

Brain–Controlled Wheelchairs: A Robotic Architecture

Tom Carlson, *Member IEEE*, and José del R. Millán, *Senior Member IEEE*

Abstract—Independent mobility is core to being able to perform activities of daily living by oneself. However, powered wheelchairs are not an option for a large number of people who are unable to use conventional interfaces, due to severe motor–disabilities. For some of these people, non–invasive brain–computer interfaces (BCIs) offer a promising solution to this interaction problem and in this article we present a shared control architecture that couples the intelligence and desires of the user with the precision of a powered wheelchair. We show how four healthy subjects are able to master control of the wheelchair using an asynchronous motor–imagery based BCI protocol and how this results in a higher overall task performance, compared with alternative synchronous P300–based approaches.

I. INTRODUCTION

Millions of people around the world suffer from mobility impairments and hundreds of thousands of them rely upon powered wheelchairs to get on with their activities of daily living [1]. However, many patients are not prescribed powered wheelchairs at all, either because they are physically unable to control the chair using a conventional interface, or because they are deemed incapable of driving safely [2]. Consequently, it has been estimated that between 1.4 and 2.1 million wheelchair users might benefit from a smart powered wheelchair, if it were able to provide a degree of additional assistance to the driver [3].

In our work with brain–actuated wheelchairs, we target a population who are—or will become—unable to use conventional interfaces, due to severe motor–disabilities. Non–invasive brain–computer interfaces (BCIs) offer a promising new interaction modality, that does not rely upon a fully–functional peripheral nervous system to mechanically interact with the world and instead uses the brain activity directly. However, mastering the use of a BCI, like with all new skills, does not come without a few challenges. Spontaneously performing mental tasks to convey one’s intentions to a BCI can require a high level of concentration, so it would result in a fantastic mental workload, if one had to precisely control every movement of the wheelchair. Furthermore, due to the noisy nature of brain signals, we are currently unable to achieve the same information rates that you might get from a joystick, which would make it difficult to wield such levels of control even if one wanted to.

Thankfully, we are able to address these issues through the use of intelligent robotics, as will be discussed. Our wheelchair uses the notion of shared control to couple the intelligence of the user with the precise capabilities of a robotic wheelchair,

given the context of the surroundings [4]. It is this synergy, which begins to make brain–actuated wheelchairs a potentially viable assistive technology of the not–so–distant future.

In this paper we describe the overall robotic architecture of our brain–actuated wheelchair. We begin by discussing the brain computer interface, since the human is central to our design philosophy. Then, the wheelchair hardware and modifications are described, before we explain how the shared control system fuses the multiple information sources in order to decide how to execute appropriate manoeuvres in cooperation with the human operator. Finally, we present the results of an experiment involving four healthy subjects and compare them with those reported on other brain–actuated wheelchairs. We find that our continuous control approach offers a very good level of performance, with experienced BCI wheelchair operators achieving a comparable performance to that of a manual benchmark condition.

II. BRAIN COMPUTER INTERFACES (BCI)

The electrical activity of the brain can be monitored in real–time using an array of electrodes, which are placed on the scalp in a process known as electroencephalography (EEG). In order to bypass the peripheral nervous system, we need to find some reliable correlates in the brain signals that can be mapped to the intention to perform specific actions. In the next two subsections, we first discuss the philosophy of different BCI paradigms, before explaining our chosen asynchronous implementation for controlling the wheelchair.

A. The BCI Philosophy

Many BCI implementations, rely upon the subject attending to visual stimuli, which are presented on a screen. Consequently, researchers are able to detect a specific event–related potential in the EEG, known as the P300, which is exhibited 300 ms after a rare stimulus has been presented. For example, in one P300–based BCI wheelchair, the user is presented with a 3×3 grid of possible destinations from a known environment (e.g. the bathroom, the kitchen etc., within the user’s house), which are highlighted in a standard oddball paradigm [5]. The user then has to focus on looking at the particular option to which they wish to drive. Once the BCI has detected their intention, the wheelchair drives autonomously along a pre–defined route and the user is able to send a mental *emergency stop* command (if required) with an average of 6 seconds delay.

Conversely, another BCI wheelchair, which is also based upon the P300 paradigm doesn’t restrict the user to navigating in known, pre–mapped environments. Instead, in this design, the user is able to select subgoals (such as close left, far right,

T. Carlson and J. d. R. Millán are with the Defitech Chair in Non–Invasive Brain Machine Interface (CNBI), École Polytechnique Fédérale de Lausanne (EPFL), Station 11, 1015 Lausanne, Switzerland. tom.carlson@epfl.ch

mid-ahead etc.) from an augmented reality matrix superimposed on a representation of the surrounding environment [6]. To minimise errors (at the expense of command delivery time), after a subgoal has been pre-selected, the user then has to focus on a *validation* option. This gives users more flexibility in terms of following trajectories of their choice, however, the wheelchair has to stop each time it reaches the desired subgoal and wait for the next command (and validation) from the user. Consequently, when driving to specific destinations, the wheelchair was stationary for more time than it was actually moving (as can be seen in Fig. 8 of [6]).

