

# HP-GAN: Probabilistic 3D human motion prediction via GAN

Emad Barsoum      John Kender  
Columbia University

116th St & Broadway, New York, NY 10027

eb2871@columbia.edu

jrk@cs.columbia.edu

Zicheng Liu  
Microsoft

One Microsoft Way, WA 98052

zliu@microsoft.com

## Abstract

*Predicting and understanding human motion dynamics has many applications, such as motion synthesis, augmented reality, security, and autonomous vehicles. Due to the recent success of generative adversarial networks (GAN), there has been much interest in probabilistic estimation and synthetic data generation using deep neural network architectures and learning algorithms.*

*We propose a novel sequence-to-sequence model for probabilistic human motion prediction, trained with a modified version of improved Wasserstein generative adversarial networks (WGAN-GP), in which we use a custom loss function designed for human motion prediction. Our model, which we call HP-GAN, learns a probability density function of future human poses conditioned on previous poses. It predicts multiple sequences of possible future human poses, each from the same input sequence but a different vector  $z$  drawn from a random distribution. Furthermore, to quantify the quality of the non-deterministic predictions, we simultaneously train a motion-quality-assessment model that learns the probability that a given skeleton sequence is a real human motion.*

*We test our algorithm on two of the largest skeleton datasets: NTURGB-D and Human3.6M. We train our model on both single and multiple action types. Its predictive power for long-term motion estimation is demonstrated by generating multiple plausible futures of more than 30 frames from just 10 frames of input. We show that most sequences generated from the same input have more than 50% probabilities of being judged as a real human sequence. We published all the code used in this paper to <https://github.com/ebarsoum/hpgan>.*

## 1. Introduction

Accurate short-term (several second) predictions of what will happen in the world given past events is a fundamental and useful human ability. Such aptitude is vital for daily activities, social interactions, and ultimately survival. For

example, driving requires predicting other cars' and pedestrians' motions in order to avoid an accident; handshaking requires predicting the location of the other person's hand; and playing sports requires predicting other players' reactions. In order to create a machine that can interact seamlessly with the world, it needs a similar ability to understand the dynamics of the human world, and to predict probable futures based on learned history and the immediate present.

However, the future is not deterministic, so predicting the future cannot be deterministic, except in the very short term. As the predictions extend further into the future, uncertainty becomes higher. People walking may turn or fall; people throwing a ball may drop it instead. However, some predictions are more plausible than others, and have a higher probability.

In this paper, we focus on creating a model that can predict multiple plausible future human (skeleton) poses from a given past. The number of poses taken from the immediate past, and the predicted number of poses in the future, which can be unrestricted, are parameters to the model. To accomplish this, we use a modified version of the improved Wasserstein generative adversarial network (WGAN-GP)[10] with a custom loss function that takes into consideration human motion and human anatomy.

The generator is a novel adaptation of sequence-to-sequence model[36] of poses derived from a Recurrent Neural Network (RNN), and the critic and discriminator are a multilayer network (MLP). We use the critic network to train the generator and the discriminator network to learn distinguishing between a real sequence of poses from a fake one. In essence, we combine some aspect of the original GAN [8] with WGAN-GP[10]. We train our model on all actions at once, so its output is not conditioned on any specific human action. Our model takes as input a sequence of previous skeletal poses, plus a random vector  $z$  from the reduced sequence space which samples possible future poses. For each such  $z$  value, the model generates a different output sequence of possible future poses.

We use an RNN for the generator because RNNs are a class of neural networks designed to model sequences, es-

pecially variable length sequences. They have been successfully used in machine translation [36], caption generation from images [17, 6], video classification and action recognition[2, 26, 6], action detection[42], video description[27, 19, 41, 6], sequence prediction[9, 35] and others.

We structure the learning by using a GAN because GANs[8] are a class of unsupervised learning algorithms, inspired by game theory[34], which allow the generation of futures that are not tied to specific ground truth. Among other domains, they have been used to generate impressive realistic images. However, GANs have weaknesses. They can be difficult to train and unstable in their learning, their loss value does not necessarily indicate the quality of the generated sample, and the training can collapse easily. Recent literature[1, 10, 25, 37, 29, 22] tries to improve GAN training and provide a theoretical guaranty for its convergence. In our work, we address this by adding a custom loss based on the skeleton physics in addition to the GAN loss, in order to stabilize and improve the training.

To quantitatively assess the quality of the non-deterministic predictions, we simultaneously train a motion-quality-assessment model that learns the probability that a given skeleton sequence is a real human motion.

We test our motion prediction model on two large datasets each captured with a different modality. The first is the NTURGB-D[32] dataset, which is the largest available RGB-D and skeleton-based dataset, with data captured by using the Microsoft Kinect v2 sensor. The second is the Human3.6M[15, 14] dataset, which is one of the largest available datasets derived from motion capture (MoCap) data.

Our main contributions are:

1. We propose a novel human motion model that can predict multiple possible futures from a single past.
2. We propose a motion-quality-assessment model to quantitatively evaluate the quality of the predicted human motions.

