

Learning User Habits for Semi-Autonomous Navigation Using Low Throughput Interfaces

Xavier Perrin, Francis Colas,
Cédric Pradalier, Roland Siegwart
Autonomous Systems Laboratory,
ETHZ, Zürich, Switzerland
xavier.perrin@a3.epfl.ch

Ricardo Chavarriaga,
José del R. Millán
Chair in Noninvasive Brain-Machine Interface
EPFL, Lausanne, Switzerland
ricardo.chavarriaga@a3.epfl.ch

Abstract—This paper presents a semi-autonomous navigation strategy aimed at the control of assistive devices (e.g. an intelligent wheelchair) using low throughput interfaces. A mobile robot proposes the most probable action, as analyzed from the environment, to a human user who can either accept or reject the proposition. In case of rejection, the robot will propose another action, until both entities agree on what needs to be done. In a known environment, the system infers the intended goal destination based on the first executed actions. Furthermore, we endowed the system with learning capabilities, so as to learn the user habits depending on contextual information (e.g. time of the day or if a phone rings). This additional knowledge allows the robot to anticipate the user intention and propose appropriate actions, or goal destinations.

Index Terms—Mobile Robot Control, Human-Robot Interaction, Goal Inference, Habit Learning.

I. INTRODUCTION

There is a growing presence of robots in our lives, not only in industrial applications, but also at home as toys and service robots. A particular case focuses on mobile assistive robots, either for user mobility (e.g. an intelligent wheelchair) or telepresence [1]. However, typical input devices like joysticks are still difficult to use by people suffering from a lack of mobility (e.g. due to quadriplegia) or people unable to perform fine movements (e.g. due to Parkinson disease). In this context, dedicated specialized interfaces are required, but they usually offer very limited bandwidth in comparison to joysticks. Examples of such devices include sip-and-puff systems, single switches, or even electroencephalogram-based (EEG) brain-computer interfaces (BCI) [2].

We present a semi-autonomous navigation strategy targeted for low throughput input devices. Instead of requiring the user to send commands continuously, the control is based on human-robot interactions at key locations along the travel. In this approach, the robot evolves autonomously in the environment using sensory inputs. However, as it does not know where the intended destination it queries the user for more information at certain places in the environment. Propositions such as local actions or possible goals are made to the user, who can accept or reject them. Given human rejection, the robot makes new propositions using this additional information. The human workload is thus reduced to only binary inputs (yes/no). When the system has knowledge about the environment, a topological map combined with sensory information is used to identify the

places where a decision is required. Furthermore, contextual information is used to propose either meaningful actions to be performed or directly goal destinations.

We have previously shown how this approach can be used to navigate in both unknown [3] and known environments [4]. In this paper we further test the proposed framework for learning user habits when navigating in known environments. Furthermore, we assessed the performance of the system by emulating a low-throughput, error-prone interface based on the accuracies reported for BCI systems aimed at the detection of brain signals related to error processing [5], [6].

II. RELATED WORKS

Most related works in goal, user, or activity recognition for navigation rely on probabilistic reasoning techniques [7] to represent the different beliefs during navigation. Two main aspect have been studied for *intention recognition*: local intention recognition (immediate action or location in the vicinity of the wheelchair) and global intention recognition (goal destination). In the first case, robots base their inference on direct human input (usually through a joystick), e.g. going left/right/forward, passing through a door, or avoiding an obstacle. Among these, dynamic Bayesian networks (DBNs) have been used to distinguish between approaching or avoiding an obstacle, or to control a wheelchair using either a joystick or brain-decoded commands [2], [8]. Alternatively, Carlson and Demiris [9] computed the probability of trajectories obtained from the relative pose and orientation between a wheelchair and two predefined possible goals.

The inference of global intentions requires the ability to localize the robot within the environment. In a grid-world environment, Partially Observable Markov Decision Process (POMDP) can be enhanced to represent goal locations [10]. After a preliminary phase where reinforcement learning is used to estimate the actions leading to a goal, the system infers the most probable goal from a set of possibly noisy observations. Similar results have been reported by fusing the robot position and the possible goals as the new state definition in a common POMDP [11]. These works aim at recognizing the goal destination based on the actions made along the travel. We want to go further by integrating contextual information closely related to the user habits, such as the time of the day

or the last visited goal. Moreover, we also consider external events, e.g. a ringing phone, that may trigger specific responses from the user, independent from her usual habits.

