

# Ham2Pose: Animating Sign Language Notation into Pose Sequences

Rotem Shalev Arkushin  
 Reichman University  
 rotemroo@gmail.com

Amit Moryossef  
 Bar-Ilan University  
 amitmoryossef@gmail.com

Ohad Fried  
 Reichman University  
 ofried@runi.ac.il

<https://rotem-shalev.github.io/ham-to-pose>

## Abstract

Translating spoken languages into Sign languages is necessary for open communication between the hearing and hearing-impaired communities. To achieve this goal, we propose the first method for animating a text written in HamNoSys, a lexical Sign language notation, into signed pose sequences. As HamNoSys is universal by design, our proposed method offers a generic solution invariant to the target Sign language. Our method gradually generates pose predictions using transformer encoders that create meaningful representations of the text and poses while considering their spatial and temporal information. We use weak supervision for the training process and show that our method succeeds in learning from partial and inaccurate data. Additionally, we offer a new distance measurement that considers missing keypoints, to measure the distance between pose sequences using DTW-MJE. We validate its correctness using AUTSL, a large-scale Sign language dataset, show that it measures the distance between pose sequences more accurately than existing measurements, and use it to assess the quality of our generated pose sequences. Code for the data pre-processing, the model, and the distance measurement is publicly released for future research.

## 1. Introduction

Sign languages are an important communicative tool within the deaf and hard-of-hearing (DHH) community and a central property of Deaf culture. According to the World Health Organization, there are more than 70 million deaf people worldwide [56], who collectively use more than 300 different Sign languages [29]. Using the visual-gestural modality to convey meaning, Sign languages are considered natural languages [40], with their own grammar and lexicons. They are not universal and are mostly independent of spoken languages. For example, American Sign Language (ASL)—used predominantly in the United States—and British Sign Language (BSL)—used predominantly in




Figure 1. **German Sign Language sign for “Haus”.** **Gloss** is a unique semantic identifier; **HamNoSys** and **SignWriting** describe the phonology of a sign: *Two flat hands with fingers closed, rotated towards each other, touching, then symmetrically moving diagonally downwards.*

the United Kingdom—are entirely different, despite English being the predominant spoken language in both. As such, the translation task between each signed and spoken language pair is different and requires different data. Building a robust system that translates spoken languages into Sign languages and vice versa is fundamental to alleviate communication gaps between the hearing-impaired and the hearing communities.

While translation research from Sign languages into spoken languages has rapidly advanced in recent years [2, 5, 6, 30, 36, 37], translating spoken languages into Sign languages, also known as Sign Language Production (SLP), remains a challenge [41, 42, 47, 48]. This is partially due to a misconception that deaf people are comfortable reading spoken language and do not require translation into Sign language. However, there is no guarantee that someone whose first language is, for example, BSL, exhibits high literacy in written English. SLP is usually done through an intermediate notation system such as a semantic notation system, *e.g.* gloss (Sec. 2.1), or a lexical notation system, *e.g.* HamNoSys, SignWriting (Sec. 2.2). The spoken language text is translated into the intermediate notation, which is then translated into the relevant signs. The signs can either be animated avatars or pose sequences later converted into videos. Previous work has shown progress in translating spoken language text to Sign language lexical notations, namely HamNoSys [54] and SignWriting [21], and in converting pose sequences into videos [8, 43, 55]. There has been some work on animating HamNoSys into avatars [3, 12, 13, 58], with unsatisfactory results (Sec. 3.1),

but no work on the task of animating HamNoSys into pose sequences. Hence, in this work, we focus on animating HamNoSys into signed pose sequences, thus facilitating the task of SLP with a generic solution for all Sign languages. To do this, we collect and combine data from multiple HamNoSys-to-video datasets [25, 26, 32], extract pose keypoints from the videos using a pose estimation model, and process these further as detailed in Sec. 4.1. We use the pose features as weak labels to train a model that gets HamNoSys text and a single pose frame as inputs and gradually generates the desired pose sequence from them. Despite the pose features being inaccurate and incomplete, our model still learns to produce the correct motions. Additionally, we offer a new distance measurement that considers missing keypoints, to measure the distance between pose sequences using *DTW-MJE* [20]. We validate its correctness using AUTSL, a large-scale Sign language dataset [44], and show that it measures pose sequences distance more accurately than currently used measurements. Overall, our main contributions are:

1. We propose the first method for animating HamNoSys into pose sequences.
2. We offer a new pose sequences distance measurement, validated on a large annotated dataset.
3. We combine existing datasets, converting them to one enhanced dataset with processed pose features.

## 2. Background

It is common to use an intermediate notation system for the SLP task. We discuss three of such notations below and show an example of them in Fig. 1.

### 2.1. Semantic Notation Systems

Semantic notation systems are the most popular form of Sign language annotation. They treat Sign languages as discrete and annotate meaning units. A well-known semantic notation system is Gloss, a notation of the meaning of a word. In Sign languages, glosses are usually sign-to-word transcriptions, where every sign has a unique identifier written in a spoken language. Some previous works focused on translating spoken language into glosses [41, 47, 48]. The glosses are then usually transformed into videos using a gloss-to-video dictionary. However, since glosses are Sign language-specific, each translation task is different from the other and requires different data.