## 2. Related work on prediction

Since the introduction of the Kinect sensor, there has been much work on recognizing human action and predicting human poses from skeleton data. For example, predicting human poses conditioned on previous poses using deep RNNs [7, 16, 23] is due in part to this availability of large human motion datasets [15, 14, 32]. In general, human motion prediction can be categorized into two categories: probabilistic and deterministic prediction.

### 2.1. Probabilistic motion prediction

Most work in probabilistic human motion prediction uses non-deep learning approaches [28, 33, 38, 39, 20, 21, 7]. In [21], the authors propose simple Markov models

that model observed data, and use the proposed model for action recognition and task completion. The limitation in this approach is that it predicts motion from a single action only, and assumes that human motion satisfies the Markov assumption. In [39], the authors introduce Gaussian process dynamical models (GPDMs) to model human pose and motion. However, they train their model on each action separately, and constrain the model to a Gaussian process. In [33], the authors map human motion to a low dimensional space, and use the coordinates as an index into a binary tree that predicts the next pose, thus casting the prediction problem into a search problem. However, this approach can not generalize to previously unseen motions. In [28], the authors use switching linear dynamic systems learned through a Bayesian network, and in [20] the author used conditional random fields (CRF) to model spatio-temporal dynamics.

In contrast to all the above, our work does not use any statistical models to constrain the motion. As far as we know, we are the first to use deep neural networks for probabilistic motion prediction.

### 2.2. Deterministic motion prediction

Recent human motion prediction, which relies on deep RNNs [7, 16, 23] or deep neural networks [3], is primarily deterministic. In [7], the authors mix both deterministic and probabilistic human motion predictions. Their deterministic aspect is based on a modified RNN called Recurrent-Decoder (ERD) that adds fully connected layers before and after an LSTM [12] layer and minimizes an Euclidean loss. Their probabilistic aspect uses a Gaussian Mixture Model (GMM) with five mixture components and minimizes the GMM negative log-likelihood. For both aspects, they predict a single future human pose at a time. To predict more, they recurrently feed the single predicted pose back to the input. One drawback of this approach is error drifting, where the prediction error of the current pose will propagate into the next pose. In contrast, we predict multiple human poses at once thus avoiding error drifting. In addition, we do not impose any statistical model constrains like GMM over the motions.

In [16], the authors develop a general framework that converts a structure graph to an RNN, called a Structure-RNN (S-RNN). They test their framework on different problem sets including human motion prediction, and show that it outperforms the current state of the art. However, they need to design the structure graph manually and task-specifically. In [23], the authors examine recent deep RNN methods for human motion prediction, and show that they achieve start-of-the-art results with a simpler model by proposing three simple changes to RNN. In [3], the authors use an encoder-decoder network based on a feed-forward network, and compare the results of three different such architectures: symmetric, time-scale, and hierarchical.

However, the main issues of deterministic prediction of human motion are two-fold. The future is not deterministic, so the same previous poses could lead to multiple possible poses. And, using an  $L_2$  norm can cause the model to average between two possible futures [24], resulting in blurred motions.

### 2.3. Non-human motion prediction

The prediction of multiple possible futures using RNNs has precedents. In [9], the authors use an LSTM [12] to generate text and handwriting from an input sequence. They generate one item at a time, by sampling the resultant probability. Then, they append the predicted item to the input sequence and remove its first item, and iterate. This creates a sequence of desired length, but the method may eventually create an input to the LSTM that does not contain any of the original input. In contrast, we pursue a method that trains our model to generate the entire desired output sequence of poses all at once.

The prediction of single or multiple possible futures using GANs also has precedents. In [24], the authors trained a convolution model for both the generator and the discriminator in order to predict future frames. They improved the predictions by adding an image gradient difference loss to the adversary loss. However, they again only predict a single possible future and, due to the use of the convolution network, the model can only predict a fixed length output. In contrast, we support variable length input and variable length output, and can also generate multiple possible futures from the same input.

In [4], the authors predict or imagine multiple frames from a single image. They generate affine transformations between each frame, and apply them to the original input image to produce their prediction. Although this can imagine multiple futures from the same input image, a single image is not sufficient to capture the temporal dynamics of a scene. Furthermore, it makes the oversimplified assumption that the change between the images can be captured by using an affine transform.

### 2.4. Human motion quality assessment

Compared to the amount of work on motion editing and synthesis in computer graphics, the research on automatic motion quality evaluation has received little attention [31, 11, 13, 30]. The existing techniques were typically designed for special types of motions, or for motions obtained from software used to edit character motion [40]. A novel aspect of our motion quality assessment model is that it is trained simultaneously together with the motion prediction model.

