

Automatic Facial Gesturing for Conversational Agents and Avatars

Goranka Zoric*, Karlo Smid** and Igor S. Pandzic*

* Department of Telecommunications
 Faculty of Electrical Engineering and Computing, University of Zagreb
 Unska 3, HR-10000 Zagreb, Croatia

** Ericsson Nikola Tesla
 Krapinska 45, HR-10002 Zagreb, Croatia
 {Goranka.Zoric, Igor.Pandzic}@fer.hr, karlo.smid@ericsson.com

Abstract— We present two methods for automatic facial gesturing of graphically embodied animated agents. In one case, conversational agent is driven by speech in automatic Lip Sync process. By analyzing speech input, lip movements are determined from the speech signal. Another method provides virtual speaker capable of reading plain English text and rendering it in a form of speech accompanied by the appropriate facial gestures. Proposed statistical model for generating virtual speaker’s facial gestures, can be also applied as addition to lip synchronization process in order to obtain speech driven facial gesturing. In this case statistical model will be triggered with the input speech prosody instead of lexical analysis of the input text.

I. INTRODUCTION

Conversational Agent is a graphically embodied animated agent (a virtual character) capable of human-like behavior, most importantly talking and natural-looking facial gesturing. Virtual characters representing users in the virtual realty world are often called avatars.

Human speech is bimodal in nature [1]. Speech that is perceived by a person depends not only on acoustic cues, but also on visual cues such as lip movements or facial expressions. The human face can express lots of information, such as emotions, intension or general condition of the person. In noisy environments, visual component of speech can compensate for a possible loss in speech signal. This combination of auditory and visual speech recognition is more accurate than auditory only or visual only. Use of multiple sources generally enhances speech perception and understanding. Consequently, there has been a large amount of research on incorporating bimodality of speech into human-computer interaction interfaces.

In this article, we concentrate on two ways for automatic facial animation – Lip sync and Visual Text-to-Speech (VTTS).

The goal is to animate the face of a Conversational Agent in such a way that it realistically pronounces the given text. For a realistic result, lip movements must be perfectly synchronized with the audio. In order to appear realistic, the produced face animation must also include facial gesturing. A very important goal of our work is to perform this gesturing automatically. We have achieved this in case of VTTS through statistical modeling of

behavior and plan to extend this method to lip sync as well.

Lip sync produces lip movements synchronized with the input speech. The most important issue in Lip Sync research is Audio to Visual (AV) mapping which consists of the speech analysis and classification in visual representatives of the speech (Figure 1). We present a new method for mapping natural speech to lip shape animation in real time. The speech signal, represented by MFCC vectors, is classified into viseme classes using neural networks. The topology of neural networks is automatically configured using genetic algorithms. This eliminates the need for tedious manual neural network design by trial and error and considerably improves the viseme classification results. This method is suitable for real-time and offline applications.




Figure 1: Schematic view of lip sync system

Classical Visual Text-to-Speech (VTTS), produces lip movements synchronized with the synthesized speech based on timed phonemes generated by speech synthesis ([9],[10],[11],[12]).




Figure 2: Training and content production processes.

We propose a new approach, Autonomous Speaker Agent, that combines the lexical analysis of input text with

the statistical model describing frequencies and amplitudes of facial gestures. The statistical model is obtained by analyzing a training data set consisting of several speakers recorded on video and stenographs of their speech. A lexical analysis of the stenograph texts allowed to correlate the lexical characteristics of a text with the corresponding facial gestures and to incorporate this correlation into a statistical model. Using a lexical analysis of input text to trigger this statistical model, a virtual speaker can perform gestures that are not only dynamically correct, but also correspond to the underlying text. Figure 2 depicts training and content production processes.

The Section 2 and Section 3 describe respectively the automatic lip sync system and autonomous speaker agent in more details. The paper closes with a conclusion and a discussion of the future work.

II. AUTOMATIC LIP SYNC

Speech sound is produced by vibration of the vocal cords and then it is additionally modeled by vocal tract. A phoneme, defined as basic unit of acoustic speech, is determined by vocal tract, while intonation characteristics (pitch, amplitude, voiced/whispered quality) are dependent on the sound source.

Lip synchronization is the determination of the motion of the mouth and tongue during speech [2]. To make lip synchronization possible, position of the mouth and tongue must be related to characteristics of the speech signal. Positions of the mouth and tongue are functions of the phoneme and are independent of intonation characteristics of speech.

