

# Enhancing Human-Robot Interaction by a Robot Face with Facial Expressions and Synchronized Lip Movements

Viktor Seib, Julian Giesen, Dominik Grüntjens, Dietrich Paulus  
Active Vision Group, AGAS

Institute for Computational Visualistics  
University of Koblenz and Landau, Germany

{vseib, jgiesen, dominik.gruentjens, paulus}@uni-koblenz.de  
<http://robots.uni-koblenz.de>

## ABSTRACT

With service robots becoming increasingly elaborate for higher level tasks, human-robot interaction is moving into the focus of robotic research. In this paper we present an animated robot face as a convenient way of interacting with robots. Our robot face can show 7 different facial expression, thus providing a robot with the ability to express emotions. This capability is crucial for robots to be accepted as everyday companions in domestic environments. Aiming towards a more realistic interaction experience our robot face moves its lips synchronously to the synthesized speech. In a broad user study with 100 subjects we test the emotions conveyed by the robot face. The results indicate that our robot face enhances human robot interaction by providing the robot with the ability to express emotions. The presented robot face is highly customizable. It is available for ROS and can be used with any robot that integrates ROS in its architecture. Further information is available at <http://ros.org/wiki/agas-ros-pkg>.

## Keywords

Robot Face, Talking Head, Animated Dialogue System, Human-Robot Interaction, ROS

## 1 INTRODUCTION

In recent years robots have found their ways into many homes around the world. As for now, most of these robots are household appliances that were designed to perform one specific task: they are able to vacuum or wipe the floor or to mow the lawn. Nevertheless, the popularity of these, single task specific, robots shows that people are willing to accept robots in their everyday lives.

Therefore, current research focuses on further improving the autonomy and generality of robots. One of the goals in mind are general purpose service robots for domestic tasks. The benefits of having such elaborate helpers at home are manifold. Not only would they take over annoying and tedious household chores, but they could also assist disabled or elderly people in helping them with their daily needs. Especially the last-mentioned aspect is becoming more important in our aging society.

These new application areas require for novel means of communication between man and machine. While it is sufficient to interact with a cleaning robot by pushing buttons on the robot itself or an a remote control, robots strongly integrated in a person's daily routine are expected to understand gestures, speech, and even facial expressions. Likewise, the robot itself has to communicate in a human-like manner using the same means of expressing itself. Since humans focus on faces when communicating with one another, a face also increases the chance of a robot to be accepted as an equal communication partner by a human. A recent psychological study shows that robots exhibiting human-like features are even ascribed more intelligence than robots with less human-likeness [Kra08].

In this paper we present an abstract, cartoon-like, animated robot face for human-robot interaction. While our robot face system possesses only the most important facial features it is able to show 7 essential face expressions that are crucial for human-robot interaction. Additionally, a text-to-speech system is used to synthesize speech by passing arbitrary input strings. The mouth moves according to the synchronized speech and thus provides an even more realistic interaction experience. All animations are generated dynamically during runtime by interpolating between previously defined shape keys. Our animated robot face is available as a package for the widely spread robotics middleware

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.



Figure 1: Robot heads designed in hardware: (a) Robot head “Kismet”, Breazeal et al. [Bre99, Bre03], (b) Emotional-display “EDDIE”, Sosnowski et al. [Sos06], (c) Cartoon-like robot head “Flobi”, Lütkebohle et al. [Lüt10], (d) Head of general-purpose social robot “Bender”, Ruiz-del-Solar et al. [Rui09].

ROS [Qui09]. It can be downloaded<sup>1</sup> and easily used on any robot equipped with a display and running a ROS-capable architecture. As it is completely designed in software, the `robot_face` is easily customizable. It is even possible to replace the whole face model by a different one without losing any of the features described in this paper. To our knowledge this is the first easy to use animated robot face that every one can adapt and integrate into an existing robot.