III. SEMI-AUTONOMOUS NAVIGATION

In the proposed approach the robot navigates autonomously until it reaches a particular location in the environment. At this location, it queries the user to obtain more information about the intended destination. These queries are in the form of potential directions of movement or intended goal locations that should be confirmed or refuted by the user. Assuming that the robot has knowledge about the environment—encoded in a topological map—, at any node in the map we can define the probability distribution over the possible actions $P(A)$. Moreover, we can also infer the intended goal location (based on the probability distribution over goals $P(G)$) based on the first decisions made along the travel. This information can be exploited to improve the system performance by directly proposing to reach the most probable goal location. Furthermore, we can learn the stereotypical user movements—habits—and exploit them in order to anticipate his next desired destination well ahead in order to decrease the human workload.

A. Human-Robot Dialog Management

We rely on a particular human-robot interaction, which consists of robot propositions and related user answers [12]. Similar to early models for human-machine dialogs (e.g. [13]) this type of dialog, based on closed questions, allows the use of low throughput interfaces with binary-like signals. Moreover, we also allow the possibility of using error-prone user interfaces (e.g. a brain-computer interface) by taking into account their estimated reliability. In this case, the true user intention has to be inferred during the whole dialog process, possibly requiring multiple propositions and answers. Inspired by previous work on spoken dialog systems relying on probabilistic reasoning techniques [13], we used a dynamic Bayesian network (DBN) in order to track the user’s intended action or goal destination.

The DBN has three inputs: first, the proposed action or goal; second, the user answer, as provided by the user interface; finally, the known interface accuracy $P(O|I)$ (i.e. the probability of receiving a human input O given the true intention I). The interface accuracy is taken into account to infer the intended user answer which, together with the proposition, allows the update of the probability distribution over the intended actions $P(A)$ or goals $P(G)$. Initially, uniform distributions are used if there is no prior knowledge. However, as we will show in the next sections, the robot is able to build such an initial distribution based on the local topology of the environment.

Finally, a criterion is needed to select which proposition will be made to the user as well as to stop the dialog and proceed to the execution of a command. Based on preliminary tests, we opt for proposing the most probable action or goal, and to use a threshold in the probability difference between the two most probable propositions

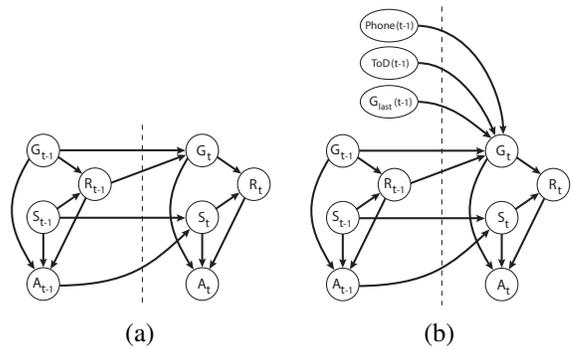


Fig. 1. Dynamic Bayesian network. (a) topological navigation and (b) habit learning. S States, A Actions, G Goals, R goal Reached, G_{last} last visited goal, ToD Time of the Day, and $Phone$.

B. Navigating in known environments

We consider here environments for which the system has a predefined representation, such as a map. In our context, besides the use of a metrical map for robot localization, we rely on a topological map, where nodes represent potential goal locations or places where a navigational decision is required; arcs represent the path between two nodes. In order to achieve low-level control (e.g. left/right action proposition), this topological map is grounded on the metrical map.

