

A Wearable Computer Based American Sign Language Recognizer

Thad Starner, Joshua Weaver, and Alex Pentland
Room E15-383, The Media Laboratory
Massachusetts Institute of Technology
20 Ames Street, Cambridge MA 02139
thad,joshw,sandy@media.mit.edu

Abstract

Modern wearable computer designs package workstation level performance in systems small enough to be worn as clothing. These machines enable technology to be brought where it is needed the most for the handicapped: everyday mobile environments. This paper describes a research effort to make a wearable computer that can recognize (with the possible goal of translating) sentence level American Sign Language (ASL) using only a baseball cap mounted camera for input. Current accuracy exceeds 97% per word on a 40 word lexicon.

1 Introduction

While there are many different types of gestures, the most structured sets belong to the sign languages. In sign language, where each gesture already has assigned meaning, strong rules of context and grammar may be applied to make recognition tractable.

To date, most work on sign language recognition has employed expensive “datagloves” which tether the user to a stationary machine [26] or computer vision systems limited to a calibrated area [23]. In addition, these systems have mostly concentrated on finger spelling, in which the user signs each word with finger and hand positions corresponding to the letters of the alphabet [6]. However, most signing does not involve finger spelling, but instead uses gestures which represent whole words, allowing signed conversations to proceed at or above the pace of spoken conversation.

In this paper, we describe an extensible system which uses one color camera pointed down from the brim of a baseball cap to track the wearer’s hands in real time and interpret American Sign Language (ASL) using Hidden Markov Models (HMM’s). The computation environment is being prototyped on a SGI Indy; however, the target platform is a

self-contained 586 wearable computer with DSP co-processor. The eventual goal is a system that can translate the wearer’s sign language into spoken English. The hand tracking stage of the system does not attempt a fine description of hand shape; studies of human sign readers have shown that such detailed information is not necessary for humans to interpret sign language [18, 22]. Instead, the tracking process produces only a coarse description of hand shape, orientation, and trajectory. The hands are tracked by their color: in the first experiment via solidly colored gloves and in the second, via their natural skin tone. In both cases the resultant shape, orientation, and trajectory information is input to a HMM for recognition of the signed words.

Hidden Markov models have intrinsic properties which make them very attractive for sign language recognition. Explicit segmentation on the word level is not necessary for either training or recognition [25]. Language and context models can be applied on several different levels, and much related development of this technology has already been done by the speech recognition community [9]. Consequently, sign language recognition seems an ideal machine vision application of HMM technology, offering the benefits of problem scalability, well defined meanings, a pre-determined language model, a large base of users, and immediate applications for a recognizer.

American Sign Language (ASL) is the language of choice for most deaf people in the United States. ASL’s grammar allows more flexibility in word order than English and sometimes uses redundancy for emphasis. Another variant, Signing Exact English (SEE), has more in common with spoken English but is not in widespread use in America. ASL uses approximately 6000 gestures for common words and communicates obscure words or proper nouns through finger spelling.

Conversants in ASL may describe a person, place, or thing and then point to a place in space to store

that object temporarily for later reference [22]. For the purposes of this experiment, this aspect of ASL will be ignored. Furthermore, in ASL the eyebrows are raised for a question, relaxed for a statement, and furrowed for a directive. While we have also built systems that track facial features [7], this source of information will not be used to aid recognition in the task addressed here.

The scope of this work is not to create a user independent, full lexicon system for recognizing ASL, the system should be extensible toward this goal. Another goal is real-time recognition which allows easier experimentation, demonstrates the possibility of a commercial product in the future, and simplifies archiving of test data. “Continuous” sign language recognition of full sentences is necessary to demonstrate the feasibility of recognizing complicated series of gestures. Of course, a low error rate is also a high priority. For this

Table 1: ASL Test Lexicon

| <i>part of speech</i> | <i>vocabulary</i> |
|-----------------------|---|
| pronoun | I, you, he, we, you(pl), they |
| verb | want, like, lose, dontwant, dontlike, love, pack, hit, loan |
| noun | box, car, book, table, paper, pants, bicycle, bottle, can, wristwatch, umbrella, coat, pencil, shoes, food, magazine, fish, mouse, pill, bowl |
| adjective | red, brown, black, gray, yellow |

recognition system, sentences of the form “personal pronoun, verb, noun, adjective, (the same) personal pronoun” are to be recognized. This sentence structure emphasizes the need for a distinct grammar for ASL recognition and allows a large variety of meaningful sentences to be generated randomly using words from each class. Table 1 shows the words chosen for each class. Six personal pronouns, nine verbs, twenty nouns, and five adjectives are included making a total lexicon of forty words. The words were chosen by paging through Humphries *et al.* [10] and selecting those which would generate coherent sentences when chosen randomly for each part of speech.