Our philosophy is to keep as much authority with the users as possible, whilst enabling them to dynamically generate natural and efficient trajectories. Rather than using external stimuli to evoke potentials in the brain, as is done in the P300 paradigm, we allow the user to spontaneously and asynchronously control the wheelchair by performing a motor imagery task. Since this does not rely on visual stimuli, it does not interfere with the visual task of navigation. Furthermore, when dealing with motor-disabled patients, it makes sense to use motor imagery, since this involves a part of the cortex, which may have effectively become redundant; i.e. the task does not interfere with the residual capabilities of the patient.

Previously, we have demonstrated that it is possible to drive a wheelchair using such a protocol [7]. However, this earlier system relied upon an expensive laser scanner to map the environment. In Section III, we show how a combination of relatively cheap sensors is sufficient to provide environmental feedback to the wheelchair controller. Moreover, the original protocol required the user to continuously deliver commands to drive the wheelchair, which resulted in a high user workload. Our current BCI protocol, coupled with shared control (see Section IV) has reduced this workload.

In our motor imagery (MI) paradigm, the user is required to imagine the kinaesthetic movement of the left hand, the right hand or both feet, yielding three distinct classes. During the BCI training process, we select the two most discriminable classes to provide a reliable mapping from the MI tasks to control actions (e.g. imagine left hand movements to deliver a turn left command and right hand movements to turn right).

To control our BCI wheelchair, at any moment, the user can spontaneously issue a high-level turn left or turn right command. When one of these two turning commands is not delivered by the user, a third implicit class of *intentional non-control* exists, whereby the wheelchair continues to travel forward and automatically avoid obstacles where necessary. Consequently, this reduces the user's cognitive workload. The implementation will be discussed in Section IV-D.

B. The BCI Implementation

Since we are interested in detecting motor imagery, we acquire monopolar EEG at a rate of 512 Hz from the motor cortex using 16 electrodes (see Fig. 1). The electrical activity of the brain is diffused as it passes through the skull, which results in a spatial blur of the signals, so we apply a Laplacian filter, which attenuates the common activity between neighbouring electrodes and consequently improves




Fig. 1: The active electrode placement over the motor cortex for the acquisition of EEG data, based on the International 10-20 system (nose at top).

our signal to noise ratio. After the filtering, we estimate the power spectral density (PSD) over the last second, in the band 4–48 Hz with a 2 Hz resolution [8]. It is well known that when one performs motor imagery tasks, corresponding parts of the motor cortex are activated, which, as a result of event related desynchronisation, yields a reduction in the mu band power (~ 8 –13 Hz) over these locations (e.g. the right hand corresponds to approximately C1 and the left hand to approximately C2 in Fig. 1). In order to detect these changes, we estimate the PSD features every 62.5 ms (i.e. 16 times per second) using the Welch method with 5 overlapped (25%) Hanning windows of 500 ms.

Every person is different, so we have to select the features that best reflect the motor-imagery task for each subject. Therefore, canonical variate analysis (CVA) is used to select subject-specific features that maximize the separability between the different tasks and that are most stable (according to cross validation on the training data) [9]. These features are then used to train a Gaussian classifier [10]. Decisions with a confidence on the probability distribution that are below a given rejection threshold are filtered out. Finally, evidence about the executed task is accumulated using an exponential smoothing probability integration framework [11]. This helps to prevent commands from being delivered accidentally.

III. WHEELCHAIR HARDWARE

Our brain-controlled wheelchair is based upon a commercially available mid-wheel drive model by Invacare that we have modified. First, we have developed a remote joystick module that acts as an interface between a laptop computer and the wheelchair's CANBUS-based control network. This allows us to control the wheelchair directly from a laptop computer. Second, we have added a pair of wheel-encoders to the central driving wheels in order to provide the wheelchair with feedback about its own motion. Third, an array of ten sonar sensors and two webcams have been added to the wheelchair to provide environmental feedback to the controller. Fourth, we have mounted an adjustable 8" display to provide visual feedback to the user. Fifth, we have built a power distribution