We design a second DBN aimed at achieving topological navigation and basic goal inference. Accordingly, we describe the motion in the environment as a Markov decision process with *Actions* A_t , such as *turning left*, *turning right*, *going forward*, or *doing a U-Turn*, and *States* S_t , a state being a robot pose in the environment (see figure 1a for the relation between variables). Note that the motion model $P(S_t|S_{t-1} A_{t-1})$ is deduced from the topological map with a graph-search-based approach. In order to infer the desired goal destination, we introduce a *Goal* variable G_t . Goals are specific locations in the environment coinciding with existing states. Both the state and the goal variables influence the action model $P(A_t|G_t S_t)$, which can also be deduced from the topological map. A boolean variable R_t , corresponding to a *goal reached status* (i.e. $R_t = 1$ if $G_t = S_t$), is added in order to trigger a specific action distribution when reaching a goal destination (i.e. the *Stop* action would be the most probable one), extending thus the action model to $P(A_t|G_t S_t R_t)$. It is also used for $P(G_t|G_{t-1} R_{t-1})$, defining the goal probabilities for the next time step: if a goal is not reached ($R_{t-1} = 0$), the probability of having the same goal in mind is much higher than changing, whereas when the user reaches a goal ($R_{t-1} = 1$), the other places become equally probable. We will refer later to this relation as the *goal persistence*.

The DBN is shown in figure 1a as a graphical model. It is able to infer the most probable actions to do at the next topological node and the intended goal by computing $P(A_t|S_{t-1} A_{t-1})$, resp. $P(G_t|S_{t-1} A_{t-1})$. Furthermore, this model, similar to [10], allows us to incrementally infer goal destinations based on the sequence of observed actions.

In order to represent user habits, the model is extended to

include contextual information (figure 1b). As an example, we present context variables representing the last visited goal, as well as the time of the day and an external event (i.e. phone ringing). The variable G_{last} stores the last visited goal and models the succession of goal visits. Then, the discretization of the day in several time intervals, such as “waking up”, “lunch”, or “evening”, is done by the variable ToD (Time of Day). Finally, the binary variable $Phone$ is an example of an external event taken into account by our DBN: hearing a phone ringing is correlated with the user going to answer the phone, thus breaking up the current sequence execution.

As can be seen in figure 1b, these additional variables influence the goal distribution $P(G_t|G_{t-1} R_{t-1} G_{last} ToD Phone)$. However, this complex probability distribution is difficult to express in such a form. Building upon our previous $P(G_t|G_{t-1} R_{t-1})$ distribution, we will rely on a Bayesian sub-model in order to model the relation between all the variables:

$$P(G_t | G_{t-1} R_{t-1} G_{last} ToD Phone) = P(G_{t-1} R_{t-1}) P(G_t | G_{t-1} R_{t-1}) P(G_{last} ToD Phone | G_t) P(G_{t-1} R_{t-1})$$

is a uniform distribution. $P(G_t | G_{t-1} R_{t-1})$ corresponds to the goal persistence term exposed previously. Finally, the term $P(G_{last} ToD Phone | G_t)$, regrouping all the newly introduced variables, represents the user habits, i.e. the dependence of the goals from the contextual information. It is a Laplace distribution, storing evidences of the user habits in a probability table. For the use of this sub-model in the main DBN, we compute at each time step the distribution $P(G_t | G_{t-1} R_{t-1} G_{last} ToD Phone)$.

The extended DBN is able to infer the user intended goal destination, and consequently the most probable actions leading to it, given the current robot situation and the contextual information. The related distributions calculated from the known variables are $P(G_t | S_{t-1} A_{t-1} G_{last} ToD Phone)$ and $P(A_t | S_{t-1} A_{t-1} G_{last} ToD Phone)$. When reaching a goal location, i.e. when $R_t = 1$, the system will update the distribution on $P(G_{last} ToD Phone | G_t)$ by adding the current observation. The robot is thus learning the habit of its user in order to improve the accuracy of the future inference processes. We will now present experiments for testing the performance of our model.

IV. EXPERIMENTAL RESULTS

A. System performance

We perform simulations in order to assess the performance of the model considering several cases including normal user behavior, as well as unusual user trajectories and its reactions to external events (i.e. phone events). To that purpose, the simulated robot evolves in the environment shown in figure 2, containing five goals and a small loop simulating a wide space. It represents a realistic environment, where ‘A’ could be the bedroom, ‘B’ the bathroom, ‘C’ the kitchen, ‘D’ the main entrance, and ‘P’ the place where a phone is. In order to evaluate the system behavior with error-prone interfaces, in all the experiments we define the interface accuracy $P(O|I)$

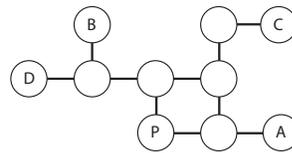


Fig. 2. Topological map of a small indoor environment, with four goals (A-D) and a phone (P).

according to the reported accuracy of BCI system decoding EEG error-related potentials (Table I).

