

First Steps Towards Emotionally Expressive Motion Control For Wheeled Robots*

Steffen Müller, Thanh Q. Trinh and Horst-Michael Gross

Abstract—During social interaction between humans and robots, body language can contribute to communicate mutual emotional states. In this paper, a method is presented, that enables wheeled robots with differential drive to express various emotional states by means of different movement styles during goal-directed motion. The motion control of the robot utilizes an objective-based motion planner optimizing the control commands in a high-dimensional search space by means of an evolutionary algorithm. The main contribution in this paper is an objective function that uses human feedback on the perceived robot emotion in order to evaluate the possible robot control sequences. First group experiments have been conducted in order to demonstrate the system and find suitable movement styles for possible emotional states.

I. INTRODUCTION

In continuation of [1], [2], the aim of our current research project SYMPARTNER (SYMBiosis of PAul and RoboT companion for Emotion sensitive caRe) [3] is the development of a social robot companion that is intended for domestic applications especially in one person households. One focus of the project is the emotion awareness of the system. On the one hand, we try to recognize the moods and emotional states of the user by means of on-board sensing capabilities and incorporation of a smart home installation, on the other hand, the emotion aware design of the system tries to evoke positive feelings in the user. Furthermore, the intended long-term interaction benefits from a strong emotional binding between human and robot. To this end, we give the robot a character with own emotional states, that are to be expressed to the user by means of different modalities. Among animated eyes, head gestures, and voice modulation, the body language of the robot seems to be a promising instrument for that.

In contrast to humanoid robots, our system is designed as a low cost platform and, therefore, has only five degrees of freedom (see Fig. 1). Besides the movable ears and a tilt servo for the head, the robot is equipped with a differential drive consisting of two driven wheels on the front side and a castor wheel in the rear. The maximum velocity of that robot is about 1 m/s. For obstacle and people detection, there are two ASUS Xtion depth cameras (one in the head facing downwards and one in the rear) as well as a Kinect 2 RGB-D camera in the head facing horizontally forward.

*This work has received funding from the German Federal Ministry of Education and Research (BMBF) to the project SYMPARTNER (grant agreement no. 16SV7218).

The authors are with the Neuroinformatics and Cognitive Robotics Lab, Ilmenau University of Technology, Germany www.tu-ilmenau.de/neurob

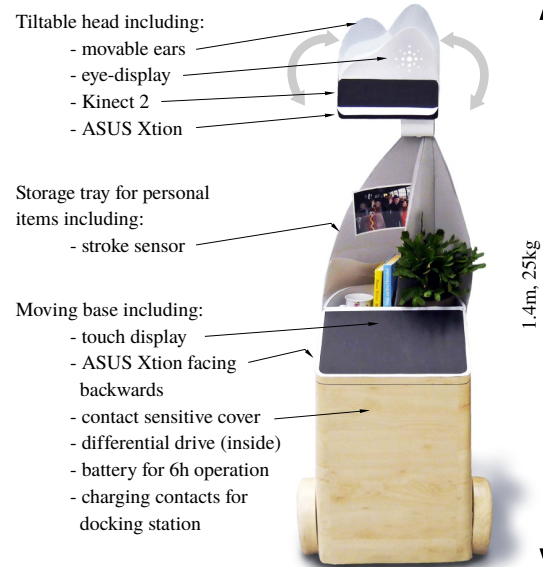


Fig. 1. Robot Alfred with its main sensors, actuators, and interaction devices; conceptual design by Hassenzahl and Welge, University of Siegen; construction and implementation by MetraLabs GmbH Ilmenau

This paper is focussed on the expression of robot internal emotions by means of specific movement patterns that are to be shown during the normal operation of the robot. That means, the already challenging navigation to goals in narrow home environments is to be augmented by different motion styles.