2 Machine Sign Language Recognition

Attempts at machine sign language recognition have begun to appear in the literature over the past

five years. However, these systems have generally concentrated on isolated signs, immobile systems, and small training and test sets. Research in the area can be divided into image based systems and instrumented glove systems.

Tamura and Kawasaki demonstrate an early image processing system which recognizes 20 Japanese signs based on matching cheremes [27]. Charayaphan and Marble [3] demonstrate a feature set that distinguishes between the 31 isolated ASL signs in their training set (which also acts as the test set). More recently, Cui and Weng [4] have shown an image-based system with 96% accuracy on 28 isolated gestures.

Takahashi and Kishino [26] discuss a user dependent Dataglove-based system that recognizes 34 of the 46 Japanese kana alphabet gestures, isolated in time, using a joint angle and hand orientation coding technique. Murakami and Taguchi [17] describe a similar Dataglove system using recurrent neural networks. However, in this experiment a 42 static-pose finger alphabet is used, and the system achieves up to 98% recognition for trainers of the system and 77% for users not in the training set. This study also demonstrates a separate 10 word gesture lexicon with user dependent accuracies up to 96% in constrained situations. With minimal training, the glove system discussed by Lee and Xu [13] can recognize 14 isolated finger signs using a HMM representation. Messing *et al.* [16] have shown a neural net based glove system that recognizes isolated finger-spelling with 96.5% accuracy after 30 training samples. Kadous [12] describes an inexpensive glove-based system using instance-based learning which can recognize 95 discrete Auslan (Australian Sign Language) signs with 80% accuracy. However, the most encouraging work with glove-based recognizers comes from Liang and Ouhyoung’s recent treatment of Taiwanese Sign language [14]. This HMM-based system recognizes 51 postures, 8 orientations, and 8 motion primitives. When combined, these constituents form a lexicon of 250 words which can be continuously recognized in real-time with 90.5% accuracy.

3 Use of Hidden Markov Models in Gesture Recognition

While the continuous speech recognition community adopted HMM’s many years ago, these techniques are just now accepted by the vision community. An early effort by Yamato *et al.* [29] uses discrete HMM’s to recognize image sequences of six different tennis

strokes among three subjects. This experiment is significant because it uses a 25x25 pixel quantized sub-sampled camera image as a feature vector. Even with such low-level information, the model can learn the set of motions and recognize them with respectable accuracy. Darrell and Pentland [5] use dynamic time warping, a technique similar to HMM's, to match the interpolated responses of several learned image templates. Schlenzig *et al.* [21] use hidden Markov models to recognize "hello," "good-bye," and "rotate." While Baum-Welch re-estimation was not implemented, this study shows the continuous gesture recognition capabilities of HMM's by recognizing gesture sequences. Closer to the task of this paper, Wilson and Bobick [28] explore incorporating multiple representations in HMM frameworks, and Campbell *et al.* [2] use a HMM-based gesture system to recognize 18 T'ai Chi gestures with 98% accuracy.

4 Tracking Hands in Video

Previous systems have shown that, given some constraints, relatively detailed models of the hands can be recovered from video images [6, 20]. However, many of these constraints conflict with recognizing ASL in a natural context, either by requiring simple, unchanging backgrounds (unlike clothing); not allowing occlusion; requiring carefully labelled gloves; or being difficult to run in real time.

In this project we have tried two methods of hand tracking: one, using solidly-colored cloth gloves (thus simplifying the color segmentation problem), and two, tracking the hands directly without aid of gloves or markings. Figure 1 shows the cap camera mount, and Figure 2 shows the view from the camera's perspective in the no-gloves case.