Two scenarios are run, one representing a user habit in the morning and one in the evening. In the morning, the user wakes up in the bedroom, goes to the bathroom, has breakfast, brushes her teeth, and leaves the apartment for work. We thus have the sequence ABCBD. In the evening, she comes back from work, has dinner, watches TV in the bedroom, goes to the bathroom, and finally goes to sleep. We thus have the sequence DCABA. At random across the two scenarios (20%), the phone will be ringing and the user will go to ‘P’ to answer. Also at random, the user will execute an unusual sequence in the environment in 10% of the time.

We run each scenario 100 times, (i.e. representing 100 consecutive days). A second experiment with 1’000 runs of the two scenarios was performed. This allows to compare with the performance when the system has learned the user habits. Two measures μ_{max} and μ_P are introduced, being respectively the average number of times the most probable goal matches the real one, and the geometrical mean of the probability values of the intended goal¹,

$$\mu_{max} = \frac{1}{N} \sum_{t=1}^N \delta(\hat{G}_t, \underset{G_t}{\operatorname{argmax}} P(G_t | S_{t-1} A_{t-1} ToD G_{last} Phone))$$

$$\mu_P = \sqrt[N]{\prod_{t=1}^N P([G_t = \hat{G}_t] | S_{t-1} A_{t-1} ToD G_{last} Phone)}$$

where N is the number of timesteps, \hat{G}_t is the intended goal, and $\delta(a, b)$ is the Kronecker function. The measure μ_{max} , evaluates the efficiency of the model in making a first correct proposition to the user. As propositions get better, less involvement is required from the user. In turn, μ_P measures the certainty of the model on the user habits. Ideally $\mu_{max} \rightarrow 1$, while μ_P would be lower due to the inherent uncertainty.

¹When comparing several models, the common way is to compare their likelihood, which is the product of the probability values. However the likelihood depends on the number of observation which prevents a comparison between datasets of different sizes. The geometrical mean allows to overcome this by normalizing the likelihood per trial.

$P(O I)$	$O = \text{no ErrP}$	$O = \text{ErrP}$
$I = \text{no ErrP}$	92.0%	8.0%
$I = \text{ErrP}$	26.5%	73.5%

TABLE I
CONFUSION MATRIX OF THE ERRP CLASSIFIER [5]. $I \equiv$ USER INTENDED RESPONSE. $O \equiv$ OBSERVED SIGNAL (I.E. EEG DECODED RESPONSE).

Situation	Global	Usual habit	Phone events	Unusual habit
Uniform	0.20	0.20	0.20	0.25
A, μ_{max}	0.85	0.89	0.65	0.55
A, μ_P	0.71	0.77	0.45	0.32
B, μ_{max}	0.85	0.90	0.82	0.48
B, μ_P	0.73	0.83	0.66	0.26

TABLE II
EVALUATION OF THE DBN FOR GOAL INFERENCE. (A) TESTS ON 100 TRIALS. (B) 1'000 TRIALS.

Situation	Global	Usual habit	Phone events	Unusual habit
Uniform	0.20	0.20	0.20	0.20
A, μ_{max}	0.87	0.88	0.70	0.74
A, μ_P	0.74	0.76	0.57	0.56
B, μ_{max}	0.90	0.92	0.69	0.76
B, μ_P	0.74	0.78	0.60	0.51

TABLE III
EVALUATION OF THE DBN FOR ACTION INFERENCE. (A) TESTS ON 100 TRIALS. (B) 1'000 TRIALS.