There are many acoustic sounds that are visually ambiguous. Therefore, there is a many-to-one mapping between phonemes and visemes, where viseme is a visual representation of phoneme [3].

The process of automatic lip sync consists of two main parts (Figure 1). The first one, audio to visual mapping, is key issue in bimodal speech processing. In this first phase speech input is analyzed and classified into viseme categories. In the second part, calculated visemes are used for animation of virtual character's face.

There are many algorithms that can be modified to perform AV mapping – Hidden Markov Model (HMM), Neural Networks (NN), Gaussian Mixture Model (GMM), Vector Quantization (VQ), etc.

A. Our lip sync system

Our system for automatic lip sync is suitable for real-time and offline applications. It is speaker independent and multilingual. We use visemes as the main classification target. Speech is first segmented into frames. For each frame most probable viseme is determined. Classification of speech in viseme classes is performed by neural networks. Then MPEG-4 compliant facial animation is produced.

Next sections present the components of the system.

1) Phoneme database

As training data, a set of phonemes is collected. For every phoneme, three different samples were recorded by nine test subjects. This gives 27 versions of each phoneme in our database. These phonemes are manually mapped onto MPEG-4 visemes, and in doing so the database is

organized in 14 classes, each corresponding to one MPEG-4 viseme. On average, each viseme class is represented by 50 samples in the database.

For fine tuning of animation, phonemes specific for certain language might be added in the database.

Accuracy of the speech classification depends greatly on the quality and the size of the recorded database.

2) Audio to Visual Mapping

In order to synchronize the lips of a computer generated face with the speech, speech must be first preprocessed and then classified into visemes.

a) Speech Analysis

The Mel-Frequency Cepstrum Coefficients (MFCC) representation of the speech is chosen as first step in preprocessing the speech.

MFCC is audio feature extraction technique which extracts parameters from speech similar to ones that are used by humans for hearing speech, while, at the same time, deemphasizes all other information. As MFCCs take into consideration the characteristics of the human auditory system, they are commonly used in the automatic speech recognition [4].

Additionally, Fisher linear discriminant transformation (FLDT) is done on MFCC vectors to separate classes. If there is no separation between classes before FLDT, transformation will not enhance separability, whereas if there is only slight distinction between classes, the FLDT will separate them satisfactory.

The overall procedure of the coefficients calculation is shown on Figure 3.



Figure 3: Audio preprocessing used

In order to use MFCCs on the speech signal, frame length and the dimension of the MFCC vectors must be determined. The frame length must be chosen, so that the frame contains enough information [4]. The choice is frame length of 256 samples and 12 dimensional MFCC vector. Overlapping of the frames is used to smooth transition from frame to frame (Figure 4).

The phoneme database is now used as a training set in order to train neural network, as it will be described later.




Figure 4: Overlapping of the frames

b) Lip Shape generator

Every frame of the speech is classified in the correct viseme class. When correct viseme is chosen, it can be sent to animated face model.

To reduce sudden discontinuous facial expressions, the outputs from neural network for four consecutive frames are analyzed [4]. One with the largest output sum is chosen as a correct viseme.

3) MPEG-4 Face Animation

Face animation (FA) is supported in MPEG-4 standard [3]. MPEG-4 FA specifies a face model in its neutral state, a number of feature points (FPs) and a set of Facial Animation Parameters (FAPs). Each FAP corresponds to a particular facial action deforming a face model in its neutral state. The first group of FAPs contains high-level parameters, visemes and expressions. Only 14 static visemes are included in the standard set.

Facial animation can be generated for any parameterized face model for speech animation if the visemes are known.

B. Training neural networks for av mapping using GA

Neural networks (NNs) are widely used for mapping between the acoustic speech and the appropriate visual speech movements [5]. Many parameters, such as weights, topology, learning algorithm, training data, transfer function and others can be controlled in neural network. A major unanswered question in NN research is how best to set a series of configuration parameters so as to maximize the network's performance. Finding a suitable network for a specific problem is proved to be difficult.

As training neural network is an optimization process where the error function of a network is minimized, genetic algorithms can be used to search optimal combination of parameters.

Genetic algorithms (GA) are a method for solving optimization or search problems inspired by biological processes of inheritance, mutation, natural selection and genetic crossover. A conventional GA consists of coding of the optimization problem and set of the operators applied on the set of possible solutions [6]. The algorithm's three main operators are selection, crossover and mutation. Only individuals that are good enough get the possibility to survive.