Heider and Simmel [4] already showed that people do associate emotions to abstract objects based on a two dimensional movement pattern. Here relations to the environment and other objects (actors) over time seem to play an important role in addition to the velocity and acceleration profile. Correlation of perceived emotions to shape and velocity of robot movements has also been investigated by Dang et al. [5]. In a Wizard of Oz experiment, they tested round vs. sharp angled shapes of movement trajectories at different velocities.

This shows that the movement patterns should not be disregarded in designing of an emotional robot. Also if mainly other modalities are used for expression of emotions, the actual movement behavior should not evoke conflicting emotions to achieve a natural and consistent impression.

Studies like the aforementioned one have a general drawback. They lack in optimality of the robots movement patterns with respect to the recognizable emotion. Usually,

the effect can only be as good as the wizard controls the robot. For future work in assistive mobile robots, the aim should be to have a system available which can generate movement patterns on its own and gets the feedback on the perceived emotion from observing people in order to improve the own motion patterns.

In order to go in that direction, the developed motion planner can be trained with arbitrary movement patterns. Afterwards, the produced patterns are labeled with an emotional state by human observers. By means of that, for a given or desired emotional state of the robot, the most appropriate pattern can be chosen among the available templates.

The remainder of the paper is structured as follows. First, a coarse sketch of the approach is given followed by some key ideas of the developed local motion planner used in the system. The main part consists of the description of the objective function used for evaluation of the suitability of potential movement trajectories according to the intended emotional impression on human observers. Finally, some results of the user experiments used for training the system are presented.

II. EMOTIONAL NAVIGATION APPROACH

For our purposes, the robot's state of affect is modeled using the ALMA model [6] considering emotions, moods, and personality. Based on external and internal events during interaction and autonomous operation, the state is described in a three-dimensional space of pleasure, arousal, and dominance. That emotional state is used by the arbitrary modalities in order to modulate the robot's expression.

For the modulation of the movement style, the 3D-state is projected to a coordinate in the 2D emotion space (2D-ES), consisting of the dimensions arousal and valence (each ranging from -1 to 1) as proposed by J. A. Russell [7] in order to make it more handy for human spectators, which are an essential component in the approach.

The basic idea is to have a motion planner that is able to propose potential movement commands and internally rate these according to several pre-defined objective functions. This approach is known as Dynamic Window Approach (DWA), which was introduced by Fox et al. in [8] as the first objective-based motion planner for differential drive robots and is now commonly used in the robotic fields. As a result, the selected control command is a compromise that is safe and satisfies various criteria. These criteria, implemented in form of objective functions, comprise a goal orientation, collision avoidance, or social constraints to people in the surroundings. In extension of that approach, we implemented a further objective function considering the emotional impression on potential observers.

We found that the quality of the resulting control commands greatly depends on the complexity of the potential movement trajectories considered by the planner during the search process. In early experiments with our DWA motion planner we found, that the bow like trajectories considered by the DWA are not sufficient for generation of more difficult and expressive motion patterns. For that purpose, we

developed a local motion planner [9] using an evolutionary optimization process for generating proposal trajectories with a high degree of freedom.

Given that framework, the idea of making the movements emotionally expressive can be broken down to an objective function, that is able to rate potential movement trajectories according to the desired emotional impression. For that purpose, we had to find a model that given an emotional state maps trajectories onto a scalar cost value. These costs have to be low for trajectories that match the intended emotional expression and high for those ones that are perceived as a different emotion. At this point, data from experiments with human observers come into play. The data are used to train the model and enable the robot to estimate the impression of trajectories correctly. Initially, the robot does not know how to modify the motion behavior due to missing training data. Therefore, in an initial experiment the robot is manually controlled by different people in order to record a representative set of expressive movement trajectories. The aim here is to cover the whole emotional repertory of the emotion space. Up to now, we do not know either if the robot will be able to reproduce these movement patterns, nor which emotional impression these patterns will have to spectators. Based on this set of template trajectories, the robot can generate movement patterns on its own in a second user experiment. Due to the restrictions of the objective functions and the motion planner these trajectories are, however, slightly different from the manually driven ones. Now a group of human observers provide their impressions as emotion labels in the 2D-ES. These labels now can be used to select an appropriate template from the set of recorded trajectories given a desired emotion to be expressed.