## 3. Our model for human motion prediction

In the human motion prediction problem, the system takes a sequence of human poses as input and predicts valid

future poses. Let  $x = \{x_1, x_2, \dots, x_m\}$  be the sequence input poses and  $z = \{z_1, z_2, \dots, z_n\}$  be the sequence of predicted poses, where each  $z_i$  and  $x_j$  corresponds to a single pose, represented by joint locations or joint angles. The goal is to learn the probability of the future sequence conditioned on the input sequence,  $P(z|x)$ .

As shown in Figure 1, our prediction model is a modified version of the sequence-to-sequence [36] network. It takes as input a sequence of human poses, plus a  $z$  vector drawn from a uniform or Gaussian distribution  $z \sim p_z$ . The drawn  $z$  value is then mapped to the same space as the output states of the encoder. We then simply add the mapped value of  $z$  to the encoder states, and use the result as the initial state of the decoder. As shown in the figure, we map  $z$  to each layer of the encoder, and then feed the last output of the encoder to the first input of the decoder. We use GRU [5] for our sequence-to-sequence network, we also tried LSTM [12] and did not notice any quality difference.

Let  $G$  be the network shown in figure 1. We have  $y = G(x, z; \theta_g)$ , where  $\theta_g$  are the network parameters that we need to learn. Each drawn value of  $z$  will sample different valid future poses from the given input  $x$ . Most recent human pose predictions using deep RNNs [7, 23] treat human motion prediction as a regression problem. However, solving human motion prediction using regression suffers from two deficiencies. First, it learns one outcome at a time, and as the predicted sequence length increases, this outcome becomes less probable. Second, using the usual  $\ell_1$  or  $\ell_2$  loss creates artifacts [24], such as predicting the average of two possible outcomes.

We believe GAN is a better way to learn multiple possible future poses from the same input poses. In particular, we use a modified version of the improved WGAN-GP [10] tailored for sequence prediction, in order to train  $y = G(x, z; \theta_g)$ .

### 3.1. Generative adversarial networks

Generative adversarial networks (GAN) was introduced by [8]. It is a unsupervised learning technique inspired by the minimax theorem [34], in which the generator network and the discriminator network try to outdo each other. The training itself alternates between both networks. In the original paper, the generator learns to generate images close to real images, and the discriminator learns to distinguish between the generated image and the real image from the dataset. In a steady state, the discriminator should predict if an image from the generator network is generated or not with 50% probability.

However, the original GAN algorithm is not stable and is difficult to train, because of its use of Jensen-Shannon (JS) divergence as its loss function. JS can result in zero or infinity due to its ratio between two probabilities that might not overlap initially, and can lead to vanishing gradients in

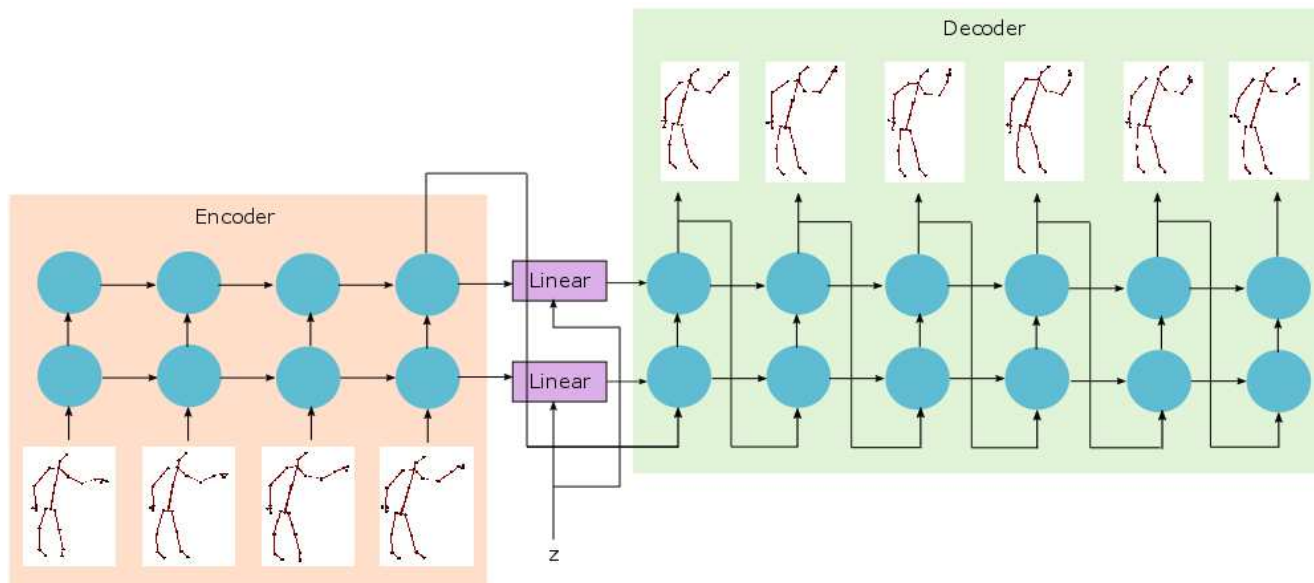


Figure 1. Sequence-to-sequence generator network. It predicts multiple future sequences from the same input by feeding in different  $z$  values.

the discriminator network. WGAN [1] replaces the JS distance with the Earth Mover Distance (EMD), which is defined and continuous almost everywhere. And according to the author, this mitigates the need to carefully maintain a balance between training the discriminator versus the generator. The discriminator in WGAN does not output a probability, and it does not discriminate between synthetic input and real input, which is why the author renamed the discriminator network to critic network.