GAs might be used to help design neural networks by determining [7]:

- *The weights.* Algorithms for setting the weights by learning from presented input/output examples with given fixed topology often get stuck in local minima. GAs avoid this by considering many points in the search space simultaneously.
- *Topology* (number of hidden layers, number of nodes in each layer and connectivity). Determining a good/optimal topology is even more difficult – most often, an appropriate structure is created by intuition and time consuming trial and error.
- *A suitable learning rule.*

However, they are generally not used in all three problems at the same time, since they are computationally very expensive.

1) Training NNs

In our approach, we use multilayer feedforward networks to map speech to lip movements. These kind of neural networks are widely used and operate in a way that an input layer of nodes is projected onto output layer through a number of hidden layers. Backpropagation algorithm is used as training algorithm for adjusting weights.

After experimenting with different number of networks, we have decided to train 15 networks, each for every viseme class. The reason is the following. A single

network tested on the entire training data gave poor results, computation time was long and recognition was not satisfactory [4]. Then the number of networks was chosen to be the same as the number of phonemes. The results were satisfactory, but it seemed like a good idea to reduce number of networks, if possible, since training of every network requested extra time.

Therefore, we took advantage of idea that phonemes that are visual ambiguous, do not need to be separated, since it does not influence animation.

The 12-dimensional MFCC vectors are used as inputs to 15 different networks. For each viseme class, a NN with 12 inputs, a number of hidden nodes and 1 output is trained. The number of hidden layers and the number of nodes per each layer should have been determined for each network. This is laborious and time consuming work since the training session must be run until the result is satisfactory. In order to avoid time consuming trial and error method, we introduced simple genetic algorithm to help find suitable topology for our NNs.

2) GA and NNs in our approach

Since the design of neural network is optimized for a specific application, we had to find suitable network for our lip sync application. As determining a good or optimal topology is even the most difficult task in design of NN, we tried to solve this problem with GAs.

In our example, given the learning rule (Levenberg-Marquardt), we used GA for training a backpropagation feedforward network to determine near optimal network topology, including the number of hidden layers and the number of units within each layer.

We use simple genetic algorithm [8], where number of genes specify the number of hidden layers (n). Gene maximum and minimum values are defined in the range from zero to m , determining the number of nodes per layer. If a value of the single gene is set to zero, the number hidden layers is decreased, so practically it ranges from zero to n . Other parameters that have to be specified are population size, maximum number of generation and mutation rate.

Experiments have shown that it is not necessary to have more than three hidden layers ($n = 3$). From the same reason, maximum number of nodes per layer is set to 30 ($m = 30$). Such configuration of GA seems suitable since larger network increases computation time, but does not give better results.

By using genetic algorithms, the process of designing neural network is automated. Once the problem to be solved is coded and GA parameters are determined, the whole process is automated. Although it is still a time consuming work, much time is saved since all work is done automatically by computer.

C. implementation

Constructing database and creation of 15 neural networks have to be done only once. In training process, network's biases and weights are extracted and saved for later use. Together with Fisher matrix (obtained by calculating FLDT), biases and weights matrix are loaded in the application.

Figure 5 shows GUI (Graphical User Interface) of our application.

Application captures speech from the microphone and segments it into frames of 256 samples. When a frame has been captured, data is stored and calculations are performed during capturing of the next frame. These calculations consist of MFCC extraction and simulation of 15 networks. The outputs are added to outputs from the previous frame. Every fourth frame, the viseme class that has the largest sum of output values from NNs is presented on the screen [4]. It is important that calculation time does not exceed time needed for recording of a frame (in case of 16 kHz, 16 bit coding takes 16 ms).




Figure 5: GUI of our application

III. AUTONOMOUS SPEAKER AGENT

A conversation consists of two domains: verbal and nonverbal. These two domains are highly synchronized because they are driven by the same forces: the prosody and lexical structure of the uttered text as well as the emotions and personality of a person that is involved in a conversation. The verbal domain deals with a human voice, while body and facial gestures (head, eyes and eyebrows movement) are part of the nonverbal domain. In this article our focus is on facial gestures and how they are synchronized and driven by the prosody and lexical structure of the uttered text. We distinguish three main classes of facial gestures [9]: head, eyes and eyebrows movement. Within each class we distinguish specific gestures, each characterized by their particular parameters. The parameters that are important for the head and eyebrows movements are amplitude and velocity. Those two parameters are in inverted proportion. A movement with a big amplitude is rather slow. Table 1 shows the types of facial gestures as identified during our data analysis. We introduce symbols incorporating both a gesture type and a movement direction.