The next Section describes related work and design concepts in some specific aspects that distinguish our animated robot face. The actual implementation is presented in Section 3. Section 4 describes the evaluation procedure of our robot face, followed by a discussion of the results in Section 5. Finally, Section 6 concludes with a summary and an outlook to future work.

## 2 RELATED WORK AND DESIGN CONCEPTS

Different talking heads were developed in the last years for research in the field of human-robot interaction. Kismet, a robot head demonstrating facial expressions is presented in [Bre99, Bre03]. It expresses emotions by moving its facial features like eyes, mouth and ears. A more recent approach, the emotional-display EDDIE [Sos06], uses the facial action coding system (FACS) [Ekm77] to depict emotions. By defining *action units*, i.e. smallest movable units, FACS describes the movements of most facial muscles and their effect on the face expression. In contrast to these two approaches, Flobi [Lüt10] was designed as a cartoon-like robot head with humanoid features. Its design completely hides the interior mechanics. Another recent approach is Bender [Rui09], which is also able to show emotions. Ruiz-del-Solar et al. conducted a study to evaluate the effect of Bender’s emotion on humans interacting with it. We compare the results of this study with the results of our own study in Chapter 4. The here mentioned robot heads are presented in Figure 1.

The robot heads of these systems are constructed in hardware, posing a challenge in designing and build-

ing these heads. Also, the costs of the different components needed might be an issue. A strong advantage, however, is the possibility to place cameras inside the head’s eyes. This allows for intuitive interaction in a way that a person can show an object to the robot by holding it in front of the robot’s head.

Although this is not possible with a face completely designed in software, we chose this approach to create our animated robot face. In our opinion the high number of advantages of an animated head outweighs its drawbacks. There is no specific hardware that needs to be added to the robot. Thus, there are no additional expenses arising from using our robot face. Moreover, it is highly customizable and can be adjusted to everyone’s individual needs. Finally, the ROS interface allows for comfortable and easy integration in existing systems.

### 2.1 Cartoon-like Appearance and Abstraction

When focusing on animated faces two main approaches can be distinguished. Human-like or even photorealistic faces are employed to convey realism and authenticity to the interacting person. On the other hand the purpose of stylized cartoon faces is to invoke empathy and emotions. Often this is achieved by exaggerated facial expressions or unrealistic proportions of eyes, mouth or other facial features.

Since our robot (like most of robots participating at the RoboCup@Home) lacks humanoid features and stature, a realistic human face is not appropriate to interact with it. Instead, we modeled an abstract cartoon face exhibiting only the most important facial features to express emotions: eyes, eyebrows, and a mouth. A second reason for the choice of a cartoon face is to avoid the risk of falling into the *uncanny valley*. According to [Mor70], the familiarity of a robot (or a doll, etc.) increases with human likeness. However, when reaching a certain point of high similarity even slight differences from natural appearance cause an uncomfortable effect in the observer. Moving entities augment the similarity with humans, but also the uncomfortable effect. We therefore aimed at creating an animated face that is able to convey familiar face expressions and emotions, but at the same time is not

<sup>1</sup> Package `robot_face` on <http://ros.org/wiki/agas-ros-pkg>




Figure 2: Animated text-to-speech systems: (a) August Dialogue System, Gustafson et al. [Gus99], (b) Facial Animation System, Albrecht et al. [Alb02], (c) Text-to-audio-visual Speech, Niswar et al. [Nis09].




Figure 3: Visemes of our robot face ((a) through (f)) and different shapes of the eyebrows ((g) through (i)).

realistic, i.e. human-like, enough to create an uncanny effect.