Nevertheless, WGAN does not address all the concerns, since the critic still must maintain a Lipschitz constraint. In order to do so, the author clips the weights of the network, which adversely affects the quality of the generator. To address this, the WGAN-GP algorithm [10] replaces weight clipping in the generator with an added penalty to the loss in the critic, which is based on the computed norm of the gradient with respect to the critic input. We were able to verify these improvements in our domain of human poses through extensive experimentation.

Therefore, we use an adversarial training scheme for three reasons. First, it allows the generation of multiple futures from a single past. Second, it allows the generator to be trained without explicitly using the ground truth of real futures. Third, it implicitly learns the cost function for the prediction based on the data.

### 3.2. Human pose prediction GAN (HP-GAN)

Figure 2 shows the high level diagram of our proposed GAN network, called HP-GAN, for Human Pose prediction. HP-GAN combines features from WGAN-GP [10], from GAN [8], and from sequence-specific optimization, in

order to generate a realistic human motion sequence, and at the same time to quantify the quality of the generated sequence.

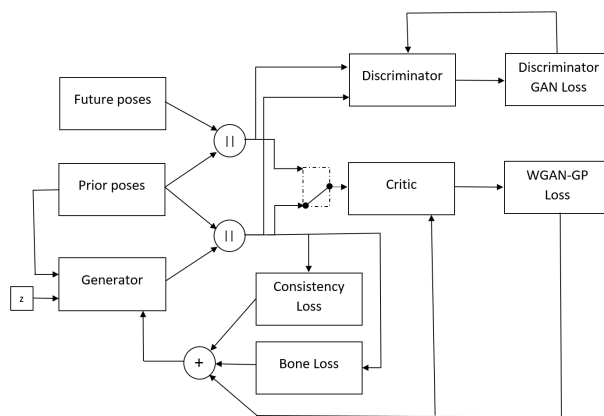


Figure 2. HP-GAN: Human Prediction GAN. The above diagram shows the critic switching back and forth between generator and critic training. The output of the generator and the prior poses are concatenated for the critic and discriminator. Similarly, ground truth is the concatenation of the prior and future poses. The only job of the discriminator is to learn to distinguish between real and fake sequences of human motion. To update the generator, we use WGAN-GP loss, in addition to consistency loss and bone loss.

The "Generator" block shown in Figure 2 is the sequence-to-sequence network defined in Figure 1. As described earlier, it takes as input previous poses and a  $z$  vector, and produces a sequence of human poses. The  $z$  vector is a 128-dimensional float vector drawn from a uniform



or Gaussian probability distribution. "Future poses" are the ground truth future poses from the dataset, and "Prior poses" are their corresponding previous poses.

Prior poses and future poses are concatenated together to form a real pose sequence. Similarly, prior poses and generated poses are concatenated together to form a fake pose sequence. Both real and fake sequences are used for the critic WGAN-GP loss and for the discriminator GAN loss. WGAN-GP loss is used to improve both the critic network and the generator network by alternating between critic loss and generator loss, whereas the discriminator GAN loss is used only to train the discriminator network. The main purpose of the discriminator network is to measure the quality of the generated human poses; it is not involved in training the generator.

The critic network is a three-layer fully connected feed-forward network that outputs a single value. This value is unbounded and is used for the WGAN-GP loss. This WGAN-GP loss is the same loss as in [10]. To train the generator to produce a realistic human pose, we add two additional losses to the WGAN-GP loss. The first one is the consistency, or pose gradient loss, which focuses on smoothing the sequence of predicted poses. The second loss is the bone loss, which focuses on reducing the changes to the bone lengths between the predicted skeleton and the ground truth. The details of each of the losses are described below.

The discriminator network is also a three-layer fully connected feedforward network that outputs a single value. However, the output is a probability value between 0 to 1, where 1 means that the input sequence is real and 0 means that the input sequence is fake. The primary goal of the discriminator network is to measure the quality of the predicted future sequence from the generator. The probability resulting from the discriminator is not used to improve the generator, in order to keep training the generator with WGAN-GP loss, for stability and the avoidance of model collapse.

### 3.2.1 Critic loss

Our critic loss is similar to WGAN-GP critic loss, with an added L2 regularization. It is defined as follows:

$$L_c = L_{wgan} + \lambda L_{gp} + \alpha L_2 \quad (1)$$

$L_{wgan}$  is the original WGAN critic loss, which is defined as:

$$L(x, y, z)_{wgan} = D(x||G(x, z)) - D(x||y) \quad (2)$$

where  $||$  indicates concatenation,  $x$  is the input sequence,  $y$  is the future sequence, and  $z$  is a random vector drawn from a uniform distribution.