At the current stage of development, the motion planner and a suitable model for the emotional objective function have been implemented, and first user experiments have been conducted showing that the approach is promising. At this point, we are able to generate visibly different motion styles with our robot, which are applied in combination with target directed motion and obstacle avoidance.

III. EVOLUTIONARY MOTION PLANNER

One essential step towards the emotionally expressive movement patterns is the objective-based local motion planner developed for optimizing the control command sequence in a very high-dimensional search space. The motion planner works in an interval of 250 ms and provides the actual velocity target for the robot's motor controllers for that period. In order to correctly estimate the outcome of a command for the next time step, all possible continuations have to be considered up to a certain planning horizon in time. This is of particular importance for objectives, like the emotional evaluation, which are looking at the exact shape rather than checking only a collision free progress in any direction. Other approaches, like the DWA, introduce vast simplifications to that search space by discretizing the command space and reducing the considered continuations

of the sequences to a constant velocity, which results in only bow-like trajectories to be tested.

In order to keep the flexibility of possible command sequences, our approach utilizes an evolutionary optimization process searching the set of complete control sequences rather than only for the next command. For that purpose, we hold a population $P = \{A_i | i = 1, \dots, m\}$ of possible future acceleration sequences as individuals. These individuals $A_i = (a^t | t = 0, \dots, f)$ consist of an acceleration vector $a^t = (a_x^t, a_\phi^t)$ for each time step in the future (cycle time of the motion planner itself) of a window of a few seconds (typically 3 to 4 sec.). a_x is the translation acceleration and a_ϕ is the rotational acceleration (see Fig. 2 for a visualization of that modeling). The acceleration values are limited to the robot's physical capabilities. The length of the planning horizon is limited by the minimum stopping distance but also increases the capability to plan over local minima in the cost function. On the other hand, a longer planning horizon also increases the complexity of the search space drastically. Although, we can benefit from reusing the population in future time steps, real-time constraints are limiting the manageable planning horizon.

An evolutionary algorithm works iteratively for g generations and consists of the following steps: evaluation of the individuals' fitness, selection, and reproduction of the best individuals. Mutation and cross-over during the reproduction ensures an exploration of the search space, and the selection brings the population closer to the optimum. At the end of a planning cycle, the first acceleration command of the best individual is executed and the population is transferred to the next time step (Fig. 2). In the following, these phases are described in more detail.

A planning cycle starts with the **initialization** of the population. Here, we can cope with the weaknesses of the random search in our algorithm, which can not guarantee, that a safe trajectory can be found. In order to enforce consideration of safe trajectories, one part of the initial population is generated deterministically, consisting of stopping trajectories of different length and direction. The remaining part is reusing knowledge from the last cycle. The best sequences from the preceding population are transferred into the current time step by cutting off the first command and filling up with $a^f = (0, 0)$, resulting in a shift in time. To take into account the deviation of the real robot velocity to the planned acceleration, the new first element a^0 is corrected by the deviation of the actual velocity $v_{odometry}$ to the planned one divided by the cycle interval Δt . $a^0 = a_{old}^1 + (v_{old}^1 - v_{odometry}) / \Delta t$

Once a population has been set up, the loop starts with the **fitness evaluation** for the individuals. Thereto, by means of a forward model of the robot's physics, the acceleration sequence A_i can be converted to a velocity sequence $V_i = (v^t | t = 0, \dots, f)$ with $v^t = (v_x^t, v_\phi^t)$. These velocity sequences in turn are limited to the robot's maximum speed parameters and integrated to a movement path in world coordinates.