### 2.2. Lexical Notation Systems

Lexical notation systems annotate phonemes of Sign languages and can be used to transcribe any sign. Given that all the details about how to produce a sign are in the transcription itself, these notations can be universal and can be used to transcribe every Sign language. Two examples of

such universal notations are Hamburg Sign Language Notation System (HamNoSys) [18], and SignWriting [49]. They have a direct correspondence between symbols (glyphs) and gesture aspects, such as hand location, orientation, shape, and movement. As they are not language-specific, it allows a Sign language invariant transcription for any desired sign. Unlike HamNoSys, SignWriting is written in a spatial arrangement that does not follow a sequential order. While this arrangement was designed for better human readability, it is less “friendly” for computers, expecting sequential text order. As demonstrated in Fig. 1, HamNoSys is very expressive, precise, and easy to learn and use. Moreover, each part of it (each glyph) is responsible for one aspect of the sign, similarly to parts-of-speech in a spoken language sentence. Therefore, each HamNoSys sequence can be thought of as a “sentence” and each glyph as a dictionary “word” from the corpus of HamNoSys glyphs.

### 2.3. Sign Language Notation Data

Many Sign language corpora are annotated with Glosses [11, 14, 24, 32, 44]. However, as there is no single standard for gloss annotation, each corpus has its own unique identifiers for each sign at different granularity levels. This lack of universality—both in annotation guidelines and in different language data—makes it difficult to use and combine multiple corpora and design impactful translation systems. Furthermore, since SignWriting is used more to “write” Sign language on the page, and not for video transcription, existing SignWriting resources [4, 50, 51] usually include parallel SignWriting and spoken language text, without including parallel videos of the performed signs. HamNoSys, on the other hand, was designed more for annotating existing Sign language videos, and as such, resources including HamNoSys [25, 26, 32] always include parallel videos. Compared to resources including gloss annotations, these resources are small, with only hundreds to thousands of signs with high-quality annotation. However, the language universality and annotation cohesion in these corpora allow grouping them together for the usage of all data without any annotation modification.

## 3. Related Work

In this section, we review related work in the field of SLP. We cover avatar approaches using HamNoSys as input and gloss approaches. Moreover, we cover HamNoSys generation work, including translating spoken language text or videos into HamNoSys, as these tasks allow for a full translation pipeline together with our work. Furthermore, we mention diffusion models as a source of inspiration for our method, and text-to-motion works, explaining how our problem and data are different from them.




Figure 2. **JASign (SigML) failure cases.** hand-inside-hand (H-in-H), hand-inside-clothes (H-in-C) artifacts, wrong signing.

### 3.1. Avatar Approaches

Since the early 2000s, there have been several research projects exploring avatars animated from HamNoSys, such as VisiCast [3], eSign [58], dicta-sign [13], and JASigning [12]. While these avatars produce sign sequences, they are not popular among the deaf community due to under-articulated and unnatural movements, making the avatars difficult to understand [48]. Furthermore, the robotic movements of these avatars can make viewers uncomfortable due to the uncanny valley phenomenon [27]. In addition, as illustrated in Fig. 2, these avatars do not perform all hand motions correctly. A later work [16], uses motion capture data to create more stable and realistic avatars. However, this method is limited to a small set of phrases due to the high data collection and annotation cost.

### 3.2. Gloss Approaches

To cope with these challenges, a recent work [48] suggests combining generative models with a motion graph (MG) [22] and Neural Machine Translation. They translate spoken language sentences into gloss sequences that condition an MG to find a pose sequence representing the input from a dictionary of poses. The sequence is then converted to a video using a GAN. A similar work [41] suggests progressive transformers for generating signed pose sequences from spoken language through glosses. Like Stoll *et al.* [48], they use a closed set of dictionary signs as the signs in their output sequence, which makes these solutions language and data-specific. A later work [42] uses learned “cheremes<sup>1</sup>” to generate signs. Similarly to glosses and phonemes, cheremes are language specific. In contrast, since our method uses a universal notation, it works for languages it was trained on and for unseen languages, as long as the individual glyphs exist in our dataset.

<sup>1</sup>A Sign language equivalent of phonemes.

### 3.3. HamNoSys Generation

Recently, different HamNoSys generation tasks have been researched. For example, Skobov *et al.* suggested a method for automatically annotating videos into HamNoSys [45]. Further research on this task could enhance the capabilities of our model, by creating more labeled data. Translating spoken language text into HamNoSys has also been researched [54], and if improved, can allow a complete translation pipeline from spoken language text into Sign languages using our model.

### 3.4. Diffusion Models

Diffusion models [19,46] recently showed impressive results on image and video generation tasks [10,35,38]. Generation is done using a learned gradual process, with equal input and output sizes. The model gradually changes the input to get the desired output. In this work, we take inspiration from diffusion models in the sense that our model learns to gradually convert the input (a sequence of a duplicated reference frame) into the desired pose sequence.

### 3.5. Text to Motion

In recent years, works on motion generation from English text [1,15,17,31,52] showed impressive results. While these works may seem related to our task, they use 3D motion capture data, which is not available for our task. As detailed in Sec. 4, our data is collected using a pose estimation model over sign videos; thus, it is both 2D and imperfect, with many missing and incorrect keypoints. Moreover, since the text in these works is written in English, recent works [15,31,52] take advantage of large pre-trained language models such as BERT [9], CLIP [33], etc. As HamNoSys is not a common language, with limited available resources, we cannot use pre-trained models as they do.

## 4. Data

Our dataset consists of 5,754 videos of Sign languages signs with their HamNoSys transcriptions. Each video is of a front-facing person signing a single sign. We collect the data from the DGS Corpus [32], Dicta-Sign [26], and the corpus-based dictionary of Polish Sign Language [25]. Together, the data contains four Sign languages, signed by 14 signers: Polish SL (PJM): 2,560 signs, 2 signers; German SL (DGS): 1,926 signs, 8 signers; Greek SL (GSL): 887 signs, 2 signers; and French SL (LSF): 381 signs, 2 signers.

