighbor generating function of the simulated annealing planner is biased towards a configuration demonstrated by the user. By controlling this seed configuration, the user controls the resulting set of candidates that will presented to them. This allows the user to find a grasp for a particular purpose by iteratively picking the grasp whose pose is nearest to the grasp that they are looking for.

3.3 One-of-many selection

The EEG interface presented in this paper is based on an "interest" detector which can be used to provide a one-of-many selection between various options. This "interest" signal paradigm is based on the work in [14]. The options are presented as a stream of images, and the subject is primed to look for particular images that suit some criterion. This paradigm is known as *Rapid Serial Visual Presentation* (RSVP). Spikes in EEG activity which correlate with "interest" are connected with the image that was presented at the time the EEG activity was evoked, which is then used to derive the user's desired input.

Previous work with this paradigm has asked the subject to look for objects of a particular category. In our system, the images represent actions that are suggested by the grasp planner, which the subject may not have had previous experience with. In this case, the subject must be given time to analyze the options and primed to find the features which make their desired option visually distinct from similar options. In Fig. 2, we illustrate the summary pane containing a grid of all of the options, which are then shuffled and presented to the user. In Fig. 1, you can see the subject reviewing the options in a summary pane before the serial presentation of them begins.

One major advantage of this paradigm is that it generalizes a single interaction across all phases of the grasp planning pipeline. The system only needs to be trained to recognize the "interest" signal for each subject. Afterwards, the subject's interaction with each phase is the same, and the system does not require phase-specific training.

3.3.1 EEG input

Our current implementation uses a B-Alert X10 EEG system from Advanced Brain Monitoring (Carlsbad, CA), which provides 9 electrodes positioned according to the 10-20 system and a pair of reference channels. The EEG data is acquired at 256 Hz, with 60 Hz notch and 0.5 Hz high-pass filters applied before any additional processing. The EEG interest metric is based on that described in [14, 17, 10], with some additional normalization and post-processing.

More information on this system can be found online at the manufacturer's website [1]. As can be seen in Fig. 1, the cap and device are relatively minimalistic, and can be comfortably worn for an hour at a time without requiring rewetting or

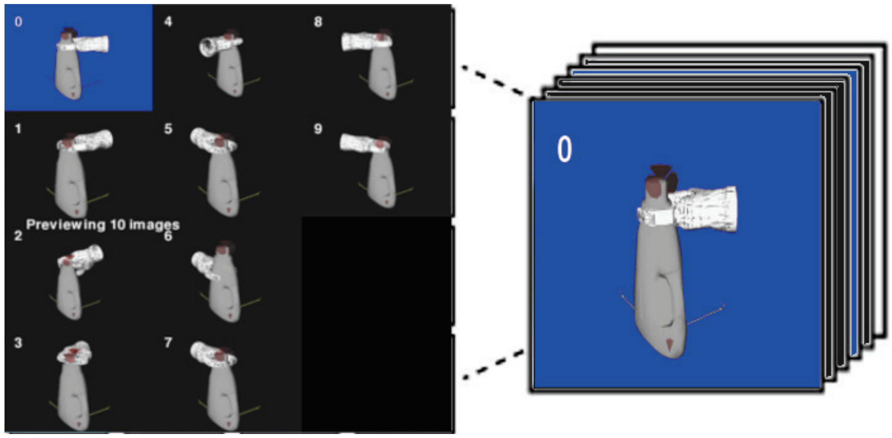

Fig. 2: The grasp planning system compiles a set of images representing potential actions, for example a set of grasps as seen in this image. The image options are tiled together to form the summary pane seen on the left, which lets the user pick out the one that reflects their desire. The images are then shuffled, with repetitions, into a stream that is serially presented to the user as described in section 3.3.3.

reseating of the electrodes. With the advent of the OpenBCI project [3] and similar efforts towards low cost, open hardware EEG devices, a low cost solution with similar capabilities may be on the horizon.