The system performs better than a random model in all situations (c.f. table II). When the test consisted in 100 trials (upper part of the table), the goals were correctly predicted (μ_{max} , column *Global*) in 85% of the time, or even in 89% if we consider only the usual sequences without any phone events (*Usual habit*). When the phone is ringing, the DBN is more than three times better than chance for inferring the correct goal destination (*Phone events*). The values for μ_P are slightly lower, but also outperform the uniform model. They indicate that the probability distributions are well shaped (rather peaked than uniform), meaning that the system acquired a precise representation of the user habits. Unsurprisingly, the system performance decreases when the user executes unusual sequences (*Unusual habit*). When we run the system for 1000 trials, the performance increases mainly for the μ_P measure indicating that the system has a more precise representation of the user habits. This is particularly true for the phone events, as now the $P(\text{ToD } G_{last} \text{ Phone} | G_t)$ term has overcome the *ToD* and G_{last} dependency. When unusual sequences are executed the system performs close to random.

When considering the action inference (table III, similar μ_{max} and μ_P computed using *A* instead of *G*) the results are similar. We can see that the DBN allows to have a proper initial probability distribution $P(A)$ for engaging a dialog with the human user in order to select an action. Interestingly, the action inference in the case of unusual sequences (*Unusual habit* column) is much better than the corresponding goal inference. This may be explained by the fact that paths towards two different goals share some of the initial actions. Then, at the node where the paths diverge, the selected action will differ from the inferred one, allowing the system to update its inference over goals. Hence, there is only one wrongly inferred action along the path, while there were several steps where the goal was wrongly inferred.

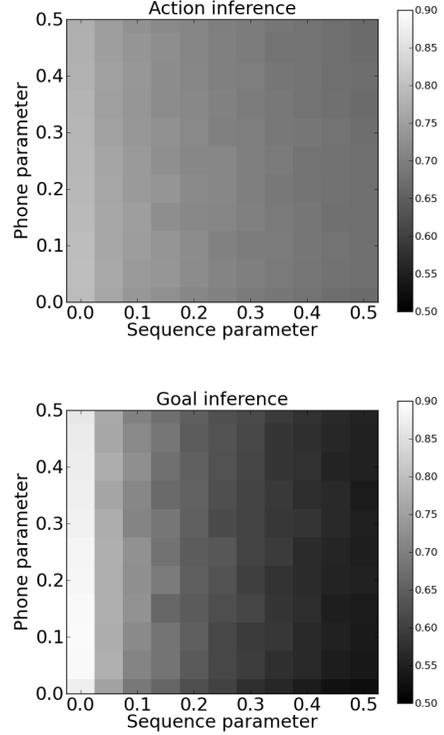


Fig. 3. Influence of unusual events on (*Top*) action and (*Bottom*) goal inference. X-axis: level of unusual sequences. Y-axis: frequency of external events (i.e. phone rings). Plotted values correspond to the geometrical mean of the probability values (μ_P).

B. Influence of unusual events

In order to test the influence of unusual events on the performance of the DBN, we ran experiments where the two parameters controlling the unusual sequences percentage and the frequency of the phone events were varied between 0% and 50%. 1'000 trials were run for each condition. Figure 3 shows the value of μ_P , corresponding to the system confidence about the user habit, for the different cases. It can be observed that external events, i.e. the phone ringing, have little influence (i.e. small variation column-wise). When there is no unusual sequence executed (first column in the graphs), both goal and action inference perform well above random levels. Notice that performance is always below 100% due to the goal persistence term (small probability to change goal). As expected, as the percentage of unusual sequences increases the performance drops, although it remains above random even for the worst-case (i.e. upper right corner of each plot).

We also test the system's adaptation to changes in the user habit. We ran a similar scenario as previously over 100 days, but after 50 days the morning sequence is changed: instead of "having breakfast at home", the user directly goes to work after a stop in the bathroom (i.e. a shorter ABD sequence, instead of ABCBD). The result is displayed in figure 4, where we separated the (a) morning and (b) evening sequence. It represents the variation of μ_P over the different simulation