In both cases color NTSC composite video is captured and analyzed at 320 by 243 pixel resolution on a Silicon Graphics 200MHz Indy workstation at 10 frames per second. When simulating the self-contained wearable computer under development, a wireless transmission system is used to send real-time video to the SGI for processing [15].

In the first method, the subject wears distinctly colored cloth gloves on each hand (a pink glove for the right hand and a blue glove for the left). To find each hand initially, the algorithm scans the image until it finds a pixel of the appropriate color. Given this pixel as a seed, the region is grown by checking the eight nearest neighbors for the appropriate color. Each pixel checked is considered part of the hand. This, in effect, performs a simple morphological dilation upon

the resultant image that helps to prevent edge and lighting aberrations. The centroid is calculated as a by-product of the growing step and is stored as the seed for the next frame. Given the resultant bitmap and centroid, second moment analysis is performed as described in the following section.




Figure 1: The baseball cap mounted recognition camera.




Figure 2: View from the tracking camera.

In the second method, the the hands were tracked based on skin tone. We have found that all human hands have approximately the same hue and saturation, and vary primarily in their brightness. Using this information we can build an *a priori* model of skin color and use this model to track the hands much as was done in the gloved case. Since the hands have the same skin tone, "left" and "right" are simply assigned to whichever hand is currently leftmost and rightmost. Processing proceeds normally except for simple rules to handle hand and nose ambiguity described in the next section.

5 Feature Extraction and Hand Ambiguity

Psychophysical studies of human sign readers have shown that detailed information about hand shape is not necessary for humans to interpret sign language [18, 22]. Consequently, we began by considering only very simple hand shape features, and evolved a more complete feature set as testing progressed [25].

Since finger spelling is not allowed and there are few ambiguities in the test vocabulary based on individual finger motion, a relatively coarse tracking system may be used. Based on previous work, it was assumed that a system could be designed to separate the hands from the rest of the scene. Traditional vision algorithms could then be applied to the binarized result. Aside from the position of the hands, some concept of the shape of the hand and the angle of the hand relative to horizontal seemed necessary. Thus, an eight element feature vector consisting of each hand's x and y position, angle of axis of least inertia, and eccentricity of bounding ellipse was chosen. The eccentricity of the bounding ellipse was found by determining the ratio of the square roots of the eigenvalues that correspond to the matrix

$$\begin{pmatrix} a & b/2 \\ b/2 & c \end{pmatrix}$$

where a , b , and c are defined as

$$a = \int \int_{I'} (x')^2 dx' dy'$$

$$b = \int \int_{I'} x' y' dx' dy'$$

$$c = \int \int_{I'} (y')^2 dx' dy'$$

(x' and y' are the x and y coordinates normalized to the centroid) The axis of least inertia is then determined by the major axis of the bounding ellipse, which corresponds to the primary eigenvector of the matrix [8]. Note that this leaves a 180 degree ambiguity in the angle of the ellipses. To address this problem, the angles were only allowed to range from -90 to +90 degrees.

When tracking skin tones, the above analysis helps to model situations of hand ambiguity implicitly. When a hand occludes either the other hand or the nose, color tracking alone can not resolve the ambiguity. Since the nose remains in the same area of the frame, its position can be determined and discounted. However, the hands move rapidly and occlude each other often. When occlusion occurs, the hands appear to the above system as a single blob of larger than normal mass with significantly different moments than either of the two hands in the previous frame. In this implementation, each of the two hands is assigned the moment and position information of the single blob whenever occlusion occurs. While not as informative as tracking each hand separately, this method still retains a surprising amount of discriminating information. The occlusion event is implicitly

modeled, and the combined position and moment information are retained. This method, combined with the time context provided by hidden Markov models, is sufficient to distinguish between many different signs where hand occlusion occurs.

6 Training an HMM network

Unfortunately, space does not permit a treatment of the solutions to the fundamental problems of HMM use: evaluation, estimation, and decoding. A substantial body of literature exists on HMM technology [1, 9, 19, 30], and tutorials on their use can be found in [9, 24]. Instead, this section will describe the issues for this application.