3.3.2 EEG interest metric

The EEG interest metric is based on the one used in [14, 17, 10]. In essence, it assumes that the P300 signal resulting from a particular image varies with a resolution of 100ms. For each block, it examines the time period from 100ms to 1200ms after the input stimulus as separate 100ms blocks, combined in a linear model:

$$y_{sn} = \sum_i w_i x_{in} \qquad y = \sum_n v_n y_{sn} \tag{1}$$

where each x_{in} is the reading at a specific electrode i at some time period n, y_{sn} is the weighted total score over a single 100 ms block and y is the combined score for the 1100ms time period following the stimulus. The weights w_i are learned from the training data so as to maximize the difference between target and non-target images in each time block using Fisher linear discriminant analysis [9]. Then, the weights v_n are determined by applying logistic regression on the training data.

In training, we additionally compute summary statistics for both target and non-target images, which are used later to normalize the individual readings per trial.

3.3.3 Option generation

To generate the RSVP sequence, the system randomly selects each option to appear between three and seven times. The sequence is then randomly shuffled, with the constraint that the same option does not appear in two consecutive image presentations. This method has, in experimental data, been sufficient to trigger the "oddball" response that is necessary for the P300 signal.

If there are less than five options, the system will automatically fill in distractor image options to make this constraint more feasible. The images are dependent on the phase and attempt to minimize the visual difference between the distractor and the original, so as to avoid unintentionally triggering the P300 signal. For example, in the object selection state, the distractor options are of the scene with no objects selected; whereas in the active refinement state they are images of the object with no visible grasp.

More formally, the grasp planner generates a set of options $Q = T \cup G$, where T is a set of strings representing textual options (e.g. "Rerun object recognition"), and G is a set of potential grasp or object images. If $|Q| < 5$, the selection system then adds $5 - |Q|$ distractor images d to result in Q'.

From Q', we generate the sequence of images I as follows:

$$I = \text{shuffle}(I_{q_1}1, I_{q_1}2, \ldots, I_{q_{1c_1}}, I_{q_2}1, \ldots, I_{q_{kc_k}}) \tag{2}$$

where $k = |Q'|$, $q_i \in Q'$ and $c_j \sim U(3,7)\ \forall\ j \in [1,k]$.

The images are each presented at 4 Hz, and preliminary EEG scores e_i are assigned. We then aggregate each of the $n = \sum_{j=1}^{k} c_j$ images by their option, and determine whether or not the user has made a selection.

To test if the user has consciously selected any of the images, we sort the images by their EEG scores, and then split it into a group of size x and $n - x$. We vary x so as to maximize the change in the average measured EEG score:

$$x^* = \arg\max_{x \in [1,n]} \left(\frac{1}{n} \sum_{i=1}^{n} e_i - \frac{1}{n-x} \sum_{i=x+1}^{n} e_i \right) \tag{3}$$

If $x^* > \max(0.2n, 7)$, we determine that the user had not made a choice. In practice, this is a highly reliable means of checking whether the user was paying attention and attempting to make a selection.

If $x^* \leq \max(0.2n, 7)$, we compute a smoothed similarity score using the top x^* positions.

3.3.4 Option scoring

The options are scored using a smoothed similarity metric, represented as a symmetric matrix $S \in R_{k \times k}$, computed such that $S_{ii} = 1$ and $S_{ij} = S_{ji} \in [-1.0, 1.0]$.

We can then construct the weighted score vector W as

$$W_q = \sum_{i=0}^{x^*} S_{q_i q} \tag{4}$$

where q_i is the option corresponding to I_i, and return

$$q^* = \arg\max_{q \in Q'} W_q \tag{5}$$

This scoring method introduces a bias towards groups of similar options, and in essence allows a near-miss selection to nonetheless help select the desired option. From our experiments, this is particularly helpful with subjects who are less experienced with the system, as they often make minor mistakes during the selection process.

We note that this is equivalent to a simple voting scheme if $S = I_{k \times k}$, i.e. the identity matrix of size k. Thus, for options where there is no obvious similarity metric, such as textual or distractor options, we use the corresponding rows and columns from the identity as default.

3.4 Grasping Pipeline

There are four states that the user progresses through when attempting to formulate a grasp, *Object Selection*, *Grasp Selection*, *Grasp Refinement*, and *Confirmation*. The pathway is illustrated in Fig. 3.