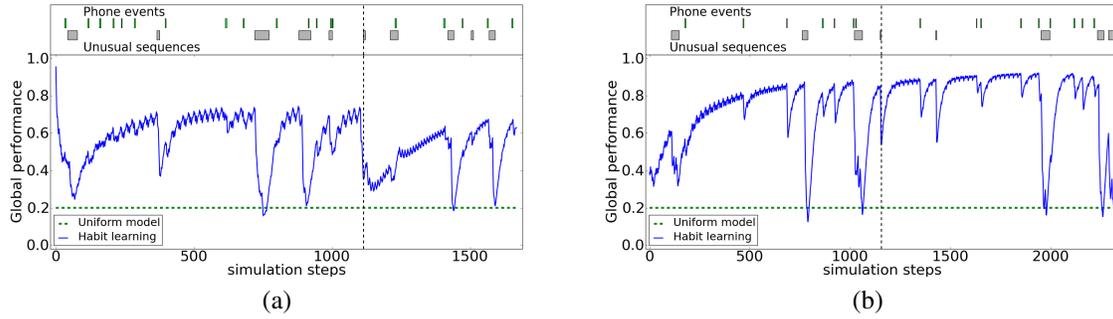


Fig. 4. Model performance (μ_P) when there is a change in the user habits. (a) morning and (b) evening scenarios. The performance, computed with a sliding window, is shown as a function of the simulation steps. The dashed vertical line indicates the moment of the change. On the upper part of each plot, a vertical bar indicates that the phone rings, a gray rectangle indicating the days where the user did not follow her usual routine.

steps, with an exponential window filter² applied for better clarity.³ As can be seen, when the morning habit changes (vertical dotted line), there is at first a drop in performance. However, it does not reach the random level (horizontal dotted line) as might be expected if the sequence would be completely different. But in this particular case, we did not change the whole sequence and kept an already learned part of it. After the habit change, the inference quickly improves again. The evening sequence is not at all influenced by the morning habit's change, as the *ToD* variable successfully separates the two situations. Figure 4 also illustrates the influence of phone events and the changes in the daily habits. The longer the habit change lasts, the greater the impact on the performance is. Of course, the DBN has some difficulties to infer the correct goal destination during these unusual situations, but the time needed to recover from the negative peaks is also due to the filtering technique we applied to the results.

C. Proposing goal destinations

In the previous experiments navigation was achieved through action selection, thus allowing to assess the accuracy of the DBN and its robustness to unusual situations. In order to reduce the user involvement, besides proposing the inferred most probable action, the robot can propose directly a goal destination. In order to decide whether a goal is proposed, we developed a simple heuristic based on a threshold t_G on the goal probability: if the probability of a particular goal exceeds t_G , it will be proposed to the user; if no goal is probable enough, the dialog will be based on actions. When a goal is selected, the robot can autonomously drive to this destination, the user having no more input to provide.

We tested 30 simulated days on a larger scenario (c.f. Figure 5) where the user went to several places (E-D3-D4-P-C-B-D3-D2-D3) with an additional random variation of 20%. We compare different values for the t_G threshold. Figure 6

$${}^2\mu_P(T) = \prod_t P([G_t = \hat{G}_t] | S_{t-1} A_{t-1} ToD G_{last} Phone)^\kappa (1-\kappa)^{(T-t)}, \quad \kappa = \frac{1}{20}$$

(recursively $\mu_P(T) = P([G_t = \hat{G}_t] | S_{t-1} A_{t-1} ToD G_{last} Phone)^\kappa \mu_P(T-1)^{1-\kappa}$)

³If we would have shown the raw data, the variations among one sequence would be dominant on the graph. These variations come from the normal evolution of our DBN, as e.g. all goals are equally probable when we leave a previously reached goal, which leads to a drop in the μ_P measure.

shows the number of questions asked to the user and their nature (i.e. action or goal proposition). It shows that allowing goal propositions decreases the number of user interactions. There is more than a two-fold decrease when $t_G < 0.4$. When $t_G = 0.2$, the numbers of action and goal propositions are similar. This is due to the fact that when leaving the actual goal, the user is first asked whether she wants to do a U-turn (or stay where she is), representing thus one action proposition, before possibly getting a goal proposition.

This approach can be extended to have a hierarchical decomposition of the environment, where goals could be grouped in sub-regions (i.e. *meta-goals*). We tested this using a simple approach consisting of summing together the goal probabilities of each goal in each region, without any normalization on the region size. We thus do not have any a-priori on the meta-goals but focus on the goals. When only a little evidence of the human habits has been accumulated, such a solution is adequate as it favors meta-goals with a higher goal density and reflects the higher probability that the user may go to such regions. A variable MG (with a dependency from the goal G_t) representing the meta-goals is added to the model in Fig. 1b. It allows a new inference process, $P(MG_t | G_{t-1} R_{t-1} G_{last} ToD Phone)$, representing the probability over the meta-goals.