$L_{gp}$  is the gradient penalty loss defined as:

$$L(x, y, z)_{gp} = (\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2 \quad (3)$$

where  $\hat{x} = \epsilon(x||y) + (1 - \epsilon)(x||G(x, z))$  and  $\epsilon \sim U[0, 1]$ .

$L_2$  is the standard  $L_2$  regularization defined as:

$$L_2 = \|\theta_d\|_2 \quad (4)$$

In all of our experiments, we set  $\gamma = 10$  and  $\alpha = 0.001$ .

### 3.2.2 Generator loss

For the generator loss, we use three components. The first is the standard WGAN-GP adversary loss. The second is a consistency loss or pose gradient loss, which measures the delta between two consecutive poses. The third is a "bone loss", which measures the delta of bone length between the predicted skeleton and the ground truth. The reason for both the consistency loss and the bone loss is that, in a valid human motion, the change in joint locations between two consecutive frames are small and the bone lengths should remain the same.

In summary, the generator network loss is as follows.

$$L_g = L_{adv} + \alpha L_{pg} + \beta L_b \quad (5)$$

where  $L_{adv}$  is the generator loss in WGAN-GP defined as:

$$L(x, z)_{adv} = -D(x||G(x, z)) \quad (6)$$

and  $L_{pg}$  is the consistency loss, or pose gradient loss, defined as:

$$L(x, z)_{pg} = \|\nabla_t y\|_p = \left[ \sum_t |y_t - y_{t-1}|^p \right]^{1/p} \quad (7)$$

In equation 7,  $\|\nabla_t y\|_p$  computes the gradient over time for the predicted sequence, we use  $p = 2$  in our training. And  $L_b$  is simply the  $L_2$  norm of the bone length differences between the predicted pose and the ground truth, that is,

$$L_b = \sum_t \left[ \sum_i |b_t^i - b_{gt}^i|^2 \right]^{1/2} \quad (8)$$

where  $b_{gt}^i$  is the ground truth bone length and  $b_t^i$  is the predicted bone length, both at time  $t$ . We iterate through all bones using index  $i$ , and sum over all the future skeleton poses using index  $t$ .

The effect of  $L_{pg}$  is critical during training. If  $\lambda$  is too large, the effect of  $L_{pg}$  loss becomes too high and we obtain less variety in future predictions; in some cases we even obtain identical predictions regardless of the input  $z$  value. Conversely, we have found that  $L_{pg}$  is not critical if we train

on a large dataset, like all the 60 classes of NTURGB-D; but for small subset, like one or two classes, then its value becomes critical in avoiding motion discontinuities at the first predicted human pose. Furthermore, to make sure that  $L_{pg}$  loss does not reach zero, we set a minimum value for the loss as follow  $L(x, z)_{pg} = \max(C, \|\nabla_x y\|_p)$ , where  $C$  is a hyper parameter, the reason is to avoid pushing two consecutive human poses to match exactly.

### 3.2.3 Discriminator loss

We use the discriminator to judge the quality of the predicted skeleton poses and to decide when to stop the training. The discriminator loss is the same loss used in GAN [8], with the exception that the generator does not use this loss in its training. It is defined as:

$$L_d = L_{GAN} + \alpha L_2 \quad (9)$$

where  $L_{GAN}$  is the standard GAN loss defined as:

$$L_{GAN} = \log(D(x||y)) + \log(1 - D(x||G(x, z))) \quad (10)$$

and  $L_2$  is the same  $L_2$  norm used by the critic network.

### 3.3. Training

For the training algorithm, we closely follow most GAN training methods. Inside the training loop, we iterate  $k$  times on the critic network, and one time on the generator and discriminator network. We use  $k = 10$ . We have tried different iteration values, and have tried to dynamically update the iteration count based on the losses of the critic and the generator, but none of those methods made any noticeable improvement.

In order to make the discriminator training procedure more stable, we reduced its learning rate by half, compared to the critic and generator learning rate. For all three networks, we use ADAM [18], and set the learning rate as  $5e - 5$  for the critic and generator network, and half of that for the discriminator network.

## 4. Experiments

To verify our model capability, we run multiple experiments on two of the largest human motion datasets: a Microsoft Kinect dataset NTURGB-D [32] and a motion capture (MoCap) dataset Human3.6M [15, 14]. The human poses in NTURGB-D dataset are based on skeleton data from Kinect, which are not perfect as shown in figure 8, and have problems due to occlusions, or to people carrying objects or interacting with another person. However, even with noisy skeletons our model generalizes well on this dataset.

The NTURGB-D action recognition dataset consists of 56,880 actions, and each action comes with the corresponding RGB video, depth map sequence, 3D skeletal data, and

infrared video. We use only the 3D skeleton data. They contain the 3D locations of 25 major body joints at each frame, as defined by the Microsoft Kinect API. NTURGB-D has 60 action classes and 40 different subjects, and each action was recorded by three Kinects from different viewpoints.