## 2.2 Lip Movement and Speech Synthesis

A key feature of a robot face designed for interaction is the ability to speak. We use a text-to-speech system *Festival*<sup>2</sup> for speech synthesis. Festival synthesizes speech by applying phonetic and linguistic rules to the input character sequence. To provide an effect of authenticity to the interacting person the lip movements have to be synchronized and animated according to the spoken words of the robot's face. The FACS [Ekm77] is not well suited for this purpose since it does not include the lower face part. We achieve this synchroniza-

tion by mapping visemes to phonemes of the synthesized text. Visemes are visually distinguishable shapes of the mouth and lips that are necessary to produce certain sounds. Phonemes are groups of similar, but not identical sounds that feel alike for the speaker. There are phonemes that produce the same viseme and some that do not alter the shape of the mouth at all. Therefore, only a few visemes are sufficient to achieve a realistic animation of the lips (Figures 3a through 3f). Several animated robot heads were developed in the recent years that possess this skill. Some examples from [Gus99, Alb02, Nis09] and are shown in Figure 2. In contrast to our approach, these animated heads were designed with the goal of modeling a realistic and human-like appearance. To our knowledge non of them was used to interact with a robot.

## 2.3 Expressing Emotions

Moving the mouth and lips is not enough to allow for comfortable interaction. The movements have to affect the whole face in order to make it appear vivid. A face capable of expressing emotions is crucial for a robot to be accepted as an equivalent communication partner. The face expressions of our robot face are depicted in Figures 5a and 5b.

Animated movies and video games often use animations created manually since the spoken text is known a priori. However, for our purpose only dynamically generated animations came into consideration, as we want to animate arbitrary text with the desired face expression. Apart from visemes we defined shape keys containing several different configurations for the eyes and eyebrows (Figure 3g through 3i).

## 3 ANIMATED ROBOT FACE

We have developed a talking head application for human-robot interaction named `robot_face`. The talking head performs synchronized lip movements with spoken language and shows 6 different emotions and a neutral face expression. Our goal was to create an application easy to use with robots and to have

<sup>2</sup> <http://www.cstr.ed.ac.uk/projects/festival/>

the possibility to customize the face. As an example for customization we provide two faces with different genders. In addition, the voice's gender, face color, iris color, and outline colors of the face can be adjusted to the needs of the individual user. With some restrictions, a completely different face can be designed with Blender and used with our application. Please refer to the `robot_face` wiki on the project's website for more information.

To accomplish this application, we used Ogre3d<sup>3</sup> as graphics engine for visualisation, Qt<sup>4</sup> as window manager, and Blender<sup>5</sup> for creating the Meshes. As mentioned before, Festival is used for speech synthesis and ROS has been chosen to allow for easy integration of our robot face with any robot using ROS.

### 3.1 Face Modelling and Animation

We designed two similar, cartoon-like faces (a male and a female one) for the presented robot face. Both faces were designed with Blender including a mouth for speaking, eyes for blinking, and eyebrows to intensify emotions. The difference between both faces are distinctive eyelashes on the female face and thicker eyebrows on the male face, as well as a different eye color.

Since we modelled our faces with Blender we used polygon models and adapted them with subdivision surface methods. According to [Par02] subdivision surfaces are a good modelling type for cartoon-like faces. We used the modelling method introduced by Jason Osipa [Osi03], where the model is created by hand and which is an excellent way to model a cartoonish face. According to this method, the mouth and eye areas are modelled separately and are connected afterwards. As we need a mouth for automatically generated animations, we modelled it slightly different than described by Osipa. Focus was put on animation during the modelling process. Thus, we created shape keys for all different face movements and emotions. An overview is given in Figure 3.

For mouth movements we limited the number to the four most important visemes namely mouth open, closed, wide, and narrow. With those four visemes, it is possible to create two clearly separated speech cycles: open and close movements together with wide and narrow movements. It is not necessary that both speech cycles are executed at the same time nor do they have to blend from one extreme into the other [Osi03].

Open and close movements occur by almost any sound as opposed to wide and narrow movements which are associated with the art of sound. There are about 38 to 45 phonemes in the English language, but only a few

<sup>3</sup> <http://www.ogre3d.org/>

<sup>4</sup> <http://qt.nokia.com/>

<sup>5</sup> <http://www.blender.org/>




Figure 4: Components and interaction of the robot face.

visual counterparts. Thus, we combined the indistinguishable phonemes into one appropriate viseme.