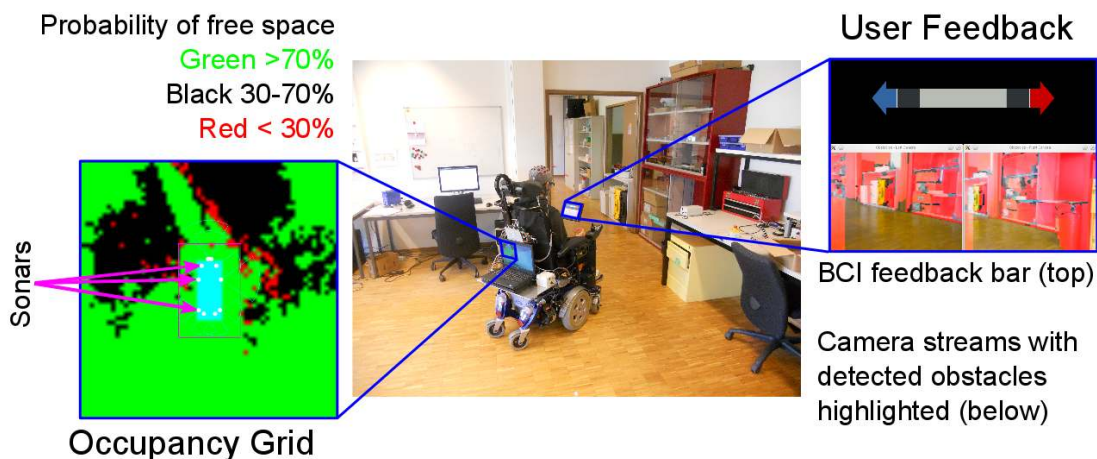


Fig. 2: The complete brain-actuated wheelchair. The wheelchair’s knowledge of the environment is acquired by the fusion of complementary sensors and is represented as a probabilistic occupancy grid. The user is given feedback about the current status of the BCI and about the wheelchair’s knowledge of the environment.

unit, to hook up all the sensors, the laptop and the display to the wheelchair’s batteries. The complete BCI wheelchair platform is shown in Fig. 2. The positions of the sonars are indicated by the white dots in the centre of the occupancy grid, whereas the two webcams are positioned forward-facing, directly above each of the front castor wheels.

A. Wheel-encoders

The encoders return 128 ticks per revolution and are geared up to the rim of the drive wheels, resulting in a resolution of 2.75×10^{-3} metres translation of the inflated drive wheel per encoder tick. We use this information to calculate the average velocities of the left and right wheels for each time-step. Not only is this important feedback to regulate the wheelchair control signals, but we also use it as the basis for dead reckoning (or estimating the trajectory that has been driven). We apply the simple differential drive model derived in [12]. To ensure that the model is always analytically solvable, we neglect the acceleration component. In practice, since in this application we are only using the odometry to update a $6\text{m} \times 6\text{m}$ map, this does not prove to be a problem. However, if large degrees of acceleration or slippage occur and the odometry does not receive any external correcting factors, the model will begin to accumulate significant errors [12].

IV. SHARED CONTROL ARCHITECTURE

The job of the shared controller is to determine the meaning of the vague, high-level user input (e.g. turn left, turn right, keep going straight), given the context of the surrounding environment [4]. We do not want to restrict ourselves to a known, mapped environment—since it may change at any time (e.g. due to human activities)—so the wheelchair must be capable of perceiving its surroundings. Then, the shared controller can determine what actions should be taken, based upon the user’s input, given the context of the surroundings. The overall robotic shared control architecture is depicted in Fig. 3 and we discuss the perception and planning blocks of the controller over then next few subsections.

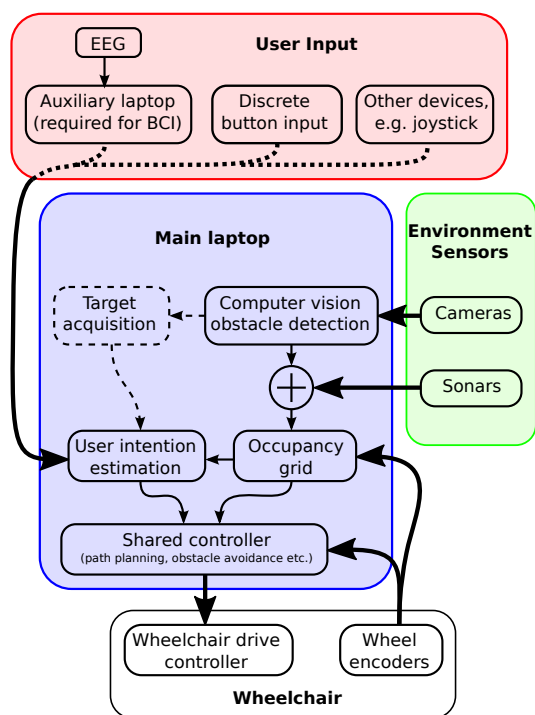


Fig. 3: The user’s input is interpreted by the shared controller given the context of the surroundings. The environment is sensed using a fusion of complementary sensors, then the shared controller generates appropriate control signals to navigate safely, based upon the user input and the occupancy grid.