Human3.6M contains 3.6 million 3D human poses and their corresponding images, captured by a Vicon MoCap system. Each of these skeletons has 32 joints. The actions were performed by 11 professional actors covering 17 action classes. Using the code from [23], we read the Human3.6m skeleton data and converted it from its angle representation to absolute 3D joint positions.

So for both datasets, we trained our model directly on the absolute 3D joint locations. We fed our model a 3D point cloud, and from this training data the model learns the relationships between the joints in order to predict a valid human pose. This is more difficult than training on the angle, which has less degrees of freedom. We train directly on the joint positions in order to use the same pipeline for both NTU-RGB-D and Human3.6m datasets, and to have a more generic model.

### 4.1. Pre-processing

For preprocessing, we normalized the (x,y,z) values of each joint to the range [-1,1] and subtracted the center of gravity. In the NTURGB-D dataset, we normalized each joint to the range of [-1, 1] by using the dimensions of the Kinect frustum at its maximum 5 meter depth. For the Human3.6M dataset, we obtained the range of the raw data by first finding the minimum and maximum values on each dimension, and then computing the minimum and maximum values over all the dimensions.

The Human3.6M dataset has fewer clips than the NTURGB-D dataset, however each clip is much longer. In order to use the same pipeline for both datasets, we split Human3.6M clips into shorter segments and only use every other frame in our training.

### 4.2. Quantifying the results

One of the main problems of GANs is that the loss does not provide any indication of the quality of the generated data. According to the authors of WGAN [1] and WGAN-GP [10], one of the improvements that WGAN made on the original GAN is that their loss value does in fact provide a quality measure. In our human motion prediction problem, the loss provides some indication of how much the generated sequence looks like a valid human pose. However, we have observed that this is not strictly monotonic; a smaller loss does not always indicate a better quality.

Therefore, to more usefully measure quality, we added a discriminator network, whose sole purpose is to learn the probability that a given sequence is a valid human motion. To find the best model, inside the training loop we generate

$N$  predictions and compute the probability, by evaluating the discriminator on each of those predictions. Then, we compute the number of predictions  $k$  that has probability more than 50%. We keep track of the model with maximum  $k$  during training, we only start tracking the best model after certain number of epochs, in order to give the discriminator a chance to learn to discriminate between a real and a fake human motion.

## 5. Results

Figure 3 shows the result of training HP-GAN on all 60 classes of the NTURGB-D dataset. The top row is the ground truth, and each subsequent row corresponds to the predicted human poses from a different  $z$  value drawn from a uniform distribution. To the left of the blue line are the input sequences. As shown in the figure, each  $z$  value generates a separate possible future sequence of human poses. The first few of the predicted poses are very close to the ground truth, which is expected. As we predict more poses, they start to differ more from the ground truth. These results did not use the consistency loss, which is more important for smaller datasets or for training on a subset of the actions. We can see that there is no discontinuity between the last input pose and the first predicted pose.

In figure 4, we plot the loss values as a function of the number of epochs, for the critic (blue), generator (green), and discriminator (red) losses. Even though the generator loss continues improving, it is difficult to compute the best parameters for either generator loss or critic loss without a visual inspection of their graphs. The discriminator loss provides a more stable signal.

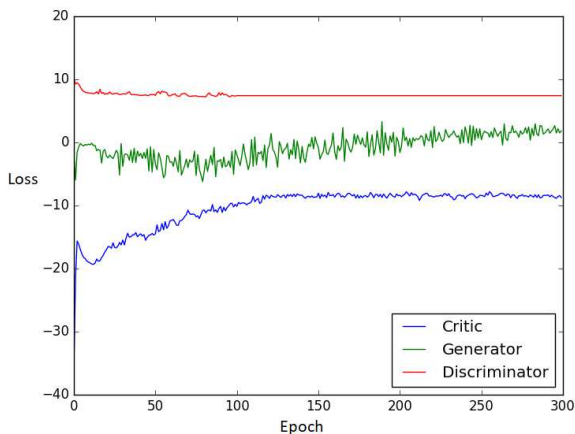


Figure 4. Critic versus Generator versus Discriminator loss.

Not all generated human poses are smooth, especially when we train on a subset of the actions. Figure 5 shows a discontinuity between the last pose of the input and the first pose of the prediction. This can be mitigated by increasing

the effect of the consistency loss in the generator loss, or by reducing the capacity of the network, or by feeding the actual last input directly into the first decoder node.

Training on the Human3.6M dataset requires different adjustments from training on NTURGB-D dataset. For example, bone loss is more important for the Human3.6M dataset than for the NTU-RGB-D dataset. Since NTU-RGB-D skeletons are generated from the Kinect, bone lengths between poses for the same subject do not match exactly. Also, Human3.6M actions are slower than NTU-RGB-D actions, so there is little variation between poses predictions even with different  $z$  values. To mitigate this, we only use every other frame. Figure 6 shows results on the Human3.6M dataset.