Table 1: The specification of facial gestures.

Head	Nod	An abrupt swing of a head with a similarly abrupt motion back. We have four nod directions: up and down (^), down and up (v), left and right (<) and right and left (>).
	Overshoot nod	Nod with an overshoot at the return, i.e. the pattern looks like an 'S' lying on its side (~).
	Swing	An abrupt swing of a head without a back motion. Sometimes rotation moves slowly, barely visible, back to the original pose, sometimes it is followed by an abrupt motion back after some delay. Five directions: up (u), down (d), left (l), right (R) and diagonal (diag).

	Reset	Sometimes follows swing movement. Returns head in central position.
Eyes	Movement in various directions	The eyes are always moving. Parameters are: gaze direction, points of fixation, the percentage of eye contact over gaze avoidance, duration of eye contact.
	Blink	Periodic blinks keep the eyes wet. Voluntary blinks support conversational signals and punctuators.
Eyebrows	Raise	Eyebrows go up and down (^).
	Frown	Eyebrows go down and up (?).

A. Lexical analysis of English text

The speech analysis module [13] performs linguistic and contextual analysis of a text written in English language with the goal of enabling the nonverbal (gestures) and verbal (prosody) behaviour assignment and scheduling. Starting from plain English text, it produces an XML document annotated with tags for each word. These tags allow the distinction of the newly introduced words, words known from the previous text and punctuation marks. The input text is first phrase-parsed, because the module needs to know the morphological, syntactic and part-of-speech information. In order to get the morphologic and semantic data about words in a sentence, in first release we used Connexor's Machineese Phrase Tagger [15] (MPT). MPT is commercial tool, its public interface is changing from version to version, and it has much more functionality that we needed. So we have made simplified version of morphologic and semantic analyzer extending WordNet 2.0 database [16]. In order to determine the correct word type based on the output queried from the extended WordNet 2.0 database, we must pass multiple times through whole sentence and apply various English grammatical rules. In the second step, we break the paragraphs (UTTERANCE) into clauses (CLAUSE). The largest unit is UTTERANCE, which represents an entire paragraph of input. The next unit is CLAUSE, which is held to represent a proposition. In order to detect clauses in an utterance, the module is searching for the punctuation marks and the placement of verb inside a phrase. The smallest unit is word with its new attribute. To determine the newness of each word, we keep track of all previously mentioned words in an utterance. We again use WordNet 2.0 to identify sets of synonyms. We tagged each noun, verb, adverb and adjective as new if itself or its synonym has not been seen in an utterance before. The other word classes are not considered for the new parameter. Since pronouns need to be tagged as new and WordNet 2.0 does not process them at all, a special algorithm to deal with pronouns is proposed.

B. Statistical model of gestures

As a training set for our analysis, we used the footage showing the newscasters presenting news. We investigated three female and two male Swedish newscasters. Observing those news casting clips, we marked the starting and ending frames for every eye blink, eyebrow raise and head movement. Analyzing those frames, the speakers Mouth-Nose Separation unit (MNS0) value, facial gesture amplitude value, facial gesture type and direction were determined. In our model, the basic

unit which triggers facial gestures is a word. We chose not to subdivide further into syllables or phonemes for simplicity reasons. Since some facial gestures last through two or more words, this level of subdivision seems appropriate. The raw data for the complete training set was statistically processed in order to build a statistical model of speaker behaviour. A statistical model consists of a number of components, each describing the statistical properties for a particular gesture type in a specific speech context. A speech context can be an old word, a new word or a punctuator. The statistical properties for a gesture type include the probability of occurrence of particular gestures and histograms of amplitude and duration values for each gesture. Such statistics exist for each gesture type and for each speech context we treated. They are built into the decision tree (Figure 6) that triggers gestures. The process is described in the following section. Note that, in the context of punctuators, only eyes gestures are used, because the statistics show that other gestures do not occur on punctuators.