The initial topology for an HMM can be determined by estimating how many different states are involved in specifying a sign. Fine tuning this topology can be performed empirically. While better results might be obtained by tailoring different topologies for each sign, a four state HMM with one skip transition was determined to be sufficient for this task (Figure 3). As

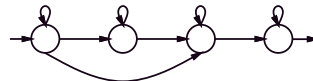


Figure 3: The four state HMM used for recognition.

an intuition, the skip state allows the model to emulate a 3 or 4 state HMM depending on the training data for the particular sign. However, in cases of variations in performance of a sign, both the skip state and the progressive path may be trained.

When using HMM's to recognize strings of data such as continuous speech, cursive handwriting, or ASL sentences, several methods can be used to bring context to bear in training and recognition. A simple context modeling method is embedded training. Initial training of the models might rely on manual segmentation. In this case, manual segmentation was avoided by evenly dividing the evidence among the models. Viterbi alignment then refines this approximation by automatically comparing signs in the training data to each other and readjusting boundaries until a minimum variance is reached. Embedded training goes on step further and trains the models *in situ* allowing model boundaries to shift through a probabilistic entry into the initial states of each model [30]. Again, the process is automated. In this manner, a more realistic model can be made of the onset and offset of a particular sign in a natural context.

Generally, a sign can be affected by both the sign in front of it and the sign behind it. For phonemes in speech, this is called “co-articulation.” While this can confuse systems trying to recognize isolated signs, the context information can be used to aid recognition. For example, if two signs are often seen together, recognizing the two signs as one group may be beneficial. Such groupings of 2 or 3 units together for recognition has been shown to halve error rates in speech and handwriting recognition [25].

A final use of context is on the inter-word (when recognizing single character signs) or phrase level (when recognizing word signs). Statistical grammars relating the probability of the co-occurrence of two or more words can be used to weight the recognition process. In handwriting, where the units are letters, words, and sentences, a statistical grammar can quarter error rates [25]. In the absence of enough data to form a statistical grammar, rule-based grammars can effectively reduce error rates.

7 Experimentation

Since we could not exactly recreate the signing conditions between the first and second experiments, direct comparison of the gloved and no-glove experiments is impossible. However, a sense of the increase in error due to removal of the gloves can be obtained since the same vocabulary and sentences were used in both experiments.

7.1 Experiment 1: Gloved-hand tracking

The glove-based handtracking system described earlier worked well. In general, a 10 frame/sec rate was maintained within a tolerance of a few milliseconds. However, frames were deleted where tracking of one or both hands was lost. Thus, a constant data rate was not guaranteed. This hand tracking process produced an 16 element feature vector (each hand’s x and y position, delta change in x and y , area, angle of axis of least inertia - or first eigenvector, length of this eigenvector, and eccentricity of bounding ellipse) that was used for subsequent modeling and recognition. Initial estimates for the means and variances of the output probabilities were provided by iteratively using Viterbi alignment on the training data (after initially dividing the evidence equally among the words in the sentence) and then recomputing the means and variances by pooling the vectors in each segment. Entropic’s Hidden Markov Model ToolKit (HTK) is used as a basis for this step and all other

Table 2: Word accuracy of glove-based system

| <i>experiment</i> | <i>training set</i> | <i>independent test set</i> |
|-------------------|---|--|
| grammar | 99.4% (99.4%) | 97.6% (98%) |
| no grammar | 96.7% (98%) (D=2, S=39, I=42, N=2500) | 94.6% (97%) (D=1, S=14, I=12, N=500) |

HMM modeling and training tasks. The results from the initial alignment program are fed into a Baum-Welch re-estimator, whose estimates are, in turn, refined in embedded training which ignores any initial segmentation. For recognition, HTK’s Viterbi recognizer is used both with and without a strong grammar based on the known form of the sentences. Contexts are not used, since a similar effect could be achieved with the strong grammar given this data set. Recognition occurs five times faster than real time.