Figure 6. Human3.6M example. Top row is the ground truth. Next two rows: to left of line is the input sequence, to right of line are the predicted sequences from two different  $z$  values.

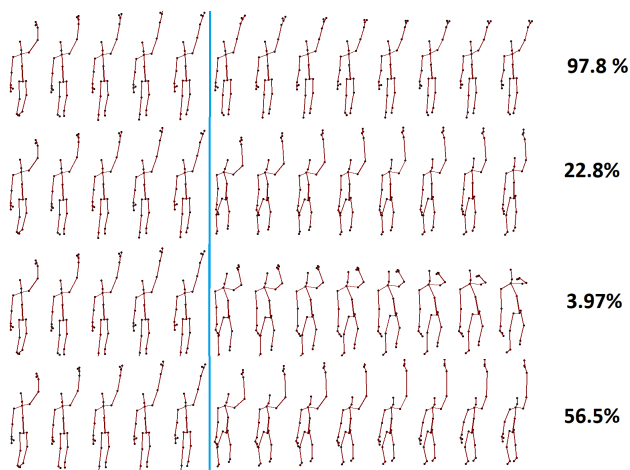


Figure 7. Sequences for the "throw" action, labeled with their corresponding probability from the discriminator network. Top row is ground truth.

Figure 7 shows predictions of the "throw" action from the NTURGB-D dataset and their corresponding estimated probabilities from the discriminator network. The top row is the ground truth, which shows a 97.8% probability of being a valid human sequence and each subsequent row correspond to a different  $z$  value. The discriminator probability output is not always accurate especially at the start of the training. We tried to use the discriminator in addition to the



Figure 3. Training on all 60 classes from the NTURGB-D dataset, and then predicting 30 poses from 10 input poses. Top row is the ground truth, and each subsequent row corresponds to a random  $z$  value drawn from a uniform distribution. Poses to the left of the blue line are the input sequences. Pose to the right of the blue line, starting in the second row, are the predicted future poses for each  $z$  value.




Figure 5. Discontinuity between input poses and predicted poses. Top row is ground truth and bottom row is the generated sequence.

critic to update the generator network parameters, however, the result was worse and the training was less stable.



Figure 8. Example of bad data from NTURGB-D dataset.

## 6. Conclusions and future work

We have developed a novel sequence-to-sequence model for probabilistic human motion prediction, which predicts multiple plausible future human poses from the same input. To our knowledge, we are the first to use deep neural networks for probabilistic motion predictions. To quantify the quality of the non-deterministic predictions, we simultaneously trained a motion-quality-assessment model that learns

the probability that a given skeleton sequence is a real human motion. We tested our architecture on two different datasets, one based on the Kinect sensor and the other based on MoCap data. Experiments show that our model performs well on both datasets.

However, using WGAN-GP and all the improvements we added, there still no good way to tell if the training has converged. Even worse, the training can diverge after it already converged if we continue the training loop. Therefore, in future work, we plan to explore better mechanism to measure model convergence and to ensure stability. Another area of research, is understanding the semantic and space of the  $z$  vector. If we can compute the reverse mapping from sequences to  $z$ , we might be able to use  $z$  values for action classification or clustering. Furthermore, HP-GAN generates a new dataset that is not in the original source, so an area of exploration is to use HP-GAN for data augmentation in order to increase diversity per action.

Lastly, we used standard GRU in our sequence-to-sequence network and we did not investigate new RNN architecture that might provide better human pose prediction such as RNN with skip connection [23] or Part-Aware LSTM [32] or something else, this is another area we plan to investigate in the future.