A. Perception

Unlike for humans, perception in robotics is difficult. To begin with, choosing appropriate sensors is a not a trivial task and tends to result in a trade-off between many issues, such as: cost, precision, range, robustness, sensitivity, complexity of post-processing and so on. Furthermore, no single sensor by it-




Fig. 4: The obstacle-detection algorithm is based upon a computer-vision approach proposed in [13], but adapted for monocular vision. The floor is deemed to be the largest region that touches the base of the image, yet does not cross the horizon.

self seems to be sufficient. For example, a planar laser scanner may have excellent precision and range, but will only detect a table’s legs, reporting navigable free space between them. Other popular approaches, like relying solely upon cheap and readily available sonar sensors have also been shown to be unreliable for such safety-critical applications [14]. To overcome these problems, we propose to use the synergy of two low-cost sensing devices to compensate for each other’s drawbacks and complement each other’s strengths. Therefore, we use an array of ten close-range sonars, with a wide detection beam, coupled with two standard off-the-shelf USB webcams, for which we developed an effective obstacle detection algorithm. We then fuse the information from each sensor modality into a probabilistic occupancy grid, as will be discussed in Section IV-C.

B. Computer Vision-Based Obstacle Detection

The obstacle detection algorithm is based on monocular image processing from the webcams, which ran at 10Hz. The concept of the algorithm is to detect the floor region and label everything that does not fall into this region as an obstacle; we follow an approach similar to that proposed in [13], albeit with monocular vision, rather than using a stereo head.

The first step is to segment the image into constituent regions. For this, we use the watershed algorithm, since it is fast enough to work in real-time [15]. We take the original image (Fig 4a) and begin by applying the well-known Canny edge-detection, as shown in Fig. 4b. A distance transform is then applied, such that each pixel is given a value that represents the minimum Euclidean distance to the nearest edge. This results in the relief map shown in Fig. 4c, with a set of peaks (the farthest points from the edges) and troughs (the edges themselves). The watershed segmentation algorithm itself is applied to this relief map, using the peaks as markers, which results in an image with a (large) number of segments (see Fig. 4d). To reduce the number of segments, adjacent regions with similar average colours are merged. Finally, the average colour of the region that has the largest number of pixels along the base of the image is considered to be the floor. All the remaining regions in the image are classified either as obstacles or as navigable floor, depending on how closely they match the newly-defined floor colour. The result is shown in Fig. 4e, where the detected obstacles are highlighted in red.

Since we know the relative position of the camera and its lens distortion parameters, we are able to build a local

occupancy grid that can be used by the shared controller, as is described in the following section.

C. Updating the Occupancy Grid

At each time-step, the occupancy grid is updated to include the latest sample of sensory data from each sonar and the output of the computer vision obstacle detection algorithm. We extend the histogram grid construction method described in [16], by fusing information from multiple sensor types into the same occupancy grid. For the sonars, we consider a ray to be emitted from each device along its sensing axis. The likelihood value of each occupancy grid cell that the ray passes through is decremented, whilst the final grid cell (at the distance value returned by the sonar) is incremented. A similar process is applied for each column of pixels from the computer vision algorithm, as shown in Fig. 5. The weight of each increment and decrement is determined by the confidence we have for each sensor at that specific distance. For example, the confidence of the sonar readings being correct in the range 3 cm to 50 cm is high, whereas outside that range it is zero (note that the sonars are capable of sensing up to 6 m, but given that they are mounted low on the wheelchair, the reflections from the ground yield a practical limit of 0.5 m). Similarly, the computer vision algorithm only returns valid readings for distances between 0.5 m and 3 m. Using this method, multiple sensors and sensor modalities can be integrated into the planning grid.

As the wheelchair moves around the environment, the information from the wheel-encoder based dead-reckoning system is used to translate and rotate the occupancy grid cells, such that the wheelchair remains at the centre of the map. In this way, the cells accumulate evidence over time from multiple sensors and sensor modalities. As new cells enter the map at the boundaries, they are set to “unknown”, or 50% probability of being occupied, until new occupancy evidence (from sensor readings) becomes available.

D. Motion Planning

All the motion planning is done at the level of the occupancy grid, which integrates the data from multiple sensors. We base our controller on a dynamical system approach to navigation, since this easily allows us to incorporate the notion of obstacles (repellers) and targets (attractors), and results in naturally smooth trajectories [17]. Previously, we have implemented




Fig. 5: Each column of pixels is scanned from bottom to top, in order to detect the nearest obstacle (assuming it intersects with the ground). The estimated distance from the wheelchair to this obstacle is a function of the (x, y) pixel coordinates and the camera distortion parameters.

such a control strategy on a circular mobile robotic platform, which was successfully controlled by motor-disabled patients using a BCI [18].

With no user input, the dynamical system causes the wheelchair to move forwards and automatically avoid any obstacles that it comes across. In practice, this is realised by adding repellers into the dynamical system according to the obstacle densities in the occupancy grid. Rather than simply looking at the densities in radial directions from the robot, as was sufficient in [18]—to account for the fact that the wheelchair’s shape and motion is more complex than the circular robot—we define a set of $N = 10$ zones within the occupancy grid, as shown in Fig. 6. These zones are split into three sets, such that if obstacles were present in them: $\Omega_c = \{RB1, LC1, LF1, LF2\}$ would cause clockwise rotations of the wheelchair, $\Omega_a = \{LB1, RC1, RF1, RF2\}$ would cause anticlockwise rotations, and $\Omega_n = \{F1, F2\}$ would not affect the rotational velocity of the wheelchair. Each zone, $z_i \in \Omega$, has a centre (z_{ix}, z_{iy}) and an associated repulsion strength $\lambda_i < 0 \in \Lambda$, which is determined according to the position of the zone relative to the wheelchair, such that $\Lambda = \Lambda_c \cup \Lambda_a \cup \Lambda_n$. The likelihood of there being an obstacle in each zone is $\varphi_i \in [0, 1]$. The rotational velocity ω is then:

$$\omega = K_\omega \sum_{i=1}^N \lambda_i S_i \varphi_i, \quad \omega \in [-\omega_{max}, +\omega_{max}], \quad (1)$$

$$K_\omega = \frac{\omega_{max}}{\sum_{\lambda_i \in \Lambda_c} |\lambda_i|} \quad (2)$$

$$S_i = \text{sgn}(-z_{ix} \times z_{iy}), \quad S_i \in \{-1, 0, +1\}, \quad (3)$$

where the constant K_ω ensures that $|\omega|$ is not greater than the maximum possible rotational velocity ω_{max} . Note that K_ω in (1) assumes that the obstacle detection zones in Ω_c and Ω_a , and their corresponding Λ_c and Λ_a values, are symmetric, as it is in our case. In general, this makes sense, since you would expect symmetric behaviour from the wheelchair for

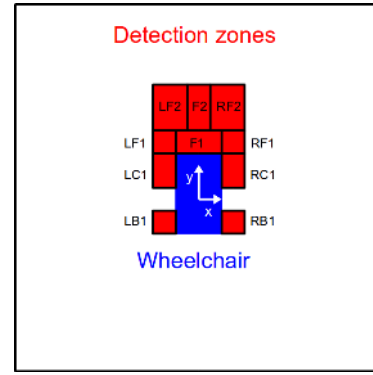


Fig. 6: The wheelchair centred in an $8\text{ m} \times 8\text{ m}$ occupancy grid (to scale). The wheelchair obstacle detection zones are labelled and the origin of the coordinate system is on the centre of the wheelchair’s driving axle.

symmetric stimuli¹. S_i simply encodes the sign (direction) of the resultant rotation, assuming that (z_{ix}, z_{iy}) is the Cartesian centre of the i -th zone, in a coordinate system whose origin is in the centre of the wheelchair’s axis (as shown in Fig. 6).

Similarly, for the translational velocity, v , each zone has an associated translational repeller, $\gamma_i < 0$:

$$v = v_{max} + \sum_{i=1}^N \gamma_i \varphi_i, \quad v \in [0, +v_{max}]. \quad (4)$$

The γ_i values are chosen empirically according to the dynamics of the wheelchair and the reliability of the sensors, such that, for example, when the zone $F1$ reaches 50% occupancy, the wheelchair will stop. Therefore we set the γ that corresponds to zone $F1$ to be $-2v_{max}$, whereas the γ that corresponds to zone $F2$ is $-0.75v_{max}$.

When the user issues a BCI command (either a high-level turn left, or turn right), the wheelchair should turn up to a maximum of 45° in the direction indicated, depending on the environmental context. To achieve this, an additional, corresponding virtual attractor zone is placed in the occupancy grid 1 m in front of the wheelchair, at an angle of 45° in the direction indicated by the BCI command. This attractor zone has a corresponding $\varphi_i = 1.0$, $\lambda_i = 0.5$ and $\gamma_i = 0.0$, such that in practice it only affects the rotational velocity dynamical system. Note that λ is a positive value when acting as an attractor. The attractor remains in the dynamical system, until the wheelchair has turned up to 45° in the corresponding direction, or a new BCI command is delivered, or until a timeout has been reached (in our case 4 seconds), at which point, it is removed.

We extend the dynamical system, by exploiting the fact that we have a human in the loop, to enable an additional *docking* behaviour. Such a behaviour is important if the system is to be useful outside of experimental lab conditions. Therefore, if the occupancy of zone $F1$ or $F2$ is greater than an empirically set activation threshold, K_T , providing there is

¹If this is not the case, one should take K_ω to be the maximum value of K_{ω_c} and K_{ω_a} , computed using Λ_c and Λ_a , respectively. However, this would result in asymmetric behaviour of the wheelchair.

no user input, the rotational velocity of the chair will be set to 0 rad^{-s} . The translational velocity will still be controlled by the dynamical system, such that the wheelchair will slow down smoothly and stop in front of the object. At any point, the user is able to deliver a left or right command to initiate an obstacle avoidance manoeuvre. If the user remains in a state of intentional non-control, once the wheelchair has completed the docking procedure, it will remain stationary and wait for further user input.

In the current implementation, the user is not able to stop the chair in free space, instead the chair will stop when it has docked to a potential target. In future this control strategy could easily be extended to include an additional BCI command (or another biosignal, in the case of a hybrid approach) to implement an explicit stop signal.