3.4.1 Object selection state

In this stage, an object recognition system is used to retrieve models from a database that fit the scene. An image representing selection of each object is generated as shown in the "summary pane" in Fig. 4a, with the target object highlighted green in each potential selection. Between the various images only the highlighted object changes. An additional state is presented that allows the user to run the recognition system again. If fewer than eight objects are detected, additional distractor images of the scene with no highlighted object are generated to act as distractor images which help establish the background level of EEG activity. The user is instructed to just look for the object they want to grasp as the image stream is shown. In this state, the similarity matrix is the identity matrix over the viable options, as the objects are highly dissimilar.

3.4.2 Grasp selection and refinement state

Once the object is selected, the system moves into the grasp selection state. The user's interaction with the grasp selection state and refinement states are very sim-

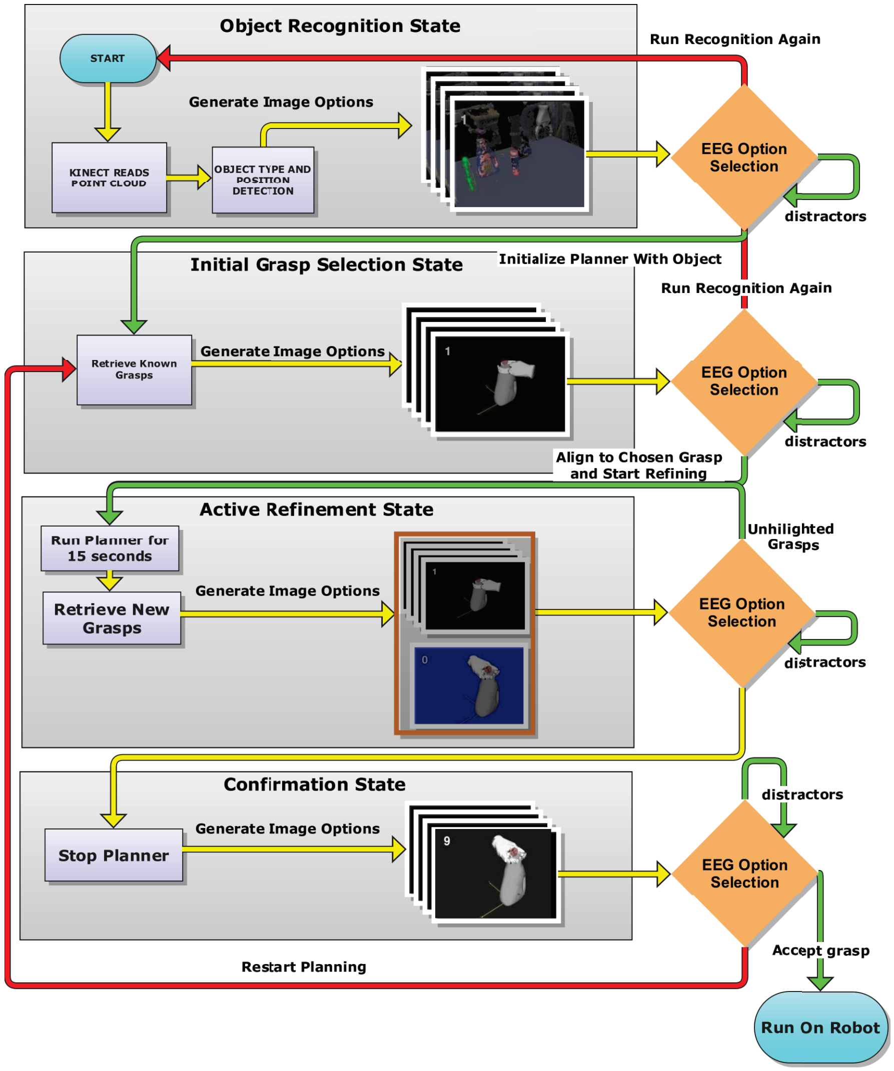

Fig. 3: A diagram outlining the EEG RSVP-driven grasping pipeline. In each phase, a series of images is generated representing the available options, as described in section 3.4. A summary pane of the image options generated at each phase is presented in more detail in Fig. 4a-4d.

ilar. Examples of the “summary pane” for these phases are shown in Fig. 4b and Fig. 4c.

In the grasp selection state, the set of preplanned grasps is retrieved and placed in an arbitrary order. Each of the grasps is visually distinct, and supplies the planner with an approach angle to start with. One additional text option is presented, which sends the user back to the object selection stage. When any grasp is detected as a valid selection, the system enters the grasp refinement state, setting the seed grasp g_s to the one just selected. If fewer than eight grasps are available, images of the object without a visible grasp are used as distractors.