## References

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017. [2](#), [4](#), [6](#)
- [2] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt. Sequential deep learning for human action recognition. In *Human Behavior Understanding - Second International Workshop, HBU 2011, Amsterdam, The Netherlands, November 16, 2011. Proceedings*, pages 29–39, 2011. [2](#)
- [3] J. Bütepage, M. J. Black, D. Kragic, and H. Kjellström. Deep representation learning for human motion prediction and classification. *CoRR*, abs/1702.07486, 2017. [2](#)
- [4] B. Chen, W. Wang, J. Wang, and X. Chen. Video imagination from a single image with transformation generation. *CoRR*, abs/1706.04124, 2017. [3](#)
- [5] J. Chung, Ç. Gülçehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*, abs/1412.3555, 2014. [3](#)
- [6] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, T. Darrell, and K. Saenko. Long-term recurrent convolutional networks for visual recognition and description. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 2625–2634, 2015. [2](#)
- [7] K. Fragkiadaki, S. Levine, P. Felsen, and J. Malik. Recurrent network models for human dynamics. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 4346–4354, 2015. [2](#), [3](#)
- [8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, pages 2672–2680, 2014. [1](#), [2](#), [3](#), [4](#), [6](#)
- [9] A. Graves. Generating sequences with recurrent neural networks. *CoRR*, abs/1308.0850, 2013. [2](#), [3](#)
- [10] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. *CoRR*, abs/1704.00028, 2017. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [11] J. Harrison, R. A. Rensink, and M. van de Panne. Obscuring length changes during animated motion. *ACM Transactions on Graphics*, 23(3):569–573, 2004. [3](#)
- [12] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. [2](#), [3](#)
- [13] J. K. Hodgins, J. F. O’Brien, and J. Tumblin. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics*, 4(4):307–316, 1998. [3](#)
- [14] C. Ionescu, F. Li, and C. Sminchisescu. Latent structured models for human pose estimation. In *International Conference on Computer Vision*, 2011. [2](#), [6](#)
- [15] C. Ionescu, D. Papava, V. Olar, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, July 2014. [2](#), [6](#)
- [16] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 5308–5317, 2016. [2](#)
- [17] X. Jia, E. Gavves, B. Fernando, and T. Tuytelaars. Guiding the long-short term memory model for image caption generation. In *ICCV*, 2015. [2](#)
- [18] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. [6](#)
- [19] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *CoRR*, abs/1411.2539, 2014. [2](#)
- [20] H. S. Koppula and A. Saxena. Learning spatio-temporal structure from RGB-D videos for human activity detection and anticipation. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 792–800, 2013. [2](#)
- [21] A. M. Lehrmann, P. V. Gehler, and S. Nowozin. Efficient nonlinear markov models for human motion. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1314–1321, June 2014. [2](#)
- [22] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, and Z. Wang. Multi-class generative adversarial networks with the L2 loss function. *CoRR*, abs/1611.04076, 2016. [2](#)
- [23] J. Martinez, M. J. Black, and J. Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017. [2](#), [3](#), [6](#), [8](#)
- [24] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *CoRR*, abs/1511.05440, 2015. [3](#)
- [25] B. Neyshabur, S. Bhojanapalli, and A. Chakrabarti. Stabilizing GAN training with multiple random projections. *CoRR*, abs/1705.07831, 2017. [2](#)
- [26] J. Y. Ng, M. Hausknecht, S. Vijayanarasimhan, R. M. Oriol Vinyals and, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [2](#)
- [27] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 4594–4602, 2016. [2](#)
- [28] V. Pavlovic, J. M. Rehg, and J. MacCormick. Learning switching linear models of human motion. In *Proceedings of the 13th International Conference on Neural Information Processing Systems, NIPS’00*, pages 942–948, Cambridge, MA, USA, 2000. MIT Press. [2](#)
- [29] G. Qi. Loss-sensitive generative adversarial networks on lipschitz densities. *CoRR*, abs/1701.06264, 2017. [2](#)
- [30] P. S. A. Reitsma and N. S. Pollard. Perceptual metrics for character animation: sensitivity to errors in ballistic motion. *ACM Transactions on Graphics*, 22(3):537–542, 2003. [3](#)
- [31] L. Ren, A. Patrick, A. A. Efros, J. K. Hodgins, and J. M. Rehg. A data-driven approach to quantifying natural human motion. *ACM Transactions on Graphics*, 24(3):1090–1097, 2005. [3](#)

- [32] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [2](#), [6](#), [8](#)
- [33] H. Sidenbladh, M. J. Black, and L. Sigal. Implicit probabilistic models of human motion for synthesis and tracking. In *Proceedings of the 7th European Conference on Computer Vision-Part I, ECCV '02*, pages 784–800, London, UK, UK, 2002. Springer-Verlag. [2](#)
- [34] M. Sion. On general minimax theorems. *Pacific J. Math.*, 8(1):171–176, 1958. [2](#), [3](#)
- [35] N. Srivastava, E. Mansimov, and R. Salakhutdinov. Unsupervised learning of video representations using lstms. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 843–852, 2015. [2](#)
- [36] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112, 2014. [1](#), [2](#), [3](#)
- [37] M. Uehara, I. Sato, M. Suzuki, K. Nakayama, and Y. Matsuo. Generative adversarial nets from a density ratio estimation perspective. *CoRR*, abs/1610.02920, 2016. [2](#)
- [38] J. M. Wang, D. J. Fleet, and A. Hertzmann. Gaussian process dynamical models. In *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada]*, pages 1441–1448, 2005. [2](#)
- [39] J. M. Wang, D. J. Fleet, and A. Hertzmann. Gaussian process dynamical models for human motion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):283–298, 2008. [2](#)
- [40] X. Wang, Q. Chen, and W. Wang. 3d human motion editing and synthesis: A survey. *Computational and Mathematical methods in medicine*, 2014, 2014. [3](#)
- [41] L. Yao, A. Torabi, K. Cho, N. Ballas, C. J. Pal, H. Larochelle, and A. C. Courville. Video description generation incorporating spatio-temporal features and a soft-attention mechanism. *CoRR*, abs/1502.08029, 2015. [2](#)
- [42] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2678–2687, 2016. [2](#)