Word recognition accuracy results are shown in Table 2; the percentage of words correctly recognized is shown in parentheses next to the accuracy rates. When testing on training, all 500 sentences were used for both the test and train sets. For the fair test, the sentences were divided into a set of 400 training sentences and a set of 100 independent test sentences. The 100 test sentences were not used for any portion of the training. Given the strong grammar (pronoun, verb, noun, adjective, pronoun), insertion and deletion errors were not possible since the number and class of words allowed is known. Thus, all errors are vocabulary substitutions when the grammar is used (accuracy is equivalent to percent correct). However, without the grammar, the recognizer is allowed to match the observation vectors with any number of the 40 vocabulary words in any order. Thus, deletion (D), insertion (I), and substitution (S) errors are possible. The absolute number of errors of each type are listed in Table 2. The accuracy measure is calculated by subtracting the number of insertion errors from the number of correct labels and dividing by the total number of signs. Note that, since all errors are accounted against the accuracy rate, it is possible to get large negative accuracies (and corresponding error rates of over 100%). Most insertion errors correspond to signs with repetitive motion.

7.2 Analysis

The 2.4% error rate of the independent test set shows that the HMM topologies are sound and that

the models generalize well. The 5.4% error rate (based on accuracy) of the “no grammar” experiment better indicates where problems may occur when extending the system. Without the grammar, signs with repetitive or long gestures were often inserted twice for each actual occurrence. In fact, insertions caused almost as many errors as substitutions. Thus, the sign “shoes” might be recognized as “shoes shoes,” which is a viable hypothesis without a language model. However, a practical solution to this problem is the use of context training and a statistical grammar instead of the rule-based grammar.

Using context modeling as described above may significantly improve recognition accuracy in a more general implementation. While a rule-based grammar explicitly constrains the word order, statistical context modeling would have a similar effect while generalizing to allow different sentence structures. In the speech community, such modeling occurs at the “triphone” level, where groups of three phonemes are recognized as one unit. The equivalent in ASL would be to recognize “trisines” (groups of three signs) corresponding to three words, or three letters in the case of finger spelling. Unfortunately, such context models require significant additional training.

In speech recognition, statistics are gathered on word co-occurrence to create “bigram” and “trigram” grammars which can be used to weight the likelihood of a word. In ASL, this might be applied on the phrase level. For example, the random sentence construction used in the experiments allowed “they like pill yellow they,” which would probably not occur in natural, everyday conversation. As such, context modeling would tend to suppress this sentence in recognition, perhaps preferring “they like food yellow they,” except when the evidence is particularly strong for the previous hypothesis.

Unlike our previous study [23] with desk mounted camera, there was little confusion between the signs “pack,” “car,” and “gray.” These signs have very similar motions and are generally distinguished by finger position. The cap-mounted camera seems to have reduced the ambiguity of these signs.

7.3 Experiment 2: Natural skin tracking

The natural hand color tracking method also maintained a 10 frame per second rate at 320x240 pixel resolution on a 200MHz SGI Indy. The word accuracy results are summarized in Table 3; the percentage of words correctly recognized is shown in parentheses next to the accuracy rates.

Table 3: Word accuracy of natural skin system

| <i>experiment</i> | <i>training set</i> | <i>independent test set</i> |
|-------------------|--|---|
| grammar | 99.3% (99%) | 97.8% (98%) |
| no grammar | 93.1% (99%) (D=5, S=30, I=138, N=2500) | 91.2% (98%) (D=1, S=8, I=35, N=500) |

7.4 Analysis

A higher error rate was expected for the gloveless system, and indeed, this was the case for less constrained “no grammar” runs. However, the error rates for the strong grammar cases are almost identical. This result was unexpected since, in previous experiments with desktop mounted camera systems [23], gloveless experiments had significantly lower accuracies. The reason for this difference may be in the amount of ambiguity caused by the user’s face in the previous experiments whereas, with the cap mounted system, the nose provided little problems.

The high accuracy rates and types of errors (repeated words) indicate that more complex versions of the experiment can now be addressed. From previous experience, context modeling or statistical grammars could significantly reduce the remaining error in the gloveless no grammar case.

8 Discussion and Conclusion

We have shown an unencumbered, vision-based method of recognizing American Sign Language (ASL). Through use of hidden Markov models, low error rates were achieved on both the training set and an independent test set without invoking complex models of the hands.