### 4.1. Data Pre-Processing

To use the collected data as ground truth (GT) for sign pose sequence generation, we extract estimated pose keypoints from each video using the OpenPose [7] pose estimation model. Each keypoint  $k_i \in K$  consists of a 2D location  $(x, y)$  and the confidence of the model,  $c_i$ . Missing

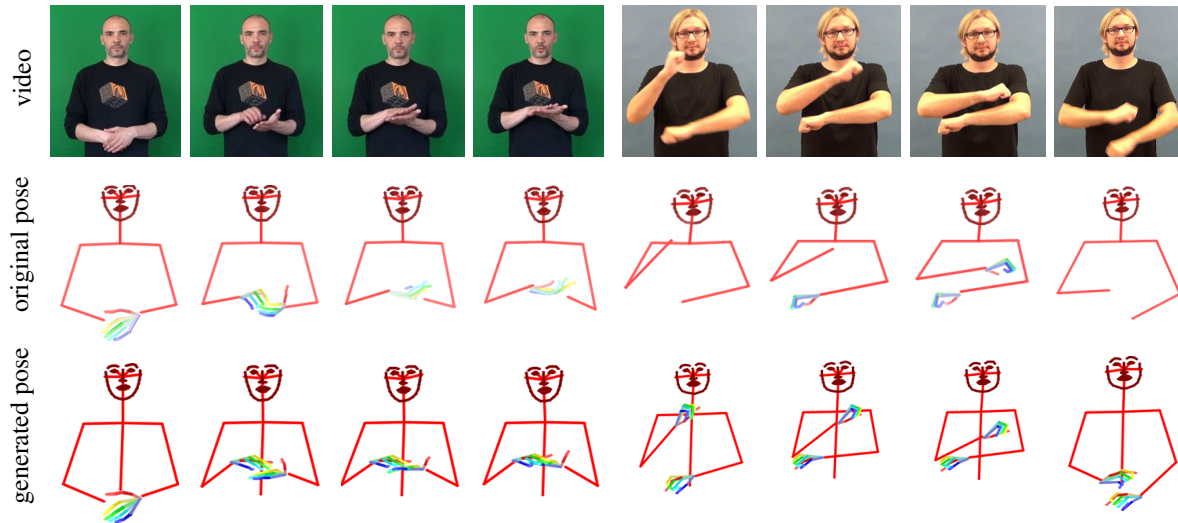


Figure 3. **Results examples:** **Top row:** original video frames, **middle row:** ground truth pose detected by OpenPose, **bottom row:** generated pose. Despite missing keypoints in the ground truth pose, our model generates a correct pose.

keypoints (keypoints with  $c_i = 0$ ) or ones with a confidence of less than 20% are filtered out. We further process the extracted keypoints as follows:

1. **Trim meaningless frames.** Some videos (*e.g.* videos fading in and out) contain leading or trailing frames that do not contain enough relevant information. Moreover, in resting position, hands are not always visible in the video frame. Hence, leading / trailing frames in which the face or both hands are not identified are removed.
2. **Mask legs and unidentified keypoints.** In addition to setting a confidence threshold, since the legs are not significant, we remove them by setting the confidence of every keypoint from the waist down to 0, allowing the model to only learn from existing and relevant keypoints.
3. **Flip left-handed sign videos.** One-hand signs are usually signed with the dominant hand. A left-handed person would mirror a sign signed by a right-handed person. Given that the signing hand is not specified in HamNoSys and that some videos in our dataset include left-handed signers, for consistency, we flip these videos to produce right-handed signs.
4. **Pose Normalization.** We normalize the pose keypoints by the pose shoulders, using the *pose\_format* library [28], so all poses have the same scale. It defines the center of a pose to be the neck, calculated as the average middle point between the shoulders across all frames. Then, it translates all keypoints, moving the center to  $(0, 0)$ , and scales the pose so the average distance between the shoulders is 1.

We note that the data is still imperfect, with many missing and incorrect keypoints, as seen in Fig. 3.

## 5. Method

Given a sign written in HamNoSys, our model generates a sequence of frames signing the desired sign. We start from a single given “reference” pose frame, acting as the start of the sequence, and duplicate it to the length of the signed video. The sign is generated gradually over  $T$  steps, where in each time step  $t \in \{T \dots 0\}$  the model predicts the required change from step  $t$  to step  $t - 1$  as described in Sec. 5.3.1. Our method takes inspiration from diffusion models in that it learns to gradually generate a sign pose from a reference pose and a HamNoSys notation guidance. Unlike diffusion models, we start from a reference pose and not from random noise to allow for the continuation of signs with the same identity. Furthermore, at each prediction step, we predict the change required to move to the previous step rather than predicting the change required to move to the final sequence. This way, the model can replace missing and incorrect keypoints with correct ones in a gradual process. We combine these ideas with transformer encoders [53] for the pose and HamNoSys, leading to meaningful representations, considering their spatial and temporal meanings.

### 5.1. Model Architecture

As shown in Fig. 4, the model is composed of two parts: the text processor (Sec. 5.2), responsible for the HamNoSys text encoding and predicting the length of the generated pose sequence; and the pose generator (Sec. 5.3), responsible for the pose sequence generation. The inputs to the model are a HamNoSys sequence and a single pose frame. The reference frame is the starting point of the generated pose. It can either be a resting pose or the last frame of a previously generated sign for continuity purposes. The pipeline of the model is as follows: the input text is tok-