The resulting movement trajectories are then used for

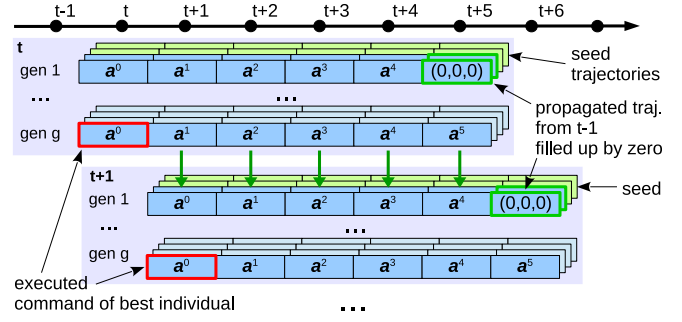


Fig. 2. Encoding of individuals as sequences of acceleration vectors which are optimized for g generations in each planner cycle, propagation of the final population of one planner cycle to the next helps reusing the optimization effort from past cycles and reduces number of generations needed in each planner cycle.

evaluating the fitness of the individuals. This is done by the set of objective functions, each yielding a cost value or a hard deny (in case of potential collisions). These individual cost values are summed up and can be compared to the costs of other individuals later on. In our setup, the following objective-functions are used:

- 1) A **path and heading objective** is responsible for navigating towards a given target position in world coordinates. This is done by evaluating a globally planned navigation function (using E* planner [10]) at the trajectory's end point and additionally evaluating the deviation of the orientation at the end point to the wanted goal orientation in proximity to the goal position. The navigation function is replanned periodically on a cost-map that incorporates the current obstacle situation as recognized by the robot's distance sensors.
- 2) A **distance objective** for avoiding collisions with static and dynamic obstacles. Trajectories are checked considering the robot's exact footprint for collisions in a static occupancy map as well as in the local map containing the sensed obstacles.
- 3) A **direction objective** preferring forward motion of the robot to account for the limited sensor capabilities in the rear.
- 4) A **personal space objective** to keep distance to people in the close proximity of the robot. For detected people, their future movements are predicted using a linear motion model, and the distances to the robot's potential movement path along the resulting trajectories is evaluated using a Gaussian-shaped model of the personal space.
- 5) The **emotional objective** that is comparing the estimated impression on the human observers to the current emotional state of the robot. Further details on this follow in the next section.

Based on the cost value, the individuals of the population are sorted in ascending order. This allows to prefer the best individuals over weaker ones during the **generation of offsprings** for the next population. The next generation is built up from the n best individuals of the last one

(**selection**), and the remaining individuals are generated each by combining two randomly drawn individuals from the last generation. The index in the sorted list thereby is drawn from a normal distribution with mean 0 and a fixed variance, which defines the selection pressure. The \mathbf{a}^t of the two parent individuals are mixed into one new sequence, while the value is taken either from the first or the second parent randomly. Last step is mutation of the \mathbf{a}^t values by means of a normal distributed random offset. After that, the loop starts over with the new generation of individuals, while over all n generations the best rated sequence is stored for execution of its first command at the end of the planning cycle.

The evolutionary motion planner has already proven its better performance compared to the DWA in a social navigation scenario. See [9] for comparative results.

IV. EMOTIONAL OBJECTIVE FUNCTION

In order to make the selection of the individuals depending on the intended emotional impression on human spectators, we developed a new objective function taking the velocity sequences V_i of the unrolled individuals and a desired emotional state in the 2D-ES and mapping it to a cost value.

The emotional impression of a robot's movement trajectory depends on various factors of influence. Besides the actual velocity profile (translation and rotation), distances and orientations to people in the surroundings seem to have an influence. Moving frontally towards people may evoke association to a curious mood, while avoiding people can support the impression of fear. Also the geometry of the environment and the robot's position in the environment may make a difference. If there is a free space, and the robot moves close to the walls, it may appear shy compared to a movement in the center of the room. Also the moving direction in relation to the robot's destination has an influence. Very goal-oriented movements indicate attention and curiosity, while undirected patterns can express boredom.