However, the cap camera mount is probably inappropriate for natural sign. Facial gestures and head motions are common in conversational sign and would cause confounding motion to the hand tracking. Instead a necklace may provide a better mount for determining motion relative to the body. Another possibility is to place reference points on the body in view of the cap camera. By watching the motion of these reference points, compensation for head motion might be performed on the hand tracking data and the head motion itself might be used as another feature.

Another challenge is porting the recognition software to the self-contained wearable computer plat-

form. The Adjeco ANDI-FG PC/104 digitizer board with 56001 DSP was chosen to perform hand tracking as a parallel process to the main CPU. The tracking information is then to be passed to a Jump 133Mhz 586 CPU module running HTK in Linux. While this CPU appears to be fast enough to perform recognition in real time, it might not be fast enough to synthesize spoken English in parallel (BT's "Laureate" will be used for synthesizing speech). If this proves to be a problem, newly developed 166Mhz Pentium PC/104 boards will replace the current CPU module in the system. The size of the current prototype computer is 5.5" x 5.5" x 2.75" and is carried with its 2 "D" sized lithium batteries in a shoulder satchel. In order to further reduce the obtrusiveness of the system, the project is switching to cameras with a cross-sectional area of 7mm. These cameras are almost unnoticeable when integrated into the cap. The control unit for the camera is the size of a small purse but fits easily in the shoulder satchel.

With a larger training set and context modeling, lower error rates are expected and generalization to a freer, user independent ASL recognition system should be attainable. To progress toward this goal, the following improvements seem most important:

- Measure hand position relative to a fixed point on the body.
- Add finger and palm tracking information. This may be as simple as counting how many fingers are visible along the contour of the hand and whether the palm is facing up or down.
- Collect appropriate domain or task-oriented data and perform context modeling both on the trisine level as well as the grammar/phrase level.
- Integrate explicit head tracking and facial gestures into the feature set.
- Collect experimental databases of native sign using the apparatus.
- Estimate 3D information based on the motion and aspect of the hands relative to the body.

These improvements do not address the user independence issue. Just as in speech, making a system which can understand different subjects with their own variations of the language involves collecting data from many subjects. Until such a system is tried, it is hard to estimate the number of subjects and the amount of data that would comprise a suitable training database. Independent recognition often places

new requirements on the feature set as well. While the modifications mentioned above may be initially sufficient, the development process is highly empirical.

So far, finger spelling has been ignored. However, incorporating finger spelling into the recognition system is a very interesting problem. Of course, changing the feature vector to address finger information is vital to the problem, but adjusting the context modeling is also of importance. With finger spelling, a closer parallel can be made to speech recognition. Trisine context occurs at the sub-word level while grammar modeling occurs at the word level. However, this is at odds with context across word signs. Can trisine context be used across finger spelling and signing? Is it beneficial to switch to a separate mode for finger spelling recognition? Can natural language techniques be applied, and if so, can they also be used to address the spatial positioning issues in ASL? The answers to these questions may be key to creating an unconstrained sign language recognition system.

Acknowledgements

The authors would like to thank Tavenner Hall for her help editing and proofing early copies of this document.

References

- [1] L. Baum. "An inequality and associated maximization technique in statistical estimation of probabilistic functions of Markov processes." *Inequalities*, 3:1-8, 1972.
- [2] L. Campbell, D. Becker, A. Azarbajejani, A. Bobick, and A. Pentland. "Invariant features for 3-D gesture recognition," *Intl. Conf. on Face and Gesture Recogn.*, pp. 157-162, 1996
- [3] C. Charayaphan and A. Marble. "Image processing system for interpreting motion in American Sign Language." *Journal of Biomedical Engineering*, 14:419-425, 1992.
- [4] Y. Cui and J. Weng. "Learning-based hand sign recognition." *Intl. Work. Auto. Face Gest. Recogn. (IWAFFGR) '95 Proceedings*, p. 201-206, 1995
- [5] T. Darrell and A. Pentland. "Space-time gestures." *CVPR*, p. 335-340, 1993.