In the grasp refinement state, the online planner begins populating the grasp list with more grasps that are similar to g_s. After allowing the planner to run for fifteen seconds, the available grasps are presented to the user. In most cases, this will be a list of at least ten potential grasps. As each grasp is generated, it is checked for reachability– while even non-reachable grasps are sent to the user, the grasp refinement state cannot be exited until a reachable grasp has been selected. The highest quality reachable grasp g_p is highlighted in blue in the user interface (Fig. 4c), so that the user has feedback as to what the planner is deciding between.

Once the user selects a grasp, the planner updates g_s to the new grasp’s hand state and approach vector. If the updated $g_s = g_p$, then the user exits grasp refinement and enters the confirmation state.

In this phase, we also take advantage of visual ambiguity. We compute the similarity matrix S as follows:

$$S_{ij} = \langle \hat{u}_i, \hat{u}_j \rangle = \cos\theta_{ij} \tag{6}$$

where $\hat{u}_i$ and $\hat{u}_j$ are the approach directions for the grasps under consideration.

As per usual, for distractor images and text, we set all of the rows and columns representing non-grasp options to the default identity matrix.

This reduces the number of ambiguous selections that occur when the user has no strong preference among a subset of very similar grasps, which is a fairly common outcome under experimental conditions.

Furthermore, within any 20° cone, we allow only five grasps to be added to the option set. When more than five are found, the oldest grasp (least-recently-generated) is penalized and moved to the end of the list, behind grasps with a more different approach direction. This forces some heterogeneity to remain in the grasps that are presented to the user, while additionally allowing the user to “walk” the grasping point and direction to a new approach direction, even if it is not presented in any of the previous options.

In the confirmation state, the user is shown an image corresponding to the selected grasp g_s, along with a set of distractor images presenting just the target object without a grasp. The user also has the text option “Go Back-Replan”, which returns the user to the object selection phase. Selecting the grasp again confirms the grasp for execution and sends the desired grasp to the robot.