All these factors make a data-driven model more complex. When designing a model, one has to consider the amount of data that is available for training. If the model is too complex, it might not be able to generalize in unseen situations. Due to this, we decided to concentrate on the velocity profiles only. Spatial relations to people and obstacles are planned for a second version of the approach, when more training data will be available.

Therefore, as a data base for the model, we have a set $T = \{(V_j, \{E_j^p\})\}$ of labeled template velocity profiles covering about 10m of robot movement each. $V_j = (\mathbf{v}^t | t = 0, \dots, t_{max})$ are the velocity sequences and $E_j^p \in 2D-ES$ are the labeled emotional impressions of observer p , when the robot performed the pattern V_j . Fig. 3 shows the velocity profiles of three exemplary templates on the right side, while the emotional labels $\{E_j^p\}$ are illustrated in the emotion space on the left.

An additional essential aspect is the efficiency for evaluating the model. The motion planner runs at 4 Hz, and in each cycle it processes 5 generations with 60 individuals each. This results in 1200 fitness evaluations per second. This is

the same amount of trajectories as the DWA has to evaluate, if the 2d velocity space is discretized in 15x20 bins.

Furthermore, in contrast to the other objective functions it is necessary to consider also a historic context when deciding on the suitability of future trajectories. For example, this is essential to continue a sinusoidal movement pattern in the correct direction rather than starting a new sinus pattern in an arbitrary direction in each time step. To that end, we added a history part $(\mathbf{v}^{-h}, \dots, \mathbf{v}^{-1})$ to the velocity profile V_i , before the voting takes place. This \mathbf{v} consists of the real odometry values measured in the past. The length h has been chosen between 2 and 4 seconds, which corresponds to 8 to 16 velocity samples in the interval of the planner cycle.

The first idea for realizing the cost function was to train a predictive model that maps the input velocity profiles V_i into the 2D-ES. Then the distance in the 2D-ES to the desired emotion yields the costs we are looking for – high costs for undesired emotional impressions and low ones for a good match. Attempts for realizing such a mapping based on the training trajectories generated during the user experiments was unsuccessful, since the space of possible trajectories is so huge that the little amount of training data is not representative. So, the evolutionary optimization in the planner each time found completely different trajectories that were mapped exactly onto the desired emotion state and, therefore, got low costs. However, the resulting movement patterns did not look like the training data at all.

The actual realization of the objective function, in contrast to the mapping attempt mentioned before, operates the other way around. First, the emotional labels E_j^p are used to find a set of suitable template trajectories in the training data. Then, the actual velocity profile to be rated is matched against these patterns, and low costs are generated if the trajectory matches the training profiles at any position. If no matching position can be found, the costs returned are high.

In more details, for each template sequence j the density of the labels $\{E_j^p\}$ at the current desired emotion E_d is evaluated (step one in Fig. 3). A high density of labels indicates a good match of the j -th template in emotion space. These matches in the second step are converted into a template specific prior cost value

$$C_j = 1 - \frac{1}{|p|} \sum_p e^{-(E_j^p - E_d)^2 / \sigma} \quad (1)$$

used as a constant cost offset for the matching process (cyan bars in Fig. 3).

After this computation, which has to be done only once per planner cycle, the input velocity profile V_i is shifted along the templates (step three in Fig. 3), and a distance to the velocities in the templates is computed and added as costs to the C_j (cyan curves in the diagrams of Fig. 3). Over all templates T_j and all time offsets in the templates, the minimum yields the final costs

$$C = \min_{j,t} \left\{ C_j + \left(1 - \frac{1}{h+f} \sum_{w=-h}^f e^{-|\mathbf{v}_j^{t+w} - \mathbf{v}_i^w|^2 / \rho} \right) \right\} \quad (2)$$