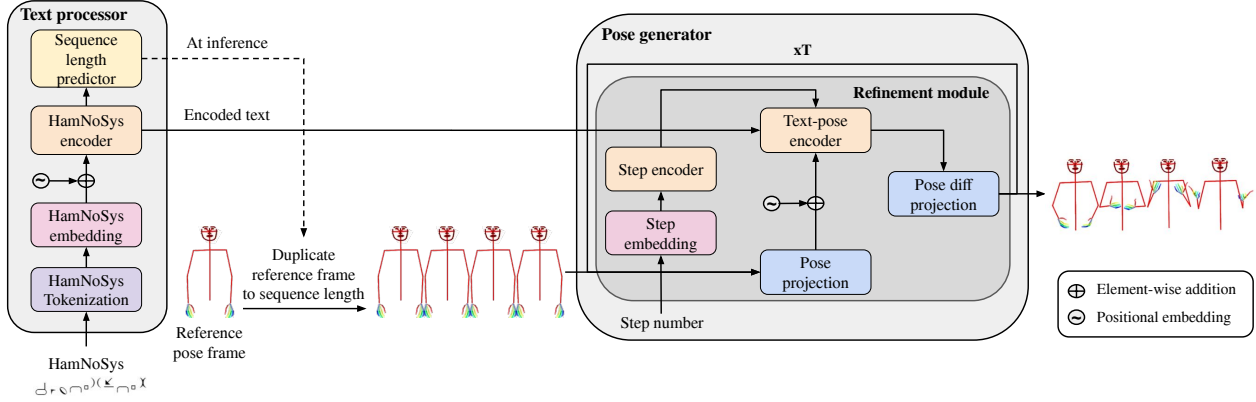


Figure 4. **Model architecture.** First, the text processor encodes the HamNoSys and predicts the sequence length. Next, the reference pose is duplicated to the sequence length and passed to the pose generator, which iteratively uses the current pose sequence and HamNoSys encoding for  $T$  steps and generates the desired pose. After  $T$  steps, the pose generator outputs the final pose sequence.

enized, embedded, and encoded by the text processor; then, it is passed to the sequence length predictor that predicts the sequence length. Next, the reference pose is duplicated to the desired pose length (during training, it is the GT length, while at inference, it is the length predicted by the sequence length predictor) and is passed to the pose generator with the encoded text. Finally, using the encoded text and extended reference pose, the pose generator (Sec. 5.3) gradually refines the pose sequence over  $T$  steps to get the desired pose. This process is summarized in Alg. 1.

## 5.2. Text Processor

This module is responsible for the HamNoSys text processing. The HamNoSys text is first tokenized into a tokens vector, so each glyph gets a unique identifier (token). Next, the tokens’ vectors are passed through a learned embedding layer, producing vector representations for each token. In addition, we use a learned embedding layer as a positional embedding to represent the positions of the sequence tokens, so the model gets information about the order of the tokens. The vector representations of the tokens’ locations are of the same dimension  $D$  of the tokens’ vector representations, to allow the summation of them. After the tokens and their locations are embedded and combined, they are passed to the HamNoSys transformer encoder [53]. Finally, the encoded text is passed to the pose generator and to the sequence length predictor—a linear layer that predicts the length of the pose sequence.

## 5.3. Pose Generator

The pose generator is responsible for the sign pose sequence generation, which is the output of the entire model. It does so gradually over  $T$  steps, where at time step  $T$  ( $s_T$ ), the sequence is the given reference frame extended to the sequence length. During training, this is the actual length of

the sign video after frame trimming (Sec. 4.1), while at inference, this is the length predicted by the sequence length predictor. To generate the desired sign gradually, we define a schedule function ( $\delta \in [0, 1]$ ) as  $\delta_t = \log_T(T - t)$ , a step size  $\alpha_t = \delta_t - \delta_{t+1}$ <sup>2</sup> and the predicted pose sequence at time step  $t$ ,  $\hat{s}_t$ , for  $t \in \{T - 1, \dots, 0\}$ , as:

$$\hat{s}_t = \alpha_t p_t + (1 - \alpha_t) s_{t+1} \quad (1)$$

where  $p_t$  is the pose value predicted by the refinement module (Sec. 5.3.1) at time step  $t$ . This way, since the previous step result is input to the current step, and the step size decreases over time, the model needs to predict smaller changes in each step. As a result, the coarser details are generated first, and the finer details are generated as the generation process proceeds. Moreover, since the result of each step is a blending between the previous step result and the current prediction and not only an addition of the prediction to the previous step, we give less weight to the initial pose sequence at each step. This way, the model can fill missing and incorrect keypoints by gradually replacing them with correct data. Additionally, since it is a gradual process, where the model predicts small changes at each step instead of predicting and replacing the whole pose sequence, its results are more smooth and accurate. Fig. 5 shows the importance of blending. To make the model more robust, we add  $\epsilon z$  noise to  $\hat{s}_t$  ( $z \sim N(0, I)$ ) at each time step during training. Finally,  $s_0$  is returned as the sign pose prediction.

### 5.3.1 Refinement Module

At each step  $t \in \{T - 1, \dots, 0\}$ , the pose generator calls the refinement module with the previously generated pose  $s_{t+1}$ , the pose positional embedding, the step number, and

<sup>2</sup>for  $t = T - 1$  we use a constant 0.1 to avoid illegal calculations




Figure 5. **Blend importance example.** Left to right: original image, original pose, pose generated by addition, by replacement, and by blend.

the encoded text. The step number is passed through an embedding layer that produces a vector representation of dimension  $D$ , which is encoded using two linear layers with activations between them. Similarly, each pose frame of  $s_{t+1}$  is projected to a vector representation of dimension  $D$ , using two linear layers with activation. The projected result is summed with the positional embedding of the pose to form a pose embedding. The pose embeddings are then concatenated with the encoded text and step, and together, as shown in Fig. 4, they are passed to the text-pose transformer encoder [53]. Finally, the result is passed to the “pose diff projection”, formed of two linear layers with activation, which generates the predicted pose for the current step,  $p_t$ , that is the output of this module.