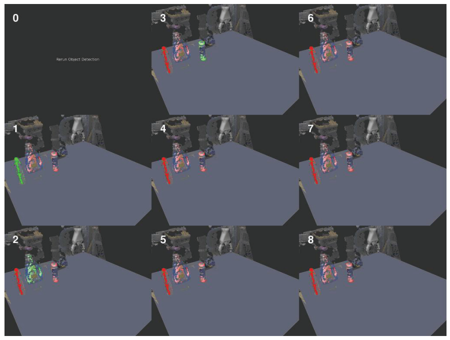

(a) Object Selection State

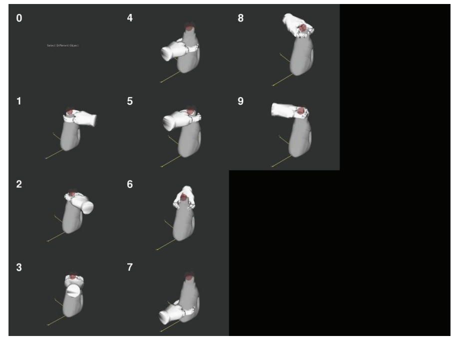

(b) Grasp Selection State

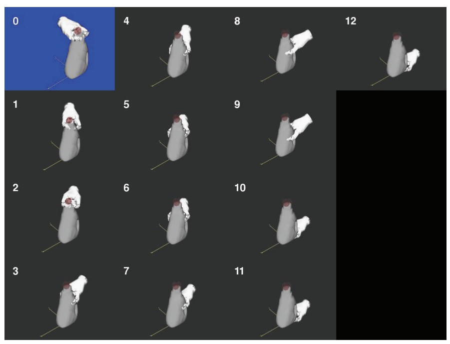

(c) Grasp Refinement State

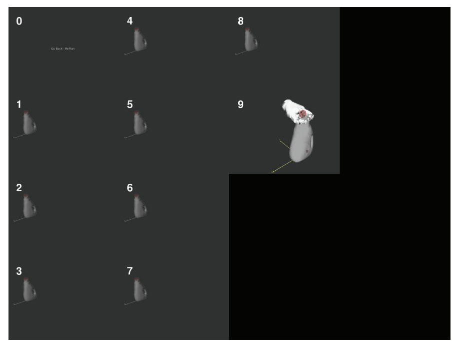

(d) Confirmation State

Fig. 4: (a) The three objects visible in the planning screen are presented to the user, with the object to be selected highlighted in green. The scene is overlaid with the pointcloud data from the Kinect, so that the user can verify that the objects have been recognized and positioned correctly within the scene.
(b) The initial set of pre-computed grasps from a database. The user may choose to go back to the previous phase and choose a different object, or to seed the online grasp planner using one of the available grasps.
(c) An updated set of grasps in the active refinement state, generated from grasp number 8 from Fig. 4b. The selected grasp g_s is reachable and highlighted in blue. Note that the generated grasps are visually distinct, but still have small groups of functionally identical grasps.
(d) There are effectively only two options in the confirmation state, which acts as a final check to determine whether the selected grasp g_s is the one that the user would like to execute.

3.4.3 Execution state

In the execution state, the user is presented with only three text options: “Restart Execution”, which restarts the execution if it has failed, telling the robot to return to its home position and attempt to grasp again; “Stop Execution”, which stops the robot from continuing the execution and returns to the confirmation state, and a set of distractor images which say “Distractor Image.”

4 Experiment

We have validated this system on three subjects, asking them to lift each of three objects visible in Fig. 5: a shaving gel bottle, a detergent bottle, and a shampoo bottle. The three objects are placed arbitrarily within the field of view of the Kinect camera, such that they do not fully occlude each other. In each case, we have verified that the object is within the reachable working area of the Mico arm, so there is at least one feasible grasp.

Fig. 5: A typical grasp of the shampoo bottle from the side in the cluttered scene. Note that the hand is just able to fit between the other objects to grasp the desired target. Note that the ability to plan this grasp in such a restricted environment is an indication that this system is very successful at handling the cluttered scene.

The subject is given the opportunity to inspect the scene, and is then asked to lift each of the objects three times from either the top, from the side, or at their own discretion. All testing was approved by the Institutional Review Board of Columbia University under Protocol IRB-AAAJ6951.

The experimental setup can be seen in Fig. 1, and a video of a subject going through the pipeline is available online[1].

4.1 Training

To demonstrate the various stages in the pipeline, the user is shown the system running under keyboard control, where each option can be selected by pressing its corresponding key. We allow the subject to walk through the stages of the grasping procedure as many times as they ask, (always less than five), while explaining what is being visualized at each step.

The EEG classifier weights described in Eq. 1 must be retrained each time the headset is placed on the user's head. This process takes approximately ten to fifteen minutes, and also serves to help familiarize the user with the RSVP paradigm.

During the training phase, the user is shown a "block" of 42 images. 40 of these images are selected uniformly from a set of fifteen object models similar to those presented during the object selection phase, while the remaining two are marked as "target" images and selected from a set of four images of bowls. Unlike the object selection phase, however, these object images are presented one at a time, without the other objects visible. There is also no "summary pane", as the images would be too small to practically see. Instead, the subject is shown the set of four potential target images.

In each block, the subject is told that there will be exactly two target images, and is asked to search for them in the sequence. After a block of images has been presented, the user is also shown the location of the two target images in the sequence and, separately, where the classifier placed those images in a list sorted by detected interest level.

The user is presented blocks at a self-paced rate until at least 20 blocks have been presented, and the user is able to consistently place the two target images in the top three sort positions. The latter condition is usually fulfilled before the former.

4.2 Results

The results of this experiment are summarized in Table 1. In all cases, the subjects were successful in selecting reasonable grasps that lifted the object. However, the system did not always detect the option that the subject wanted correctly. In the table, the third column describes the number of misselections, which represents the number of times that the user inadvertently selected an option. Because the pipeline allows stepping back through each phase, this is not fatal, though it does result in a longer task duration (shown in the fifth column in fifteen-second increments). Detected selections of known "distractor" images are not considered misselections, as they do not elicit any actual action that changes the state of the system. The fourth