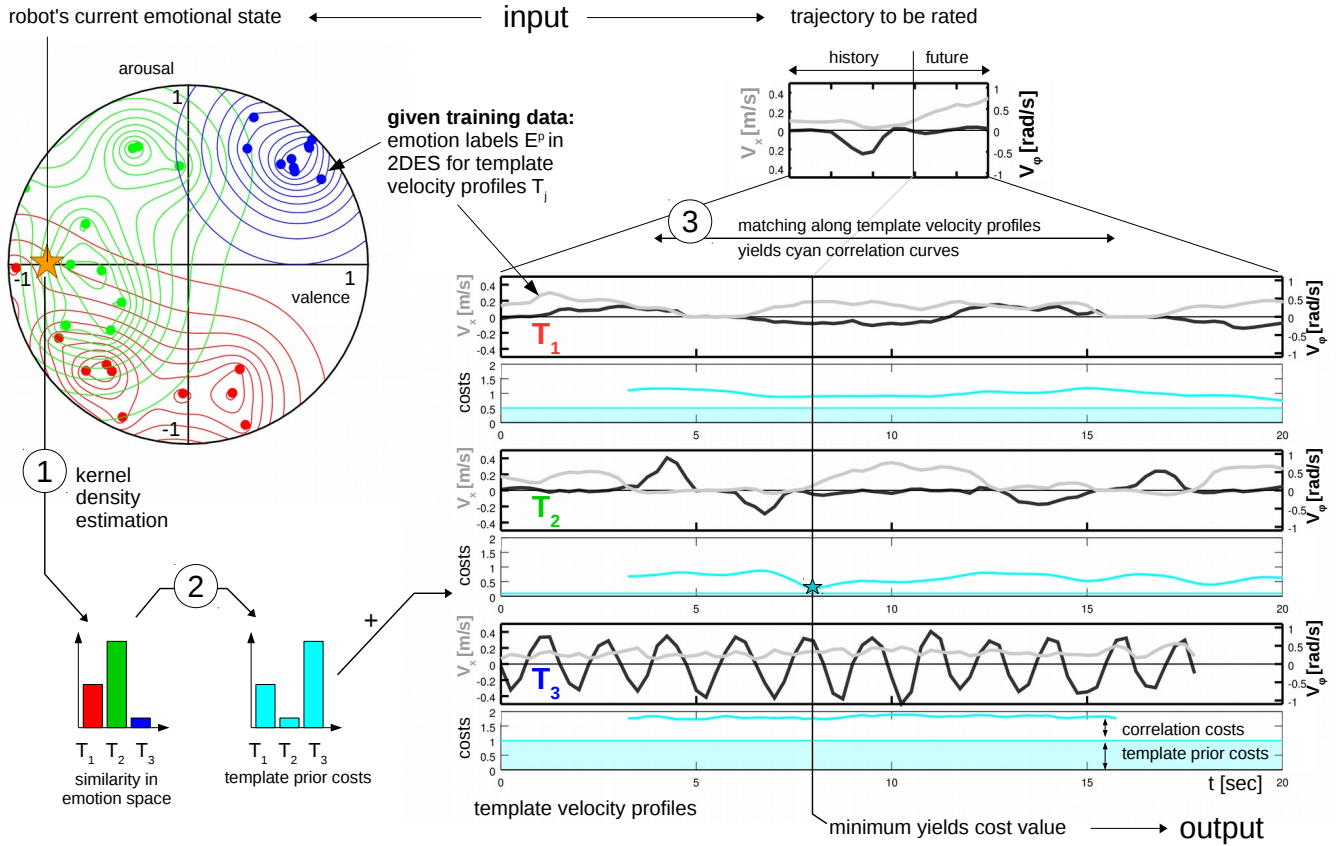


Fig. 3. Illustration of cost computation in the emotional objective function, for a given proposal trajectory and a desired emotional expression (input). First, the suitability of the templates (T_1 red, T_2 green, and T_3 blue) in the emotion space is evaluated. The colored samples in the 2D-ES on the left are the labels given by human observers during the training experiment. The sample density in a second step leads to template-specific prior costs, which are added to correlation results during step three. The training velocity profiles on the right (only three are shown) are used to compute the costs by matching the input velocity profiles to the templates resulting in the cyan curves. The global minimum of these over all time offsets and templates yields the output. T_1 is a slow, curved forward motion with occasional hesitation (perceived as fear or boredom), T_2 is a slow motion with stops to look left and right (fear and anger), T_3 is forward motion in narrow sine like pattern (excited, happy).