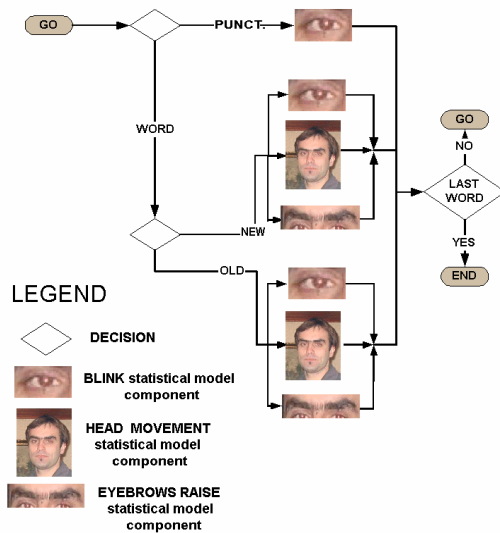


Figure 6: Decision tree with components of the statistical model

C. The System




Figure 7: The data flow through the Autonomous Speaker Agent system.

The input to the system (Figure 7) is plain English text. It is processed by a lexical analysis which converts it into an XML format with lexical tags. The facial gesture module is the core of the system – it actually inserts appropriate gestures into text in the form of special bookmark tags that are read by the TTS/MPEG-4 Encoding module. While the Microsoft Speech API (SAPI) [18] Text To Speech (TTS) engine generates an audio stream, the SAPI notification mechanism is used to catch the timing of phonemes and bookmarks the containing gesture information. Based on this

information, an MPEG-4 FBA bit stream is encoded with the appropriate viseme and facial gestures animation. For MPEG-4 FBA bit stream generation, we are using Visage SDK API [17]. The facial gesture module is built upon the statistical model described in the previous section. The statistical model is built into the decision tree illustrated in Figure 6.

The first branch point classifies the current context as either a word or a punctuation mark. Our data analysis showed that only eye blink facial gesture had occurred on the punctuation marks. Therefore only the blink component of the statistical model is implemented in this context. The words could be new or old (Section Lexical analysis of text) in the context of uttered text – this is the second branch point. All facial gestures occurred in both cases but with different probabilities. Because of that, in each case we have different components for facial gestures parameters.




Figure 8: Results of subjective evaluations. Average score and standard deviation.

1) Results

We conducted a subjective test in order to compare our proposed statistical model to simpler techniques. We synthesized facial animation on our face model using three different methods. In the first (Type 1), head and eye movements were produced playing animation sequence that was recorded by tracking movements of a real professional speaker. In the second (Type 2), we produced a facial animation using the system described in this paper. In the third (Type 3), only the character's lips were animated.

We conducted a subjective test to evaluate the three types of facial animation. The three characters (Type 1, Type 2 and Type 3) were presented in random order to 29 subjects. All three characters presented the same text. The presentation was conducted in the Ericsson Nikola Tesla and all subjects were computer specialists. However, most of the subjects were not familiar with virtual characters, and none of the subjects were authors of the study. The subjects were asked the following questions:

Q1: Did the character on the screen appear interested in (5) or indifferent (1) to you?

Q2: Did the character appear engaged (5) or distracted (1) during the conversation?

Q3: Did the personality of the character look friendly (5) or not (1)?

Q4: Did the face of the character look lively (5) or deadpan (1)?

Q5: In general, how would you describe the character?

Note that higher scores correspond to more positive attributes in a speaker. For questions 1 to 4, the score was graded on a scale of 5 to 1.

Figure 8 summarizes the average score and standard deviation (marked with a black color) for the first four questions. From the figure, we can see that the character of type 2 was graded with the highest average grade for all questions except for the Q2. The reason for that is because type 3 character only moves its lips and its head is static. This gave the audience the impression of engagement in the presentation. A Kruskal-Wallis ANOVA indicated that the three character types had significantly different scores ($p = 0.0000$).

According to general remarks in Q5, the subjects tended to believe the following:

1. Type 1 looked boring and uninteresting, it seemed to have cold personality. Also, implemented facial gestures were not related to the spoken text.
2. Type 2 had a more natural facial gesturing and facial gestures were coarticulated to some extent. Head movements and eye blinks are related to the spoken text. However, eyebrow movements were with unnatural amplitudes and were not related to the spoken text.
3. Type 3 looked irritating, stern and stony. However, it appeared to be concentrated and its lips animation was the best.