[1] `http://robotics.cs.columbia.edu/jweisz/eegGraspingISRR2015.mov`

Table 1: Experimental results from three subjects

Grasp	Subject	Misselections	Refinement Iterations	Time(s)
Detergent Bottle Top	1	0	1	120
	2	2	3	150
	3	1	2	135
Detergent Bottle Side	1	0	1	120
	2	1	2	135
	3	0	1	120
Detergent Bottle Choice	1	0	10	270
	2	0	2	135
	3	3	5	180
Shampoo Bottle Top	1	0	1	135
	2	0	1	120
	3	0	1	150
Shampoo Bottle Side	1	0	1	120
	2	1	1	135
	3	0	2	135
Shampoo Bottle Choice	1	1	1	210
	2	1	3	120
	3	0	1	150
Shaving Gel Top	1	0	2	180
	2	1	1	120
	3	0	2	135
Shaving Gel Side	1	1	2	135
	2	0	1	120
	3	0	2	150
Shaving Gel Choice	1	0	2	120
	2	0	1	120
	3	0	2	180

column describes the number of iterations of the "grasp refinement" stage the user stepped through in order to find an acceptable grasp.

When grasping the detergent bottle, Subject 1 chose to attempt to grasp the handle of the object starting from a top grasp, eliciting the "walking" behavior described in Sec. 3.4.2. This necessitated a correspondingly large number of iterations of the refinement state.

The largest number of misselections came from the users accidentally selecting the option to rerun object detection in the "object selection" phase. Misselections of the wrong grasp during the "grasp refinement" stage when the user actually wanted to accept the current best grasp (g_s, above) and continue to the confirmation state also occurred, but these mistakes were quickly recoverable because similar grasps were very likely to be present in the next set of presented grasps.

5 Conclusions

These results are encouraging, and demonstrate that a relatively fast and effective pipeline based off of only EEG data is workable. The experiment revealed some issues, specifically in terms of how images are generated when representing abstract concepts (e.g. text-based image options, and the “selected” grasp in the “grasp refinement” stage).

The most common misselection was the command to redo the object detection during the “object selection” phase. This is probably because the difference between the images representing object selections and the text option image is large and somewhat startling, which elicits a reaction from the subject. A similar issue was seen in the initial attempt to use the system with Subject 2, who selected the blue image option every time it was presented in the “grasp refinement” stage, until the brightness of the background was reduced by 50%. After this modification, the subject had no trouble making the correct selection in the “grasp refinement” stage. Subject 2 sometimes had trouble making the correct selection in the “confirmation” stage, possibly because the distractor images were too similar, which makes the task too different from the one that the classification system is trained on.

There is an enormous space of design parameters that can be explored to potentially resolve some of these issues to produce a more robust system. One option would be adapting the thresholds used by the classifier based on the content of the image options. Another option would be to modify the training set of images to be more similar to the images presented in the task stage. While the current training regime has proven to be somewhat generalizable, it may not be adequately representative of the responses that are elicited by large stimuli like the blue background of the “selected image.” Finally, some calibration procedure for modifying the images based on the responses they elicit from the user may need to be incorporated into the training regime.

The extension of the online Human-in-the-Loop planner to this EEG based image streaming paradigm has just begun. In its current implementation, the subject decisions are elicited at fixed points of the pipeline. Future work will move towards attempting to integrate the EEG data in a more real-time strategy, perhaps being fully embedded into the augmented reality environment. Although this system is primarily designed as a component of an assistive robotic manipulation platform, the real time system would be useful even for able-bodied users as a fast, passive filter for eliciting feedback from the user.

Finally, another approach to be explored in the future may be to combine this method with the sEMG method from [22]. This multimodal strategy would incorporate an EEG-based classifier for directing the planner towards the user’s preferences while a facial EMG input is used to signal discrete decisions. Such a system would address the shortcomings of each individual modality – allowing the system to quickly filter reasonable options, where accuracy may be less important so long as it is somewhat conservative, while making the final selections, which affect the state of the robot and may result in inappropriate or potentially damaging behaviors, more robust.

Acknowledgements The authors would like to gratefully acknowledge the help of Dr. Paul Sajda and Dr. David Jangraw in this work. This work was supported by NSF Grant IIS-1208153.