for the input trajectory with ρ defining the sensitivity for deviations in the velocity space. In order to speed up the evaluation, dynamic programming is used for the minimum search. Additionally, the velocity samples from the history of the trajectory have to be matched only once per planner cycle since they are the same for all proposal trajectories coming from the evolutionary motion planner. Therefore, increasing the context does not increase the computational effort drastically.

In future work, we plan to further investigate the forward mapping approach, which failed so far, due to the bad generalization properties in the high-dimensional input space. For consideration of additional context information, like spatial relations to people and objects along the movement path, the template matching approach used so far seems to be unusable. We plan to apply artificial neural networks as a predictive model, which maps the input trajectory and the context information into the 2D-ES. The approach of Leibig and Wahl [11] is promising, since it enables a network to know if the input is close to the training data or not, which helps to identify those inputs for which the network still generalizes insufficiently.

V. HUMAN LABELING EXPERIMENTS

As already mentioned, the proposed emotional navigation system is based on a dataset of emotional labels for distinct movement templates presented by the robot to human spectators.

Two different group experiments have been conducted in order to get these data. Initially, there are no distinct movement patterns that could be autonomously performed by the robot. Therefore, a first group experiment with 12 participants provided a spectrum of possible movement patterns and additionally aimed to prove the existence of a common sense of emotional interpretation among the participants. For recording the manually driven template trajectories, an easy to use control mechanism is required, since steering the robot with the keyboard is not that handy and only useful for smooth straight trajectories. A remote control by means of a 13 cm small model of the robot has been implemented for that purpose. The model can be moved on a small glass table and is tracked by means of a webcam capturing a QR-code like pattern from below. The velocity observations of that model are sent to the real robot, where

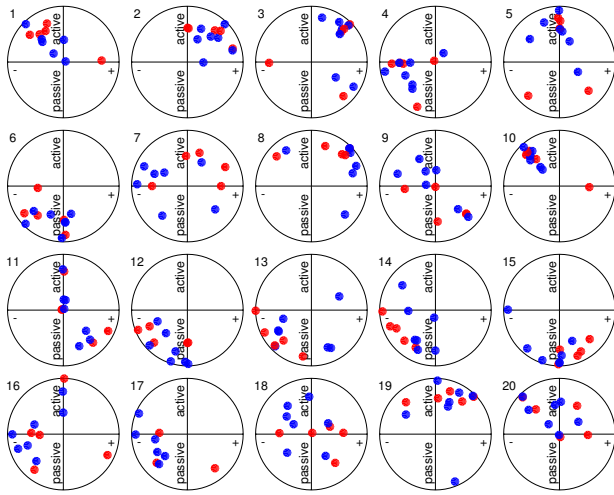


Fig. 4. Emotion labels of the autonomous navigation experiment (blue - male, red - female). The robot performed emotional movement patterns while targeting a goal in the hallway of our institute building while 11 participants observed and rated their impression in the 2D-ES.

the velocity commands are directly relayed to the motor controllers resulting in a relatively direct replication of the model robot's movements.

During the experiment, in each case, one of the participants controlled the robot along a hallway of our institute building, while the remaining audience observed it and had to note their impression by means of a cross in the 2D emotion circle on a form. During the runs, the actual movement trajectory of the robot was recorded. The group consisted of staff members and students of our lab in age between 24 and 50 with 5 female and 7 male participants.