- [6] B. Dorner. "Hand shape identification and tracking for sign language interpretation." IJCAI Workshop on Looking at People, 1993.
- [7] I. Essa, T. Darrell, and A. Pentland. "Tracking facial motion." IEEE Workshop on Nonrigid and articulated Motion, Austin TX, Nov. 94.
- [8] B. Horn. Robot Vision. MIT Press, NY, 1986.
- [9] X. Huang, Y. Ariki, and M. Jack. Hidden Markov Models for Speech Recognition. Edinburgh Univ. Press, Edinburgh, 1990.
- [10] T. Humphries, C. Padden, and T. O'Rourke. A Basic Course in American Sign Language. T. J. Publ., Inc., Silver Spring, MD, 1980.
- [11] B. Juang. "Maximum likelihood estimation for mixture multivariate observations of Markov chains." AT&T Tech. J., 64:1235-1249, 1985.
- [12] W. Kadous. "Recognition of Australian Sign Language using instrumented gloves." Bachelor's thesis, University of New South Wales, October 1995.
- [13] C. Lee and Y. Xu, "Online, interactive learning of gestures for human/robot interfaces." IEEE Int. Conf. on Robotics and Automation, pp 2982-2987, 1996.
- [14] R. Liang and M. Ouhyoung, "A real-time continuous gesture interface for Taiwanese Sign Language." Submitted to UIST, 1997.
- [15] S. Mann "Mediated reality", "MIT Media Lab, Perceptual Computing Group TR# 260"
- [16] L. Messing, R. Erenshteyn, R. Foulds, S. Galuska, and G. Stern. "American Sign Language computer recognition: Its Present and its Promise" Conf. the Intl. Society for Augmentative and Alternative Communication, 1994, pp. 289-291.
- [17] K. Murakami and H. Taguchi. "Gesture recognition using recurrent neural networks." CHI '91 Conference Proceedings, p. 237-241, 1991.
- [18] H. Poizner, U. Bellugi, and V. Lutes-Driscoll. "Perception of American Sign Language in dynamic point-light displays." J. Exp. Psychol.: Human Perform., 7:430-440, 1981.
- [19] L. Rabiner and B. Juang. "An introduction to hidden Markov models." IEEE ASSP Magazine, p. 4-16, Jan. 1996.
- [20] J. Rehg and T. Kanade. "DigitEyes: vision-based human hand tracking." School of Computer Science Technical Report CMU-CS-93-220, Carnegie Mellon Univ., Dec. 1993.
- [21] J. Schlenzig, E. Hunter, and R. Jain. "Recursive identification of gesture inputers using hidden Markov models." *Proc. Second Ann. Conf. on Appl. of Comp. Vision*, p. 187-194, 1994.
- [22] G. Sperling, M. Landy, Y. Cohen, and M. Pavel. "Intelligible encoding of ASL image sequences at extremely low information rates." *Comp. Vision, Graphics, and Image Proc.*, 31:335-391, 1985.
- [23] T. Starner and A. Pentland. "Real-Time American Sign Language Recognition from Video Using Hidden Markov Models." MIT Media Laboratory, Perceptual Computing Group TR#375, Presented at ISCV'95.
- [24] T. Starner. "Visual Recognition of American Sign Language Using Hidden Markov Models." Master's thesis, MIT Media Laboratory, Feb. 1995.
- [25] T. Starner, J. Makhoul, R. Schwartz, and G. Chou. "On-line cursive handwriting recognition using speech recognition methods." ICASSP, V-125, 1994.
- [26] T. Takahashi and F. Kishino. "Hand gesture coding based on experiments using a hand gesture interface device." SIGCHI Bul., 23(2):67-73, 1991.
- [27] S. Tamura and S. Kawasaki. "Recognition of sign language motion images." *Pattern Recognition*, 21:343-353, 1988.
- [28] A. Wilson and A. Bobick. "Learning visual behavior for gesture analysis." *Proc. IEEE Int'l. Symp. on Comp. Vis.*, Nov. 1995.
- [29] J. Yamato, J. Ohya, and K. Ishii. "Recognizing human action in time-sequential images using hidden Markov models." *Proc. 1992 ICCV*, p. 379-385. IEEE Press, 1992.
- [30] S. Young. HTK: Hidden Markov Model Toolkit V1.5. Cambridge Univ. Eng. Dept. Speech Group and Entropic Research Lab. Inc., Washington DC, Dec. 1993.