During navigation, if no goal is probable enough, the system proposes a meta-goal that exceed a threshold t_{MG} . If there is none, the most probable action is proposed. When a meta-goal is selected the robot goes to the first node in the meta-goal region and from then on it proposes either actions or goals as described above. Fig. 6b shows the results using this approach.

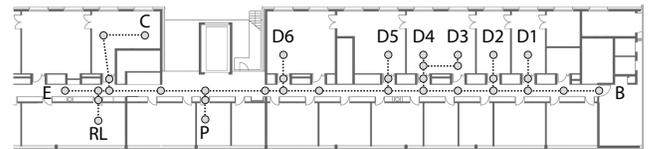


Fig. 5. Topological map of the environment used for the action and goal proposition comparison, with an entrance (E), desks (D1-D6), a printer room (P), a cafeteria (C), a robot lab (RL), and the bathroom (B).

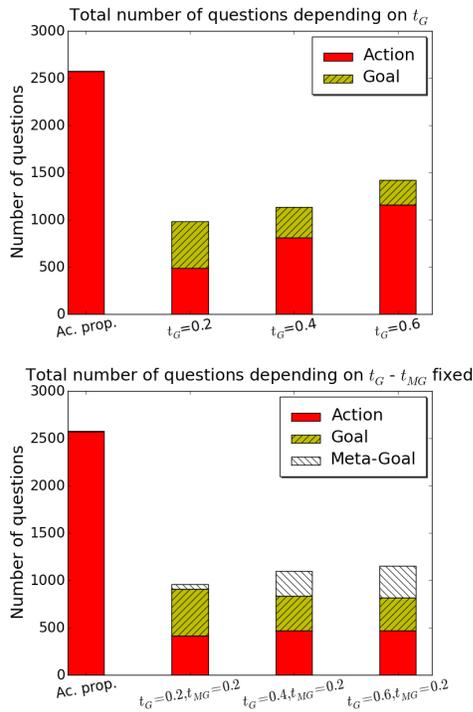


Fig. 6. Propositions made to the user. (Top) Action and goal propositions. (Bottom) Action, goal, and meta-goal (hierarchy) propositions. t_G and t_{MG} : thresholds on the goal, resp. the meta-goal, probabilities.

V. CONCLUSION

We described a semi-autonomous navigation system compatible with low throughput, error-prone interfaces (e.g. brain-machine interfaces). Relying on a human-robot interaction where propositions are made to the user, this system requires low user involvement, as meaningful propositions are made at key locations during the travel. In known environments, experiments showed that the learning of the user habits allows goal directed navigation after just a few interactions. In unknown environments, meaningful action propositions are made thanks to the analysis of the local topology [3].

The proposed approach implies less user effort when compared with other approaches to semi-autonomous navigation. Shared-control systems normally require constant user involvement in order to provide navigational commands [8], [14]; while standard semi-autonomous systems also rely on several input signals for selecting a command through some user interface (such as going up or down in a predefined menu or selecting an item, [15]) or for providing directions to the system (left, right, forward) at a given point in the travel [16].

The human-robot interaction we use, consisting of context-based propositions, requires only a single binary signal and requires needs globally less information from the user. This may be of importance when dealing with low throughput interfaces such as BCIs or alternative specialized interfaces. Other BCI navigation systems have been proposed, either based on a continuous shared control [2], or the P300 paradigm [17]. Compared to the proposed approach, the former case

requires constant involvement from the user, thus requiring a large effort on her part. In turn, the P300 paradigm requires the averaging over several iterations, slowing down the goal selection process.

Previous work using this approach for navigation of a real robot yielded encouraging results using speech or emulated BCI user interfaces [4], [3]. We plan to further test the proposed approach using similar setups. Moreover, the inference system could also be combined with other frameworks of shared control in order to offer different levels of assistance, according to the particular context and the user preferences.

ACKNOWLEDGMENT

Work supported by the Swiss National Science Foundation NCCR-IM2 and by the EC project BACS FP6-IST-027140.