Beside those visemes we have designed other mouth shape keys for emotional representation. We used 6 different emotions namely happy, sad, angry, surprised, scared, disgusted, and also a neutral expression. These emotions are shown in Figures 5a and 5b. Additionally, we added shape keys for eyebrows (up, middle up, middle down). These are shown in Figure 3. To achieve movement, the shape keys are interpolated in our developed application with the use of Ogre3D.

### 3.2 Structure of `robot_face`

Our `robot_face` application consists of two ROS-nodes. The `TalkingHead` node manages both the mesh and the animation. To get even better feedback on what the robot says it also displays the spoken text under the robot face. Furthermore, emoticons that are used to specify the robot's face expression are removed from the displayed text. The creation of phonetic features including speech and voice is handled by the `FestivalSynthesizer` node. An overview is given in Figure 4.

We use the messaging system of ROS to communicate with `robot_face`. In order to do this, a string needs to be published on a specific ROS topic. It is directly delivered to the application where it gets synthesized, animated, as well as displayed. In detail, if a given text is sent via the message system to `robot_face` it arrives at the two ROS-nodes `TalkingHead` and `FestivalSynthesizer`. The `TalkingHead` displays the text for the duration of the animation. It is also capable of displaying additional information (i.e. robot state, recognized speech) published as string to a different topic.

`FestivalSynthesizer` synthesizes the speech. It generates phonemes and speech corresponding to the provided text using Festival. We use PulseAudio<sup>6</sup> as sound system for audio output. Apart from the phonemes corresponding timestamps are generated by the `FestivalSynthesizer` node. This information is used by the node `TalkingHead` for animation.