---

#### Algorithm 1 text2pose

---

**Input: text:** HamNoSys tokens  
**ref\_pose:** a single reference pose frame

**Output:  $s_0$ :** pose sequence prediction

- 1:  $et, seq\_len = process\_text(text)$
  - 2:  $s_T = [ref\_pose] * seq\_len$
  - 3:  $pe = positional\_embedding([0, \dots, seq\_len])$
  - 4: **for**  $step = T - 1, \dots, 0$  **do**
  - 5:      $se = encode\_step(step)$
  - 6:      $p_t = refine(s_{t+1}, et, pe, se)$
  - 7:      $s_t = \alpha_t p_t + (1 - \alpha_t) s_{t+1} + \epsilon z$
  - 8: **end for**
- 

### 5.4. Loss Function

At every refinement step we want to compare the predicted sequence to an interpolation between the real sequence  $s_0$  and the starting sequence  $s_T$ . However, as our model uses the prediction of the previous step, at time step  $t$  we interpolate between  $s_0$  and the previous step  $s_{t+1}$ . For that purpose, we define:

$$s_t = \delta_t s_0 + (1 - \delta_t) s_{t+1} \quad (2)$$

We mark the  $i$ th joint in the  $j$ th frame in  $s_t$  by  $s_t^i[j]$ . The refinement loss function  $L_p$  at time step  $t$  is a weighted MSE with weight  $c_i[j]$  for each joint  $i \in K$  in frame  $j \in \{0, \dots, N\}$ :

$$L_p(s_t, \hat{s}_t) = \frac{1}{N} \sum_{j=0}^N \frac{1}{|K|} \sum_{i=0}^{|K|} c_i[j] (s_t^i[j] - \hat{s}_t^i[j])^2 \quad (3)$$

To avoid affecting the learning rate when experimenting with different step numbers, we scale the loss by  $\ln(T)^2$  (See full derivation in the supplementary material).

To train the sequence length predictor, we calculate MSE loss between the predicted sequence length  $\hat{N}$  and the real one  $N$ , and add it to the final loss with a small weight of  $\gamma$ . The complete loss term is then:

$$L = \ln(T)^2 L_p(s_t, \hat{s}_t) + \gamma \cdot L_{len}(N, \hat{N}) \quad (4)$$

## 6. Implementation Details

We provide a detailed model architecture in the supplementary material. We use learned embedding layers with dimension  $D = 128$  to embed the HamNoSys text, the step number, and the text and pose positions. For the pose and text encoding layers, we use a Transformer encoder [53] with two heads, with depths 2 and 4 for the text and pose respectively. We project the pose frames using two linear layers with swish activation [34] between them to a vector with the same dimension  $D = 128$  for each frame. The step encoder and the pose diff projection are also formed of two linear layers with swish activations between them. After experimenting with three step number options, we set  $T = 10$  as the number of steps for the pose generation process in all our experiments. We discuss the experiments and their results in the supplementary material. We train using the Adam Optimizer for 2000 epochs, setting the learning rate to  $1e-3$  empirically, and teacher forcing with probability of 0.5. For the noise addition we use  $\epsilon = 1e-4$ , and for the sequence length loss weight we use  $\gamma = 2e-5$ . We use this value because a lower value prevents the sequence length predictor from learning, while a higher value prevents the refinement module from learning.

### 6.1. Teacher Forcing

During training, we use teacher forcing [57] in the pose generation process with a probability of  $p = 0.5$ . Meaning, with a probability of  $p$ , we feed the refinement module (Sec. 5.3.1) with  $s_t$  as defined in Eq. (2), and with probability of  $1 - p$  we feed it with the predicted  $\hat{s}_t$  as defined in Eq. (1), to help the model learn faster. Hence, during training,  $s_t$  is defined by:

$$f \sim Ber(p); z \sim N(0, I)$$

$$s_t = \begin{cases} \alpha_t p_t + (1 - \alpha_t) s_{t+1} + \epsilon z & f = 0 \\ \delta_t s_0 + (1 - \delta_t) s_{t+1} + \epsilon z & \text{otherwise} \end{cases} \quad (5)$$

Method	prec@1 ↑	prec@5 ↑	prec@10 ↑	mAP ↑
MSE	0.5	0.36	0.34	0.27
nMSE	0.82	0.68	0.62	0.4
APE	0.82	0.55	0.45	0.29
nAPE	0.93	0.78	0.7	0.44
DTW-MJE	1	0.78	0.66	0.33
nDTW-MJE	<b>1</b>	<b>0.9</b>	<b>0.84</b>	<b>0.58</b>

Table 1. Distance measurements results over AUTSL

## 7. Evaluation

Currently, there is no suitable evaluation method for SLP in the literature. APE (Average Position Error) is a distance measurement used to compare poses in recent text-to-motion works [1, 15, 31]. It is the average L2 distance between the predicted and the GT pose keypoints across all frames and data samples. Since it compares absolute positions, it is sensitive to different body shapes and slight changes in timing or position of the performed movement. Huang *et al.* [20] suggested *DTW-MJE* (Dynamic Time Warping - Mean Joint Error), which measures the mean distance between pose keypoints after aligning them temporally using *DTW* [23]. However, it is unclear from the original *DTW-MJE* definition how to handle missing keypoints, hence we suggest a new distance function that considers missing keypoints and apply *DTW-MJE* with our distance function over normalized keypoints. We mark this method by *nDTW-MJE*. We validate the correctness of *nDTW-MJE* by using AUTSL [44], a large-scale Turkish Sign Language dataset, showing that it measures pose sequences distance more accurately than existing measurements. Then, we evaluate the results of our model using *nDTW-MJE* in two ways: Distance ranks (Sec. 7.3) and Leave-one-out (Sec. 7.4). We test the sequence length predictor using absolute difference between the real and predicted sequence length. The mean difference is 3.61, which usually means more resting pose frames in one of the poses. Finally, we provide qualitative results in Fig. 3 and in the supplementary material.