IV. CONCLUSION AND FUTURE WORK

In this paper we present two methods for automatic facial gesturing of graphically embodied animated agents. One approach is an autonomous speaker agent with full facial animation produced automatically from the plain text. With statistical data that we have gathered during our work, we have confirmed some of the conclusions of other papers. We confirmed that, on average, the amplitude of a faster head nod is lesser than the amplitude of a slower nod. Furthermore, we concluded that the words, that bring something new in the utterance context, are very often accompanied by some facial gesture. An extension to Embodied Conversational Characters is a logical item for future work, adapting and extending the statistical model to include more complicated gesturing modes and speech prosody that occur in a conversation. Coarticulation occurrence of facial displays should be also taken into account. In order to get more natural head movements, the velocity dynamics [14] of those movements must be implemented.

As well, we propose our approach for lip sync system by speech signal analysis. Speech is classified into viseme classes by neural networks and GA is used for obtaining optimal NN topology. By introducing segmentation of the speech directly into viseme classes instead of phoneme classes, computation overhead is reduced, since only visemes are used for facial animation. Automatic design of neural networks with genetic algorithms saves much time in the training process. Moreover, better results are achieved than with manual search of network configuration.

However, a face that only moves the lips, looks extremely unnatural because natural speech always involves facial gestures. Our next step will be to extend automatic lip sync system with the similar statistical model for facial gestures as proposed here in order to

generate facial expressions in addition to lip movements from the speech signal. But in case of speech driven facial gesturing, statistical model will be based on the input speech prosody, instead of lexical analysis of the text.

ACKNOWLEDGMENT

The initial version of this lip sync system has been implemented by A. Axelsson and E. Björhall as part of their master thesis of Linköping University [4] and in collaboration with Visage Technologies AB, Linköping, Sweden. This work is also partly supported by Visage Technologies.

Research on autonomous speaker agent is partly supported by Ericsson Nikola Tesla (ETK).

REFERENCES

- [1] T.Chen and R.Rao, "Audio-visual integration in multimodal communication", *Proceedings of IEEE, Special Issue on Multimedia Signal Processing*, pp. 837-852, May 1998.
- [2] D.F. McAllister, R.D. Rodman, D.L. Bitzer, A.S. Freeman, "Lip synchronization for Animation", *Proceedings of SIGGRAPH 97*, Los Angeles, CA, 1997.
- [3] I.S. Pandžic, R. Forchheimer, Editors, "MPEG-4 Facial Animation - The Standard, Implementation and Applications", *John Wiley & Sons Ltd*, ISBN 0-470-84465-5, 2002.
- [4] A.Axelsson and E. Björhall, "Real time speech driven face animation", *Master Thesis at The Image Coding Group*, Dept. of Electrical Engineering at Linköping University, Linköping 2003.
- [5] J.J.Dávila, "Genetic optimization of neural networks for the task of natural language processing", *dissertation*, New York, 1999.
- [6] R.Rojas, "Neural networks, A Systematic Introduction", *Springer-Verlag Berlin Heidelberg*, 1996.
- [7] A.J.Jones, "Genetic algorithms and their applications to the design of neural networks", *Neural Computing & Applications*, 1(1):32-45, 1993.
- [8] Black Box Genetic algorithm, <http://fdtd.rice.edu/GA/>
- [9] C. Pelachaud, N. Badler and M. Steedman. Generating Facial Expressions for Speech. *Cognitive, Science*, 20(1), 1-46, 1996.
- [10] B. Legoff and C. Benoît. A French speaking synthetic head. *ESCA Workshop on Audio-Visual Speech Processing 1997*, Rhodes, Greece, 145-148, 1997.
- [11] J.P. Lewis and F.I. Parke. Automated lipsynch and speech synthesis for character animation. *Human Factors in Computing Systems and Graphics Interface 1987*, 143-147, 1987.
- [12] K. Smid and I.S. Pandzic. A Conversational Virtual Character for the Web. *Computer Animation 2002, Geneva, Switzerland*, 240 - 248, 2002.
- [13] V. Radman. *Leksicka analiza teksta za automatsku proizvodnju pokreta lica*, Graduate work no. 2472 on Faculty of Electrical Engineering and Computing, University of Zagreb, 2004.
- [14] Hadar, U., Steiner, T., Grant, E., and Rose, F. C. 1983. Kinematics of head movements accompanying speech during conversation. *Human Movement Science*, 2:35-46.
- [15] Connexor, <http://www.connexor.com>
- [16] <http://www.cogsci.princeton.edu/~wn/>
- [17] Visage Technologies AB <http://www.visagetechologies.com/>
- [18] Microsoft Speech Technologies, <http://www.microsoft.com/speech>