V. EVALUATION

We demonstrate that both naïve and experienced BCI wheelchair operators are able to complete a navigation task successfully. Furthermore, unlike in P300-based systems, not only was the user in continuous spontaneous control of the wheelchair, but the resultant trajectories were smooth and intuitive (i.e. no stopping, unless there was an obstacle, and users could voluntarily control the motion at all times).

A. Participants

Mastering a motor imagery BCI requires extensive training, over a period of weeks or months to generate stable volitional control; it is not simply a case of putting a cap on and starting to drive. Therefore, we have performed an initial evaluation with four healthy male subjects, aged 23–28. All subjects were experienced BCI users, who had participated in at least 12 hours of online motor imagery BCI training and other BCI experiments over the previous few months. They all had some previous experience of driving a BCI-based tele-presence mobile robot, which requires a better level of performance, compared to simply moving a cursor on a screen [18]. Subjects s1 and s2 had no previous experience of driving a BCI-controlled wheelchair, whereas subjects s3 and s4 had each clocked-up several hours of driving the BCI wheelchair. Subject s1 used motor imagery of both feet to indicate *turn left* and of the right hand to mean *turn right*; all the other subjects used left hand motor imagery to *turn left* and right hand motor imagery to *turn right*.

B. Experiment Protocol

As a benchmark, the subject was seated in the wheelchair and was instructed to perform an online BCI session, before actually driving. In this online session, the wheelchair remained stationary and the participant simply had to perform the appropriate motor imagery task to move a cursor on the wheelchair screen in the direction indicated by a cue arrow. There was a randomized balanced set of 30 trials, separated by short resting intervals, which lasted around 4–5 mins, depending on the performance of the subject.

After the online session, participants were given 15–30 minutes to familiarise themselves with driving the wheelchair




Fig. 7: Trajectories followed by subject s3 on one of the manual benchmark trials (left), compared with one of the BCI trials (right). These trajectories were reconstructed from odometry using the independent reconstruction method [19].

using each of the control conditions: a two button *manual* input, which served as a benchmark, and the BCI system. Both input paradigms allowed the users to issue left and right commands at an inter-trial interval of one second.

The actual task was to enter a large open-plan room through a doorway from a corridor, navigate to two different tables, whilst avoiding obstacles and passing through narrow openings (including other non-target tables, chairs, ornamental trees and a piano), before finishing by reaching a second doorway exit of the room (as shown in Fig 7). When approaching the target tables, the participants were instructed to wait for the wheelchair to finish docking to the table, then once it had stopped they should issue a turning command to continue on their journey. The trials were counter-balanced, such that users began with a manual trial, then performed two BCI trials and finished with another manual trial.

C. Results and Discussion

All subjects were able to achieve a remarkably good level of control in the stationary online BCI session, as can be seen in Table I. Furthermore, the actual driving task was completed successfully by every subject, for every run and no collisions occurred. A comparison between the typical trajectories followed under the two conditions is shown in Fig 7. The statistical tests reported in this section are paired Student's *t*-tests.

A great advantage that our asynchronous BCI wheelchair brings, compared with alternative approaches like the P300-based chairs, is that the driver is in continuous control of the wheelchair. This means that not only does the wheelchair follow natural trajectories, which are determined in real-time by the user (rather than following predefined ones, like in [5]), but also that the chair spends a large portion of the

navigation time actually moving (see Fig. 8). This is not the case with some state-of-the-art P300-controlled wheelchairs, where the wheelchair has to spend between 60% and 80% of the manoeuvre time stationary, waiting for input from the user (c.f. Fig. 8 of this article with Fig. 8 of [6]).

In terms of path efficiency, there was no significant difference ($p = 0.6107$) across subjects between the distance travelled in the manual benchmark condition (43.1 ± 8.9 m) and that in the BCI condition (44.9 ± 4.1 m). Although the actual environments were different, the complexity of the navigation was comparable to that of the tasks investigated on a P300-based wheelchair in [6]. In fact, the average distance travelled for our BCI condition (44.9 ± 4.1 m), was greater than that in the longest task of [6] (39.3 ± 1.3 m), yet on average our participants were able to complete the task in 417.6 ± 108.1 s, which was 37% faster than the 659 ± 130 s reported in [6]. This increase in speed might (at least partly) be attributed to the fact that our wheelchair was not stationary for such a large proportion of the trial time.

Across subjects, it took an average of 160.0 s longer to complete the task under the BCI condition (see Fig. 8, $p = 0.0028$). This is probably due to a combination of subjects issuing manual commands with a higher temporal accuracy and a slight increase in the number of turning commands that were issued when using the BCI (c.f. Fig. 7), which resulted in a lower average translational velocity. It should be noted that in the manual benchmark condition, the task completion time varied slightly from subject to subject, as the experiments were carried out on different days, where the changes in lighting conditions affected the computer vision system. On brighter days, some shadows and reflections from the shiny wooden floor caused the wheelchair to be cautious and slow down earlier than on dull days, until the sonars confirmed that actually there was not an obstacle present. Therefore, it makes more sense to do a within subjects comparison, looking at the performance improvement or degradation on a given day, rather than comparing absolute performance values between subjects on different days.