### 7.1. Distance measurement (*nDTW-MJE*)

We suggest a new method for measuring the distance between two pose sequences using *DTW-MJE* [20] over normalized pose keypoints, that considers missing keypoints. To measure the distance between two pose sequences, we use a variant of *DTW* [23]. *DTW* is an algorithm that can measure the distance between two temporal sequences, because it is able to ignore global and local shifts in the time dimension. We use a linear-time approximation of *DTW* suggested by Salvador *et al.* [39] while considering missing keypoints using the following distance function for each keypoints pair:

$$dist(ref, other) = \begin{cases} 0 & ref = NaN \\ \frac{\|ref\|_2}{2} & ref \neq NaN \ \& \ other = NaN \\ \|ref - other\|_2 & otherwise \end{cases} \quad (6)$$

This way, we only compare keypoints existing in both poses and punish keypoints that exist in the *reference* pose but not in the *other* pose.

To measure the distance from *ref* to *other*, we normalize them as described in Sec. 4 and calculate the *DTW-MJE* distance using Eq. (6) between the normalized pose sequences.

### 7.2. nDTW-MJE Validation

To validate *nDTW-MJE* with our suggested distance measurement, we test it against AUTSL [44] — a large-scale Turkish Sign Language dataset comprising 226 signs performed by 43 different signers in different positions and postures and 38,336 isolated sign video samples in total. We collect all samples for every unique sign  $s$ , and randomly sample  $4 \times |s|$  other signs from the dataset to keep a 1 : 4 ratio of *same* : *other* samples. Then, we measure the distance from the reference sign to each one of the samples using *nDTW-MJE* as explained in Sec. 7.1, and calculate *mAP* (mean Average Precision) and *mean precision@k* for  $k = 1, 5, 10$ , across all signs, measuring how many of the  $k$  most similar poses were of a video signing the query sign. In Tab. 1, we compare the results of using *nDTW-MJE* vs. using *MSE* and *APE* over normalized and unnormalized pose keypoints and of using *DTW-MJE* over unnormalized pose keypoints (with our distance function). The results show that *DTW-MJE* with our distance function better suits pose sequences distance measurement than the more commonly used *APE*. Moreover, using normalized keypoints improves the performance of all measurements, while *nDTW-MJE* measures the distances more accurately than all other options.

### 7.3. Distance Ranks

To evaluate our animation method, we calculate the distance between each prediction and GT pair using *nDTW-MJE*. Then, for each pair, we create a gallery of 20 random GT samples from the dataset and 20 predictions for other samples from the dataset. For each gallery sample, we calculate its distance from the prediction and the GT as references. Finally, we calculate rank 1, 5, and 10, where rank  $k$  is the percentage of test samples for which the tested pose (GT pose for the distance to prediction case, prediction for the distance to GT case), was in the  $k$  most similar poses to the reference pose and report them in Tab. 2. For comparison, we use the most similar previous work to ours,

Reference	Model	Rank 1 $\uparrow$	Rank 5 $\uparrow$	Rank 10 $\uparrow$
Prediction	PT	0.04	0.19	0.34
	Resting	0.05	0.17	0.31
	<b>Ours</b>	<b>0.08</b>	<b>0.2</b>	<b>0.35</b>
Ground Truth	PT	0.0	0.001	0.02
	Resting	0.008	0.08	0.16
	<b>Ours</b>	<b>0.21</b>	<b>0.44</b>	<b>0.56</b>

Table 2. **Distance ranks.** Top: distance to prediction. Bottom: distance to GT pose. We compare our results to the Progressive Transformers (PT), and to a sequence of constant resting positions.

Ref.	Lang.	Rank 1 $\uparrow$	Rank 5 $\uparrow$	Rank 10 $\uparrow$
Pred.	PJM	0.03 / 0	0.15 / 0.005	0.27 / 0.02
	DGS	0.14 / 0.02	0.23 / 0.08	0.39 / 0.14
	GSL	0 / 0.01	0.1 / 0.06	0.25 / 0.14
	LSF	0.22 / 0.18	0.6 / 0.41	0.8 / 0.58
Ground Truth	PJM	0.08 / 0.06	0.25 / 0.3	0.41 / 0.5
	DGS	0.26 / 0.21	0.49 / 0.49	0.58 / 0.64
	GSL	0.27 / 0.53	0.67 / 0.84	0.8 / 0.92
	LSF	0.7 / 0.75	0.95 / 0.91	0.97 / 0.95

Table 3. **Distance ranks by language** (full / leave-one-out model). Top: distance to prediction, bottom: distance to ground truth.

Progressive Transformers (PT) [41], which aims to translate glosses into pose sequences. We adjust their model to take HamNoSys sequences as input instead of glosses and train it over our data. We also present the results of using  $s_T$  of each GT (*i.e.* a sequence of the first pose frame of the GT in the length predicted by the sequence length predictor) as the prediction for comparison. As shown, our model outperforms both PT and the “resting” option in both settings.

#### 7.4. Leave One Out

To check if our method truly is generic, we perform a “leave-one-out” experiment: we train our model with all languages but one and test it on the left-out language. Then, we report rank 1, 5, and 10 results for each model in Tab. 3. For comparison, we also show the rank results of the full model per language. For the leave-one-out experiment, all ‘other’ samples are taken from the left-out language. As demonstrated, the results when testing on unseen languages are only slightly worse than the results of the full model, and in some cases, they are even better. The degradation in results might be due to insufficient data or rare glyphs that only appear in one language.