<sup>6</sup> <http://www.pulseaudio.org/>

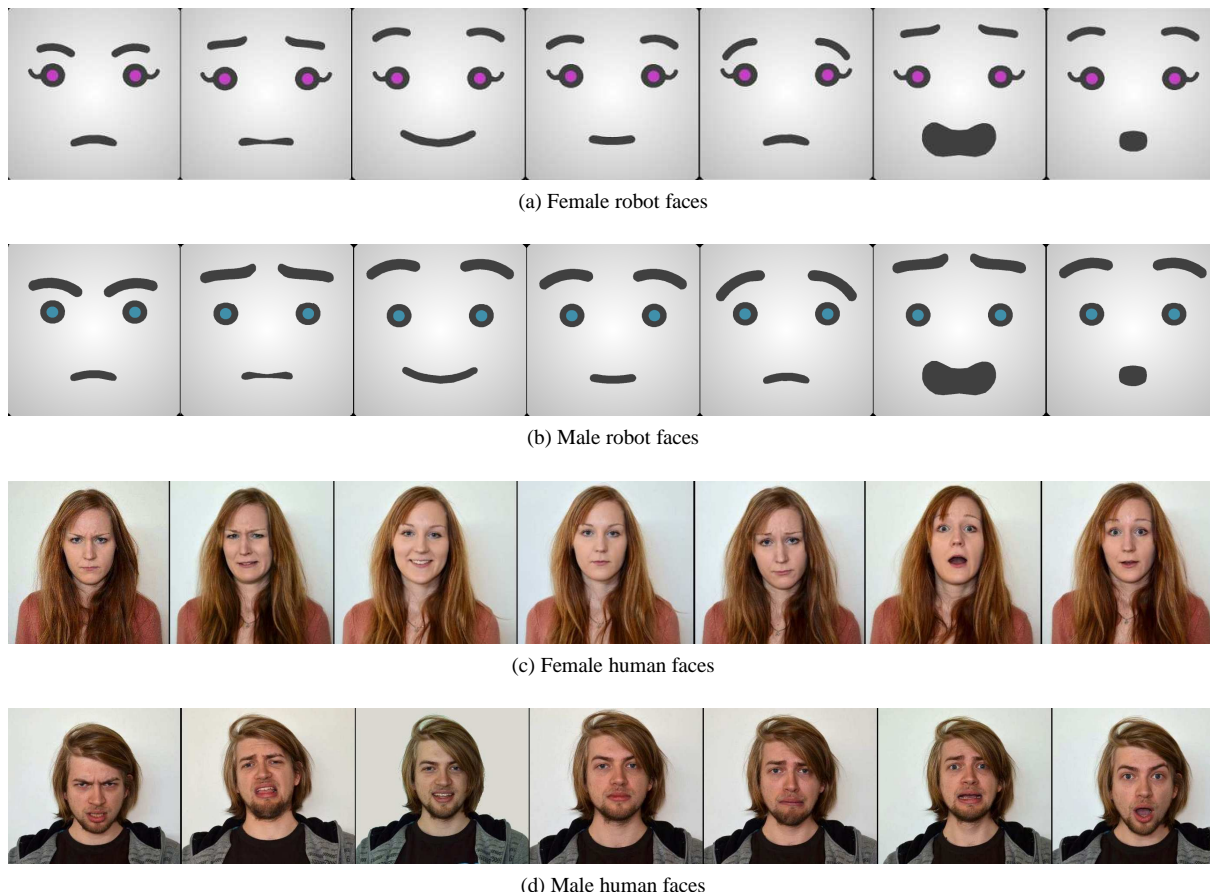


Figure 5: Face expressions that can be displayed by our robot face and the corresponding face expressions of our human models for evaluation (from left to right): angry, disgusted, happy, neutral, sad, frightened, and surprised.

In the TalkingHead node the face mesh is animated by Ogre3D. The main structure of TalkingHead is organized into the creation of the scene, creation of animation, and play-back of animation.

The submeshes of the loaded mesh are counted and the same number of animations is created. These animations need to get filled with keyframes to represent movement. By default, incidental blinking and wiggle animations are active. Keyframes are generated with the phonemes and timestamps mentioned before. We build a predefined phoneme-viseme-map to associate phonemes with visemes. A keyframe is generated for every viseme and emotion using the timestamps. The keyframes are then connected to a whole animation. As soon as the animation starts the spoken text is displayed below the robot's face.

#### 4 EVALUATION

Similar to the evaluation presented in [Rui09], we evaluated the presented robot face to determine how the intended face expressions are perceived by people and whether the intended emotions could be conveyed. Further, we tested how comfortable people were when looking at the developed robot face. The results of both evaluations are compared and discussed in Section 5.

The evaluation was performed as an online questionnaire. The test was divided into two parts, each having 14 questions. In the first part the test persons were presented all 7 face expression of our robot face (Figure 5a) and a photo of a human face expressing one of these emotions (one of the photos in Figure 5c). The probands had to select the robot face that best matched the face expression of the human. Although the presented human face always was intended to show one of the displayed robot faces, the test subjects also had the possibility to select *unknown* and thus skip the question if they could not decide. This test was performed once for each of the 7 face expression in Figure 5c, each time with a different photo. Subsequently, all 7 questions were repeated in a different order with a robot face depicting a male face (Figure 5b) and photos of a male human (one of the photos in Figure 5d). In this part of the test no adjectives describing or naming any of the face expressions were involved.

In the second part the probands were presented one of the robot faces and had to select from a list with 14 adjectives which described the displayed face best. The 14 adjectives contained the 7 available expressions, 6 expressions that were not depicted by the robot face,

Table 1: Results of the first part of the evaluation. Each line represents a photo of a human face with the indicated expression. The numbers show which robot faces were matched to the displayed photo (in percent). Matches above 10% are printed in bold, the maximum of each line is marked gray.