REFERENCES

- [1] L. Tonin, R. Leeb, M. Tavella, S. Perdakis, and J. del R. Millán, "The role of shared-control in BCI-based telepresence," in *IEEE Int Conf Systems Man and Cybernetics*, 2010.
- [2] G. Vanacker, J. Millán, E. Lew, P. W. Ferrez, F. Galán, J. Philips, H. V. Brussel, and M. Nuttin, "Context-Based Filtering for Assisted Brain-Actuated Wheelchair Driving," *Computational Intelligence and Neuroscience*, 2007.
- [3] X. Perrin, R. Chavarriaga, F. Colas, R. Siegwart, and J. Millán, "Brain-coupled interaction for semi-autonomous navigation of an assistive robot," *Robotics and Autonomous Systems*, vol. 58, pp. 1246–1255, 2010.
- [4] X. Perrin, F. Colas, C. Pradalier, and R. Siegwart, "Learning to Identify Users and Predict their Destination in a Robotic Guidance Application," in *Field and Service Robotics (FSR)*, 2009.
- [5] R. Chavarriaga, P. Ferrez, and J. Millán, "To Err Is Human: Learning from Error Potentials in Brain-Computer Interfaces," in *Int Conf Cognitive Neurodynamics (ICCN 2007)*, Shanghai, China, 2007.
- [6] R. Chavarriaga and J. Millán, "Learning from EEG error-related potentials in noninvasive brain-computer interfaces," *IEEE Trans Neural Syst Rehabil Eng*, vol. 18, no. 4, pp. 381–388, Aug 2010.
- [7] S. Carberry, "Techniques for Plan Recognition," *User Modeling and User-Adapted Interaction*, vol. 11, no. 1-2, pp. 31–48, 2001.
- [8] E. Demeester, A. Hütemann, D. Vanhooydonck, G. Vanacker, H. Van Brussel, and M. Nuttin, "User-Adapted Plan Recognition and User-Adapted Shared Control: A Bayesian Approach to Semi-Autonomous Wheelchair Driving," *Auton Robots*, vol. 24, no. 2, pp. 193–211, 2008.
- [9] T. Carlson and Y. Demiris, "Human-Wheelchair Collaboration Through Prediction of Intention and Adaptive Assistance," in *IEEE Int Conf Robotics & Automation (ICRA '08)*, 2008, pp. 3926–3931.
- [10] D. Verma and R. Rao, "Goal-Based Imitation as Probabilistic Inference over Graphical Models," in *NIPS 18*. MIT Press, 2006, pp. 1393–1400.
- [11] T. Taha, J. V. Miró, and G. Dissanayake, "POMDP-Based Long-Term User Intention Prediction for Wheelchair Navigation," in *IEEE Int Conf Robotics & Automation (ICRA '08)*, 2008, pp. 3920–3925.
- [12] X. Perrin, R. Chavarriaga, C. Ray, R. Siegwart, and J. Millán, "A Comparative Psychophysical and EEG Study of Different Feedback Modalities for HRI," in *Human-Robot Interaction 08*, 2008.
- [13] E. Levin, R. Pieraccini, and W. Eckert, "Using Markov Decision Process for Learning Dialogue Strategies," in *Int Conf Acoustics, Speech and Signal Processing (ICASSP)*, 1998, pp. 201–204.
- [14] S. Levine, D. Bell, L. Jaros, R. Simpson, Y. Koren, and J. Borenstein, "The NavChair Assistive Wheelchair Navigation System," *IEEE Trans Rehab Eng*, vol. 7, no. 4, pp. 443–450, 1999.
- [15] H. A. Yanco, "Shared User-Computer Control of a Robotic Wheelchair System," Ph.D. dissertation, MIT, 2000.
- [16] A. Argyros, P. Georgiadis, P. Trahanias, and D. Tsakiris, "Semi-Autonomous Navigation of a Robotic Wheelchair," *J Intelligent and Robotic Systems*, vol. 34, no. 3, pp. 315–329, 2002.
- [17] I. Iturrate, J. Antelis, J. Minguez, and A. Kübler, "Non-Invasive Brain-Actuated Wheelchair Based on a P300 Neurophysiological Protocol and Automated Navigation," *IEEE Trans Robotics*, vol. 25, no. 2, pp. 614–627, 2009.