The visual analysis of the label data showed, that there are patterns that make a similar impression on the observers, but others are ambivalent (also visible in the data of the second experiment shown in Fig. 4). Since the "puppeteer" had a certain emotion in mind for each trial, for most of the parts of the 2D-ES a significant pattern could be generated at the end. Only the lower right quadrant representing the relaxed mood is not covered sufficiently.

After the velocity patterns were recorded, the robot was able to perform the movement patterns autonomously. Therefore, the recorded velocity profiles have been trimmed manually to only contain the desired pattern going in a straight direction, and the emotional objective function was used as described above but only one template T_j was active at a time. Finally, 19 patterns have been generated using the earlier recorded velocity profiles.

In a second experiment, the participants had to label the perceived emotion again while the robot drove autonomously, resulting in the data shown in Fig. 4. During trial 20, the emotional objective function was disabled as a reference, which results in the "normal" navigation behavior. The label data shows that the impression is slightly biased towards the active side of the 2D-ES in that particular case, which is due to the faster and goal-directed motion compared to the patterns with the activated emotion objective.

An analysis of the emotion labels showed, that the autonomously generated movement patterns are of a similar quality as the manually driven patterns with respect to perceivable emotions. A comparison of the average standard deviation of the emotion labels given by the observers for manually driven trajectories (0.471) and the autonomously driven (0.493) shows that there is no significant difference in the interpretability of the patterns. The concentration of labels in different sectors of the 2D-ES shows that there is a common sense of emotional interpretation of movements in the group as supposed in the beginning.

VI. CONCLUSIONS

We could show that it is possible to realize an autonomous motion controller for a mobile robot, that is able to combine safe, goal directed navigation with emotionally expressive movement patterns. In order to overcome arbitrariness in the design of emotional expressive patterns, we suggested to involve opinions of a group of people voting the performed movement patterns directly in a real-valued emotion space and demonstrated the feasibility of such an approach. Open issues to be considered in future developments are alternative models for the emotional objective function also taking into account further context information (e.g. spatial distances to users and obstacles). Nevertheless, the approach of using only the velocity profiles without any context information showed a good performance in our user experiments.

REFERENCES

- [1] C. Schröter, S. Müller, M. Volkhardt, E. Einhorn, C. Huijnen, H. van den Heuvel, A. van Berlo, A. Bley, and H.-M. Gross, "Realization and user evaluation of a companion robot for people with mild cognitive impairments," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2013, pp. 1145–1151.
- [2] H.-M. Gross, S. Müller, C. Schröter, M. Volkhardt, A. Scheidig, K. Debes, K. Richter, and N. Döring, "Robot companion for domestic health assistance: Implementation, test and case study under everyday conditions in private apartments," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 5992–5999.
- [3] Sympartner project page. [Online]. Available: <http://www.sympartner.de>
- [4] F. Heider and M. Simmel, "An experimental study of apparent behavior," *The American Journal of Psychology*, vol. 57, no. 2, pp. 243–259, 1944.
- [5] T. H. H. Dang, G. Hutzler, and P. Hoppenot, "Mobile robot emotion expression with motion based on mace-grace model," in *2011 15th International Conference on Advanced Robotics (ICAR)*, June 2011, pp. 137–142.
- [6] P. Gebhard, "Alma: A layered model of affect," in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM, 2005, pp. 29–36.
- [7] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, pp. 1161–1178, Dec. 1980.
- [8] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics Automation Magazine*, vol. 4, no. 1, pp. 23–33, Mar 1997.
- [9] S. Müller, T. Q. Trinh, and H.-M. Gross, "Local real-time motion planning using evolutionary optimization," in *Proceedings of Towards Autonomous Robotic Systems (TAROS)*, 2017.
- [10] R. Philippsen and R. Siegwart, "An interpolated dynamic navigation function," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*. IEEE, 2005, pp. 3782–3789.
- [11] C. Leibig and S. Wahl, "Discriminative bayesian neural networks know what they do not know," in *NIPS Workshop: Deep Learning and Representation Learning*, 2016.