	matched with	angry	disgusted	happy	neutral	sad	frightened	surprised	unknown
angry	<b>85</b>	8	0	0.5	0.5	1	0	5	
disgusted	6.5	<b>34.5</b>	0	1	7.5	<b>26.5</b>	2	<b>22</b>	
happy	0.5	0	<b>76.5</b>	<b>19</b>	0	1	0	3	
neutral	4.5	3	0.5	<b>87.5</b>	2.5	0	0.5	1.5	
sad	1	<b>75</b>	0.5	1.5	<b>21</b>	0.5	0	0.5	
frightened	0.5	5.5	0	0	3.5	<b>76</b>	<b>12.5</b>	2	
surprised	0	2	0.5	5	1.5	<b>12</b>	<b>77</b>	2	

Table 2: Results of the second part of the evaluation split in two halves. The upper half contains presented face expressions, while the lower part contains face expressions that were not shown to the test subjects. Each line represents the robot face with the indicated expression. The numbers show which expression was matched to the displayed robot face (in percent). Matches above 10% are printed in bold, the maximum of each expression assigned is marked gray.

	identified as	angry	disgusted	happy	neutral	sad	frightened	surprised
angry	<b>81.5</b>	0.5	0	0	8	0	0.5	
disgusted	2	1.5	0	0	<b>19.5</b>	1	0	
happy	0	0	<b>94</b>	2.5	0	0	0.5	
neutral	0	0	3.5	<b>88.5</b>	1	0	0.5	
sad	0.5	0	0	0	<b>87.5</b>	0.5	0	
frightened	3	8	0.5	0	0	<b>70.5</b>	6	
surprised	0	0	0	0	0	4.5	<b>90.5</b>	

	identified as	anxious	tired	bashful	bored	arrogant	hurt	none of these
angry	0	0	0.5	1.5	1	6	0.5	
disgusted	<b>39.5</b>	1	<b>12</b>	1	0.5	<b>19.5</b>	2.5	
happy	0	0	0	0	0	0	3	
neutral	0.5	1	0.5	2	0	0	2.5	
sad	2	2	1.5	1.5	0	4.5	0	
frightened	7	0.5	0.5	0	0.5	1	2.5	
surprised	0.5	0	0	0	0	4	0.5	

and the option *none of these*. Again, this was tested for each of the 7 robot face expressions, first with female then with male robot faces.

A total of 100 persons (62 male, 38 female) aged between 19 and 58 years (average 26.2 years) participated in our evaluation. To 53 persons the face of our robot was unknown before the evaluation. 34 people stated to have seen the face before, but to have never interacted

with the robot. The remaining 13 persons knew the face and also had interacted with the robot.

## 5 RESULTS AND DISCUSSION

Each part of the evaluation was performed with male and female faces (either human or robot). The results of both genders were averaged for each part of the evaluation and are presented in Table 1 for the first part and

Table 3: Comfort of the test subjects when looking at the robot faces presented in the evaluation.

very uncomfortable	uncomfortable	undecided	comfortable	very comfortable
2 %	5 %	38 %	46 %	9 %

in Table 2 for the second part. Each line in Table 1 represents a photo of a human face showing the face expression indicated in the first column. Accordingly, every line in Table 2 stands for a robot face with the given expression. The numbers are percentage values and indicate which robot faces were matched to the displayed photo (Table 1) or the expression the robot face was identified as (Table 2). Every case above 10 % is printed in bold, the maximum of each line is marked gray. Ideally, the diagonal would show 100 % at each position in Table 1 and in the first half of Table 2.

Most of the elements in the diagonal of Table 1 have high values: 5 have values of over 75 % and 2 of them have 85 % or more. Only 2 of 7 photos were not matched well with the provided robot face expressions. This is a strong indication for the fact that the key facial features of our robot face are able to recreate the face expressions of humans correctly. The misclassifications in the first part can also result from misclassification of the presented human face. Thus, in the second part of the evaluation no human faces were presented to the probands. The diagonal of Table 2 has 6 elements with more than 70 %, 3 of these have more than 80 %, and the other 2 even over 90 % identification rates. When the robot faces are evaluated on their own without being compared to human faces, only 1 of 7 does not match the intended expression.