References

1. Advanced Brain Monitoring. http://www.advancedbrainmonitoring.com/xseries/x10/
2. Kinova Robotics Mico. http://kinovarobotics.com/products/mico-robotics/
3. OpenBCI. http://www.openbci.com
4. Bell, C.J., Shenoy, P., Chalodhorn, R., Rao, R.P.: Control of a humanoid robot by a noninvasive brain–computer interface in humans. Journal of neural engineering **5**(2), 214 (2008)
5. Bryan, J., Thomas, V., Nicoll, G., Chang, L., Rao, R.: What you think is what you get: Brain-controlled interfacing for the pr2. IROS (2011)
6. Ciocarlie, M., Clanton, S., Spalding, M., Allen, P.: Biomimetic grasp planning for cortical control of a robotic hand. In: Proc. IROS, pp. 2271–2276 (2008)
7. Ciocarlie, M.T., Allen, P.K.: Hand posture subspaces for dexterous robotic grasping. The International Journal of Robotics Research **28**(7) (2009)
8. Ferrari, C., Canny, J.: Planning optimal grasps. In: Proc. of the Int. Conf. on Robotics and Automation (1992)
9. Fisher, R.A.: The use of multiple measurements in taxonomic problems. Annals of Eugenics **7**(2)
10. Gerson, A., Parra, L., Sajda, P.: Cortically coupled computer vision for rapid image search. IEEE Trans. Neural Syst. Rehabil. Eng. **14**(2), 174–179 (2006)
11. Horki, P., Solis-Escalante, T., Neuper, C., Müller-Putz, G.: Combined motor imagery and ssvep based bci control of a 2 dof artificial upper limb. Medical & biological engineering & computing **49**(5) (2011)
12. Lampe, T., Fiederer, L.D., Voelker, M., Knorr, A., Riedmiller, M., Ball, T.: A brain-computer interface for high-level remote control of an autonomous, reinforcement-learning-based robotic system for reaching and grasping. In: Proc. of Int. Conf. on Intelligent User Interfaces, IUI '14 (2014)
13. Muelling, K., Venkatraman, A., Valois, J.S., Downey, J., Weiss, J., Javdani, S., Hebert, M., Schwartz, A.B., Collinger, J.L., Bagnell, J.A.: Autonomy infused teleoperation with application to bci manipulation. arXiv preprint arXiv:1503.05451 (2015)
14. Pohlmeyer, E.A., Wang, J., Jangraw, D.C., Lou, B., Chang, S.F., Sajda, P.: Closing the loop in cortically-coupled computer vision: a brain-computer interface for searching image databases. J. Neural Eng. **8**(3) (2011)
15. Postelnicu, C.C., Talaba, D., Toma, M.I.: Controlling a robotic arm by brainwaves and eye movement. In: Technological Innovation for Sustainability. Springer (2011)
16. Royer, A.S., Rose, M.L., He, B.: Goal selection versus process control while learning to use a brain–computer interface. Journal of neural engineering **8**(3), 036,012 (2011)
17. Sajda, P., Pohlmeyer, E., Wang, J., Parra, L., Christoforou, C., Dmochowski, J., Hanna, B., Bahlmann, C., Singh, M., Chang, S.F.: In a blink of an eye and a switch of a transistor: Cortically coupled computer vision. Proceedings of the IEEE **98**(3), 462–478 (2010)
18. Sucan, I.A., Chitta, S.M.: (2013). URL `http://moveit.ros.org`
19. Vogel, J., Haddadin, S., Simeral, J.D., Stavisky, S.D., Bacher, D., Hochberg, L.R., Donoghue, J.P., van der Smagt, P.: Continuous control of the dlr light-weight robot iii by a human with tetraplegia using the braingate2 neural interface system. In: Experimental Robotics, pp. 125–136. Springer (2014)
20. Waytowich, N., Henderson, A., Krusienski, D., Cox, D.: Robot application of a brain computer interface to staubli tx40 robots-early stages. In: World Automation Congress. IEEE (2010)
21. Weisz, J., Barszap, A.G., Joshi, S.S., Allen, P.K.: Single muscle site semg interface for assistive grasping. IROS (2014)
22. Weisz, J., Elvezio, C., Allen, P.K.: A user interface for assistive grasping. 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (2013)
23. Weisz, J., Shababo, B., Dong, L., Allen, P.K.: Grasping with your face. Springer Tracts in Advanced Robotics pp. 435–448 (2013)