From Fig. 8, it can be seen that for the inexperienced users (s1 and s2), there was some discrepancy in the task completion time between the benchmark manual condition and the BCI condition. However, for the experienced BCI wheelchair users (s3 and s4), the performance in the BCI condition is much closer to the performance in the manual benchmark condition. This is likely to be due to the fact that performing a motor-imagery task, whilst navigating and being seated on a moving wheelchair, is much more demanding than simply moving a cursor on the screen (c.f. the stationary online BCI session of Table I). In particular, aside from the increased workload, when changing from a task where one has to deliver a particular command as fast as possible following a cue, to a task that involves navigating asynchronously in a continuous control paradigm, the timing of delivering commands becomes very important. In order to drive efficiently, the user needs to develop a good mental model of how the entire system behaves (i.e. the BCI, coupled with the wheelchair) [20]. Clearly, through their own experience, subjects s3 and s4 had developed such mental models and were therefore able to

TABLE I: Confusion matrices of the left and right classes and accuracy for the online session, for each subject, before actually controlling the wheelchair.

	s1		s2		s3		s4	
	L	R	L	R	L	R	L	R
Left class	13	2	12	3	14	1	15	0
Right class	0	15	0	15	0	15	0	15
Accuracy (%)	93.3		90.0		96.7		100.0	




Fig. 8: The average time required to complete the task for each participant in a benchmark manual condition (left bars) and the BCI condition (right bars). The wheelchair was stationary, waiting user input, only for a small proportion of the trial.

anticipate when they should begin performing a motor imagery task to ensure that the wheelchair would execute the desired turn at the correct moment. Furthermore, they were also more experienced in refraining from *accidentally* delivering commands (intentional non-control) during the periods where they wanted the wheelchair to drive straight forwards and autonomously avoid any obstacles. Conversely, despite the good online BCI performance of subjects s1 and s2, they had not developed such good mental models and were less experienced in controlling the precise timing of the delivery of BCI commands. Despite this, the use of shared control ensured that all subjects, whether experienced or not, could achieve the task safely and at their own pace, enabling continuous mental control over long periods of time (> 400 s, almost 7 minutes).

VI. CONCLUSION

In this article, we have seen how a viable brain-actuated wheelchair can be constructed by combining a brain computer interface with a commercial wheelchair, via a shared control layer. The shared controller couples the intelligence and desires of the user with the precision of the machine. We have found that this enabled both experienced and inexperienced users alike to safely complete a driving task that involved docking to two separate tables along the way.

Furthermore, we have compared our results with those published on other state-of-the-art brain-controlled wheelchairs that are based on an alternative synchronous stimulus-driven protocol (P300). Our asynchronous motor-imagery approach

gives users greater flexibility and authority over the actual trajectories driven, since it allowed users to interact with the wheelchair spontaneously, rather than having to wait for external cues as was the case with [5], [6]. Moreover, combining our BCI with a shared control architecture allowed users to dynamically produce intuitive and smooth trajectories, rather than relying on predefined routes [5] or having to remain stationary for the majority of the navigation time [6].

Although there was a cost in terms of time for inexperienced users to complete the task using the BCI input compared with a manual benchmark, experienced users were able to complete the task in comparable times under both conditions. This is probably as a result of them developing good mental models of how the coupled BCI-shared control system behaves.

In summary, the training procedure for spontaneous motor imagery-based BCIs might take a little longer than that for stimulus-driven P300 systems, but ultimately it is very rewarding. After learning to modulate their brain signals appropriately, we have demonstrated that both experienced and inexperienced users were able to master a degree of continuous control that was sufficient to safely operate a wheelchair in a real world environment. They were always successful in completing a complex navigation task using mental control over long periods of time. One participant remarked that the motor-imagery BCI learning process is similar to that of athletes or musicians training to perfect their skills: when they eventually succeed they are rewarded with a great sense of self-achievement.

VII. THE FUTURE

We have already begun evaluating our brain-actuated wheelchair with motor-disabled patients in partnership with medical practitioners and rehabilitation clinics, but this is an arduous process that will take significantly longer than the initial trials with healthy subjects. This is for a number of reasons, not least that patients tend to take part in fewer sessions per week and generally tire more quickly than healthy participants. This leads us to another one of the exciting new challenges for the future of such shared control systems. Since each user's needs are not only different, but also change throughout the day (e.g. due to fatigue, frustration etc.), it is not sufficient that a shared control system offers a constant level of assistance. Furthermore, if this assistance is not well-matched to the user, it could lead to degradation or loss of function. Therefore we are developing shared control systems that adapt to the user's evolving needs, given not only the environmental context, but also the state of the user. This will allow people to use intelligent assistive devices in their day-to-day lives for extended periods of time.