In Table 1 the expressions *angry* and *neutral* have the best matches and were not falsely related to other robot faces (i.e. no other columns with 10 % or above). Table 2 confirms this findings. Thus, these two face expressions can be classified well on their own and even pass the comparison with a human photograph.

The *happy* photo was matched correctly with the corresponding robot face in 76.5 % of cases. However, almost every fifth proband assigned the neutral robot face to this photo. Comparing this result to Table 2 shows on the other hand that the *happy* robot face has the highest correct classification result of 94 %. Thus, the high misclassification rate when directly compared to a human photo stems from the human face expression and not from the robot face.

A look at the expressions *frightened* and *surprised* shows a duality in Table 1. Both have very similar correct matches, but were at the same time misclassified with one another - again with very similar rates. Table 2 shows again that this error must result from the human face expression on the photo since the robot faces were misclassified with a significantly lower rate.

The expressions *disgusted* and *sad* have bad matching results in the first part of the evaluation. When presented on its own, the *sad* robot face has excellent classification results (Table 2). However, the *sad* human photo was mostly matched with the robot face that shows a *disgusted* face. Thus, while the *sad* robot face is indeed perceived as sad the *disgusted* robot face seems to resemble better the features of sad human faces. On the other hand, the *disgusted* photo was matched to the correct robot face in only 34.5 %. Over one fourth of all test subjects matched it with the *frightened* robot face. Further, the high number of probands that selected *unknown* indicate that non of our robot face expressions can resemble the features of disgusted human faces. This findings are confirmed by the results in Table 2 where almost no correct identifications for the *disgusted* face are present (only 1.5 %). The *disgusted* robot face was mostly classified as *anxious* (39.5 %), *sad* (19.5 %), *hurt* (19.5 %) or *bashful* (12 %). The various maxima in the classification of this robot face show that it is difficult to identify and to be assigned a feeling to. However, considering that sad and hurt are similar expressions, it can be stated that the *disgusted* robot face resembles an anxious or a sad face expression.

In contrast to Bender [Rui09], who can show 4 different face expressions, our robot face can show 7. Compared to the results of the evaluation of Bender, our robot face achieves higher recognition rates by the test subjects. The highest difference occurs with the *happy* face expression, where our application was recognized correctly in 94 % of cases (compared to 51 % of Bender). The other 3 face expressions compare as follow (results for Bender given in brackets): surprised 90.5 % (76.5 %), sad 87.5 % (78.4 %), and angry 81.5 % (76.5 %). One needs to take into account that Bender is a hardware robot head and looks more technically compared to our cartoonish animated robot face. It is obvious that designing a robot head in hardware with several facial expressions is more challenging than in software.

Apart from the classification of the presented face expressions the test subjects were asked to rate their comfort when looking at the robot's faces. The results are shown in Table 3. While only 7 % of the probands experience discomfort, 55 % feel comfortable when looking at the presented robot face. Although, the number of undecided test subjects is high the results indicate that our robot face does not fall into the *uncanny valley*.

## 6 CONCLUSION AND FUTURE WORK

We presented an animated robot face that is able to show 7 different face expressions and whose lips are synchronized to the synthesized speech. This robot face is highly customizable and can be used with any robot running ROS.

An evaluation with 100 test subjects shows that 5 of 7 robot faces were correctly assigned to a presented human face in 80% (average) of all cases. Also, 6 of 7 robot face expressions are classified correctly in 85% on average. This is a strong indication that our robot face enhances human robot interaction by providing the robot with the ability to express emotions. Compared to a similar evaluation of a state-of-the-art robot face in hardware, the presented approach performs significantly better in a user study.

The only face expression not classified correctly by most users was the face expression that we intended to show disgust. According to the results of the user study this expression conveys a mixture of anxiety and sadness and thus should be used accordingly.

The evaluation also shows that most probands (55%) feel comfortable when looking at the robot face, while 38% are undecided. This and the reason that it is a cartoon face leads to the assumption that it does not fall into the uncanny valley, although more investigation in this area is desirable.