#### 7.5. Qualitative Results

We present results in Fig. 3 and in the supplementary material. Although some frames of the GT pose often have missing or incorrect keypoints, our model generates a complete and correct pose.




Figure 6. **Hand location failure case.** Note that the pointing finger is correct in our prediction, not in the ground truth pose.

## 8. Limitations and Future Work

Despite making a big step toward SLP, much work still needs to be done. Predicted hand shapes or movements are not always entirely correct, which we attribute to missing or incorrect keypoints in our data. The pose estimation quality is a significant bottleneck, making it almost impossible for the model to learn the correct meanings of some glyphs. Figs. 3 and 6 show examples of missing, incorrect keypoints. Future work may improve hand pose estimation to address this issue. Moreover, our model generates some movements correctly but not exactly in the right location due to local proximity (*e.g.* hand over mouth instead of chin example in Fig. 6). Note: the pointing finger is correct in our prediction, not in the GT pose). We present more failure examples in the supplementary material. Finally, more annotated data is needed. Some glyphs are rare and only appear once or a handful of times in our dataset, making them difficult, if not impossible, for the model to learn.

## 9. Conclusion

In this work, we propose the first method for animating HamNoSys into pose sequences. As demonstrated, our model can generate signs even when trained on partial and inaccurate data. Additionally, we introduce a new distance function, considering missing keypoints, for measuring the distance between two pose sequences using *DTW-MJE* over normalized pose keypoints. We validate its correctness using a large-scale sign language dataset and show that it better suits pose sequences evaluation than existing methods. We hope our method leads the way toward developing an end-to-end system for SLP, that will allow communication between the hearing and hearing-impaired communities.

## Acknowledgements

This project was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No. 802774 (iEXTRACT), and by the Israel Science Foundation (grant No. 1574/21).



## References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, pages 719–728. IEEE, 2019. 3, 7
- [2] Samuel Albanie, Gül Varol, Liliane Momeni, Triantafyllos Afouras, Joon Son Chung, Neil Fox, and Andrew Zisserman. Bsl-1k: Scaling up co-articulated sign language recognition using mouthing cues. In *European conference on computer vision*, pages 35–53. Springer, 2020. 1
- [3] J Andrew Bangham, SJ Cox, Ralph Elliott, John RW Glauert, Ian Marshall, Sanja Rankov, and Mark Wells. Virtual signing: Capture, animation, storage and transmission-an overview of the visicast project. In *IEE Seminar on speech and language processing for disabled and elderly people (Ref. No. 2000/025)*, pages 6–1. IET, 2000. 1, 3
- [4] Ingo Barth. sign2mint. Retrieved online at: <https://sign2mint.de/>, 2021. 2
- [5] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7784–7793, 2018. 1
- [6] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10023–10033, 2020. 1
- [7] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 3
- [8] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5933–5942, 2019. 1
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. 3
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 3
- [11] Amanda Duarte, Shruti Palaskar, Lucas Ventura, Deepti Ghadiyaram, Kenneth DeHaan, Florian Metze, Jordi Torres, and Xavier Giro-i Nieto. How2sign: A large-scale multi-modal dataset for continuous american sign language. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2735–2744, 2021. 2
- [12] Sarah Ebling and John Glauert. Building a swiss german sign language avatar with jasingning and evaluating it among the deaf community. *Universal Access in the Information Society*, 15(4):577–587, 2016. 1, 3
- [13] Eleni Efthimiou, Stavroula-Evita Fontinea, Thomas Hanke, John Glauert, Rihard Bowden, Annelies Braffort, Christophe Collet, Petros Maragos, and François Goudenove. Dicta-sign–sign language recognition, generation and modelling: a research effort with applications in deaf communication. In *Proceedings of the 4th Workshop on the Representation and Processing of Sign Languages: Corpora and Sign Language Technologies*, pages 80–83, 2010. 1, 3
- [14] Jens Forster, Christoph Schmidt, Thomas Hoyoux, Oscar Koller, Uwe Zelle, Justus Piater, and Hermann Ney. Rwth-phenix-weather: A large vocabulary sign language recognition and translation corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 3785–3789, 2012. 2
- [15] Anindita Ghosh, Noshaba Cheema, Cennet Oguz, Christian Theobalt, and Philipp Slusallek. Synthesis of compositional animations from textual descriptions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1396–1406, 2021. 3, 7
- [16] Sylvie Gibet, François Lefebvre-Albaret, Ludovic Hamon, Rémi Brun, and Ahmed Turki. Interactive editing in french sign language dedicated to virtual signers: Requirements and challenges. *Universal Access in the Information Society*, 15(4):525–539, 2016. 3
- [17] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. 3
- [18] Thomas Hanke. Hamnosys-representing sign language data in language resources and language processing contexts. In *LREC*, volume 4, pages 1–6, 2004. 2
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 3
- [20] Wencan Huang, Wenwen Pan, Zhou Zhao, and Qi Tian. Towards fast and high-quality sign language production. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3172–3181, 2021. 2, 7
- [21] Zifan Jiang, Amit Moryossef, Mathias Müller, and Sarah Ebling. Machine translation between spoken languages and signed languages represented in signwriting, 2022. 1
- [22] Lucas Kovar, Michael Gleicher, and Frédéric H. Pighin. Motion graphs. In *SIGGRAPH ’08*, pages 1–10, 2008. 3
- [23] Joseph B Kruskal. An overview of sequence comparison: Time warps, string edits, and macromolecules. *SIAM review*, 25(2):201–237, 1983. 7
- [24] Dongxu Li, Cristian Rodriguez, Xin Yu, and Hongdong Li. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1459–1469, 2020. 2
- [25] Jadwiga Linde-Usiekiewicz, Małgorzata Czajkowska-Kisil, Joanna Łacheta, and Paweł Rutkowski. A corpus-based dictionary of polish sign language (pjm). In *Proceedings of the XVI EURALEX International Congress: The user in focus*, pages 365–376, 2014. 2, 3