VIII. ACKNOWLEDGEMENT

This work is supported by the European ICT Project TOBI (FP7-224631) and the Swiss National Science Foundation through the NCCR Robotics. This paper only reflects the authors' views and funding agencies are not liable for any use that may be made of the information contained herein.

REFERENCES

- [1] A. van Drongelen, B. Roszek, E. S. M. Hilbers-Modderman, M. Kallewaard, and C. Wassenaar, "Wheelchair incidents," Rijksinstituut voor Volksgezondheid en Milieu RIVM, Bilthoven, NL, Tech. Rep., November 2002, accessed February, 2010. [Online]. Available: <http://hdl.handle.net/10029/9183>
- [2] A. Frank, J. Ward, N. Orwell, C. McCullagh, and M. Belcher, "Introduction of a new NHS electric-powered indoor/outdoor chair (EPIOC) service: benefits, risks and implications for prescribers," *Clinical Rehabilitation*, no. 14, pp. 665–673, 2000.
- [3] R. C. Simpson, E. F. LoPresti, and R. A. Cooper, "How many people would benefit from a smart wheelchair?" *Journal of Rehabilitation Research and Development*, vol. 45, no. 1, pp. 53–71, 2008.
- [4] T. Carlson and Y. Demiris, "Collaborative control for a robotic wheelchair: Evaluation of performance, attention, and workload," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 42, no. 3, pp. 876–888, 2012.
- [5] B. Rebsamen, C. Guan, H. Zhang, C. Wang, C. Teo, M. Ang, and E. Burdet, "A brain controlled wheelchair to navigate in familiar environments," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 6, pp. 590–598, dec. 2010.
- [6] I. Iturrate, J. Antelis, A. Kübler, and J. Minguez, "A noninvasive brain-actuated wheelchair based on a P300 neurophysiological protocol and automated navigation," *IEEE Transactions on Robotics*, vol. 25, no. 3, pp. 614–627, june 2009.
- [7] J. d. R. Millán, F. Galán, D. Vanhooydonck, E. Lew, J. Philips, and M. Nuttin, "Asynchronous non-invasive brain-actuated control of an intelligent wheelchair," in *Proc. 31st Annual Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2009, pp. 3361–3364.
- [8] J. d. R. Millán, F. Renkens, J. Mourião, and W. Gerstner, "Noninvasive brain-actuated control of a mobile robot by human EEG," *IEEE Trans Biomed Eng*, vol. 51, no. 6, pp. 1026–1033, 2004.
- [9] F. Galán, P. W. Ferrez, F. Oliva, J. Guàrdia, and J. d. R. Millán, "Feature extraction for multi-class BCI using canonical variates analysis," in *IEEE Int Symp Intelligent Signal Processing*, 2007.
- [10] J. d. R. Millán, P. W. Ferrez, F. Galán, E. Lew, and R. Chavarriaga, "Non-invasive brain-machine interaction," *Int J Pattern Recognition and Artificial Intelligence*, vol. 22, no. 5, pp. 959–972, 2008.
- [11] S. Perdakis, H. Bayati, R. Leeb, and J. d. R. Millan, "Evidence accumulation in asynchronous BCI," *International Journal of Bioelectromagnetism*, vol. 13, no. 3, pp. 131–132, 2011.
- [12] G. Lucas, "A tutorial and elementary trajectory model for the differential steering system of robot wheel actuators," The Rossum Project, Tech. Rep., May 2000. [Online]. Available: <http://rosum.sourceforge.net/papers/DiffSteer/>
- [13] E. Fazl-Ersi and J. Tsotsos, "Region classification for robust floor detection in indoor environments," in *Image Analysis and Recognition*, M. Kamel and A. Campilho, Eds. Springer Berlin / Heidelberg, 2009, vol. 5627, pp. 717–726.
- [14] T. Dutta and G. Fernie, "Utilization of ultrasound sensors for anti-collision systems of powered wheelchairs," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 13, no. 1, pp. 24–32, March 2005.
- [15] S. Beucher, "The watershed transformation applied to image segmentation," *Scanning Microscopy International*, vol. 6, pp. 299–314, 1992.
- [16] J. Borenstein and Y. Koren, "The vector field histogram - fast obstacle avoidance for mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, pp. 278–288, 1991.
- [17] G. Schöner, M. Dose, and C. Engels, "Dynamics of behavior: Theory and applications for autonomous robot architectures," *Robot. Autonomous Syst.*, vol. 16, pp. 213 – 245, 1995.
- [18] L. Tonin, T. Carlson, R. Leeb, and J. d. R. Millán, "Brain-controlled telepresence robot by motor-disabled people," in *Proc. Annual International Conference of the IEEE Engineering in Medicine and Biology Society EMBC 2011*, 2011, pp. 4227–4230.
- [19] S. ten Hagen and B. Krose, "Trajectory reconstruction for self-localization and map building," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, 2002, pp. 1796 – 1801 vol.2.
- [20] D. Norman, *The Design of Everyday Things*. Doubleday Business, 2002.