Our future work will concentrate on improving the ability of our robot face to express emotions. For instance, the appearance of the robot's eyes can be changed depending on the presented emotion. Also, a new face expression for disgust needs to be found as the current one will be used as anxiety and sadness in the future. Further, with the fact in mind that our robot face can express emotions as is shown by the presented evaluation, we want to evaluate whether it can invoke empathy in humans interacting with a robot that is equipped with the presented robot face.

## 7 ACKNOWLEDGEMENTS

The authors would like to thank Alruna Veith and Lubosz Sarnecki for posing as face models for the evaluation photos.

## 8 REFERENCES

- [Alb02] Albrecht, I., Haber, J., Kahler, K., Schroder, M., and Seidel, H.P.; May I talk to you?:-) - facial animation from text. In *Computer Graphics and Applications*, 2002. Proceedings. 10th Pacific Conference on, pages 77-86. IEEE, 2002.
- [Bre03] Breazeal, C.; Toward sociable robots. *Robotics and Autonomous Systems*, 42(3):167-175, 2003.
- [Bre99] Breazeal, C. and Scassellati, B.; How to build robots that make friends and influence people. In *Intelligent Robots and Systems, 1999. IROS'99. Proceedings. 1999 IEEE/RSJ International Conference on*, volume 2, pages 858-863. IEEE, 1999.
- [Ekm77] Ekman, P. and Friesen, W.V.; Facial action coding system. 1977.
- [Gus99] Gustafson, J., Lundeberg, M., and Liljen-crants, J.; Experiences from the development of August - a multi-modal spoken dialogue system. In *ESCA Workshop on Interactive Dialogue in Multi-Modal Systems (IDS-99)*, 1999.
- [Kra08] Krach, S., Hegel, F., Wrede, B., Sagerer, G., Binkofski, F., and Kircher, T.; Can Machines Think? Interaction and Perspective Taking with Robots Investigated via fMRI. *PLoS ONE*, 3(7), 2008.
- [Lüt10] Lütkebohle, I., Hegel, F., Schulz, S., Hackel, M., Wrede, B., Wachsmuth, S., and Sagerer, G.; The Bielefeld Anthropomorphic Robot Head Flobi. In *2010 IEEE International Conference on Robotics and Automation*, Anchorage, Alaska, 5 2010. IEEE, IEEE.
- [Mor70] Mori, M.; Bukimi no tani [The uncanny valley]. 1970.
- [Nis09] Niswar, A. and Ong, E.P. and Nguyen, H.T. and Huang, Z.; Real-time 3D talking head from a synthetic viseme dataset. In *Proceedings of the 8th International Conference on Virtual Reality Continuum and its Applications in Industry*, pages 29-33. ACM, 2009.
- [Osi03] Osipa, J.; *Stop Staring - Facial Modeling and Animation Done Right™*. Sybex, 2003.
- [Par02] Parent R.; *Computer Animation - Algorithm and Techniques*. The Morgan Kaufmann Series in Computer Graphics and Geometric Modeling. Academic Press, 2002.
- [Qui09] Quigley, M., Gerkey, B., Conley, K., Faust, J., Foote, T., Leibs, J., Berger, E., Wheeler, R. and Ng, A.; Ros: an open- source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.
- [Rui09] Ruiz-del-Solar, J., Mascaró, M., Correa, M., Bernuy, F., Riquelme, R., and Verschae, R.; Analyzing the human-robot interaction abilities of a general-purpose social robot in different naturalistic environments. In *Lecture Notes in Computer Science (RoboCup Symposium 2009)*, volume 5949 of LNCS, 2009.
- [Sos06] Sosnowski, S., Bittermann, A., Kuhnlenz, K. and Buss, M.; Design and evaluation of emotion-display eddie. In *Intelligent Robots and Systems, 2006 IEEE/RSJ Int. Conf. on*, 2006.