- [26] Silke Matthes, Thomas Hanke, Anja Regen, Jakob Storz, Satu Worseck, Eleni Efthimiou, Athanasia-Lida Dimou, Annelies Braffort, John Glauert, and Eva Safar. Dicta-sign-building a multilingual sign language corpus. In *5th Workshop on the Representation and Processing of Sign Languages: Interactions between Corpus and Lexicon. Satellite Workshop to the eighth International Conference on Language Resources and Evaluation (LREC-2012)*, 2012. 2, 3
- [27] Masahiro Mori, Karl F MacDorman, and Norri Kageki. The uncanny valley [from the field]. *IEEE Robotics & automation magazine*, 19(2):98–100, 2012. 3
- [28] Amit Moryossef and Mathias Müller. pose-format: Library for viewing, augmenting, and handling .pose files. <https://github.com/AmitMY/pose-format>, 2021. 4
- [29] United Nations. International day of sign languages. <https://www.un.org/en/observances/sign-languages-day>, 2022. 1
- [30] Maria Parelli, Katerina Papadimitriou, Gerasimos Potamianos, Georgios Pavlakos, and Petros Maragos. Exploiting 3d hand pose estimation in deep learning-based sign language recognition from rgb videos. In *European Conference on Computer Vision*, pages 249–263. Springer, 2020. 1
- [31] Mathis Petrovich, Michael J. Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *ECCV*, 2022. 3, 7
- [32] Siegmund Prillwitz, Thomas Hanke, Susanne König, Reiner Konrad, Gabriele Langer, and Arvid Schwarz. Dgs corpus project—development of a corpus based electronic dictionary german sign language/german. In *sign-lang@ LREC 2008*, pages 159–164. European Language Resources Association (ELRA), 2008. 2, 3
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 3
- [34] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Swish: a self-gated activation function. *arXiv preprint arXiv:1710.05941*, 7(1):5, 2017. 6
- [35] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022. 3
- [36] Razieh Rastgoo, Kourosh Kiani, and Sergio Escalera. Hand sign language recognition using multi-view hand skeleton. *Expert Systems with Applications*, 150:113336, 2020. 1
- [37] Razieh Rastgoo, Kourosh Kiani, and Sergio Escalera. Real-time isolated hand sign language recognition using deep networks and svd. *Journal of Ambient Intelligence and Humanized Computing*, 13(1):591–611, 2022. 1
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Seyedeh Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *ArXiv*, abs/2205.11487, 2022. 3
- [39] Stan Salvador and Philip Chan. Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, 11(5):561–580, 2007. 7
- [40] Wendy Sandler and Diane Lillo-Martin. *Sign language and linguistic universals*. Cambridge University Press, 2006. 1
- [41] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. Progressive transformers for end-to-end sign language production. In *European Conference on Computer Vision*, pages 687–705. Springer, 2020. 1, 2, 3, 8
- [42] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. Mixed signals: Sign language production via a mixture of motion primitives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1919–1929, 2021. 1, 3
- [43] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. Signing at scale: Learning to co-articulate signs for large-scale photo-realistic sign language production. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5141–5151, 2022. 1
- [44] Ozge Mercanoglu Sincan and Hacer Yalim Keles. Autsl: A large scale multi-modal turkish sign language dataset and baseline methods. *IEEE Access*, 8:181340–181355, 2020. 2, 7
- [45] Victor Skobov and Yves Lepage. Video-to-hamnosys automated annotation system. In *sign-lang@ LREC 2020*, pages 209–216. European Language Resources Association (ELRA), 2020. 3
- [46] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [47] Stephanie Stoll, Necati Cihan Camgöz, Simon Hadfield, and Richard Bowden. Sign language production using neural machine translation and generative adversarial networks. In *Proceedings of the 29th British Machine Vision Conference (BMVC 2018)*. British Machine Vision Association, 2018. 1, 2
- [48] Stephanie Stoll, Necati Cihan Camgoz, Simon Hadfield, and Richard Bowden. Text2sign: towards sign language production using neural machine translation and generative adversarial networks. *International Journal of Computer Vision*, 128(4):891–908, 2020. 1, 2, 3
- [49] Valerie Sutton. Signwriting. Retrieved online at: <http://www.signwriting.org/about/what/what02.html>, 1974. 2
- [50] Valerie Sutton. Signbank. Retrieved online at: <https://www.signbank.org>, 2002. 2
- [51] Valerie Sutton. Asl wikipedia. Retrieved online at: <https://incubator.wikimedia.org/wiki/Wp/ase>, 2015. 2
- [52] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H. Bermano. Human motion diffusion model. *ArXiv*, abs/2209.14916, 2022. 3
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4, 5, 6

- [54] Harry Walsh, Ben Saunders, and Richard Bowden. Changing the representation: Examining language representation for neural sign language production. In *Proceedings of the 7th International Workshop on Sign Language Translation and Avatar Technology: The Junction of the Visual and the Textual: Challenges and Perspectives*, pages 117–124, 2022. [1](#), [3](#)
- [55] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. [1](#)
- [56] WHO. Deafness and hearing loss. <https://www.who.int/news-room/fact-sheets/detail/deafness-and-hearing-loss>, 2021. [1](#)
- [57] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989. [6](#)
- [58] Inge Zwitterlood, Margriet Verlinden, Johan Ros, Sanny Van Der Schoot, and T Netherlands. Synthetic signing for the deaf: Esign. In *Proceedings of the conference and workshop on assistive technologies for vision and hearing impairment (CVHI)*. Citeseer, 2004. [1](#), [3](#)