

Temporal Parameter-free Deep Skinning of Animated Meshes^{*}

Anastasia Moutafidou¹, Vasileios Toulatzis¹, and Ioannis Fudos¹

University of Ioannina, Greece

Abstract. In computer graphics, animation compression is essential for efficient storage, streaming and reproduction of animated meshes. Previous work has presented efficient techniques for compression by deriving skinning transformations and weights using clustering of vertices based on geometric features of vertices over time. In this work we present a novel approach that assigns vertices to bone-influenced clusters and derives weights using deep learning through a training set that consists of pairs of vertex trajectories (temporal vertex sequences) and the corresponding weights drawn from fully rigged animated characters. The approximation error of the resulting linear blend skinning scheme is significantly lower than the error of competent previous methods by producing at the same time a minimal number of bones. Furthermore, the optimal set of transformation and vertices is derived in fewer iterations due to the better initial positioning in the multidimensional variable space. Our method requires no parameters to be determined or tuned by the user during the entire process of compressing a mesh animation sequence.

Keywords: Animation, Skinning, Deep Learning

1 Introduction

Nowadays an animator may produce a realistic character animation by following either of the two modern workflows:

- i rigging a static mesh (i.e. define a bone structure and associate the bones with the mesh vertices by weight painting), apply transformations to bones along a time line, correct erroneous deformations by adding bones, introduce additional per frame deformations to simulate non linear effects, or
- ii use recent developments of computer vision and tracking techniques to derive mesh sequences that are reconstructed by markerless capture or by motion capture with dense markers (see e.g.[10]).

Both workflows produce sequences of animated meshes. These mesh animation sequences must subsequently be converted to a representation that allows for streaming and editing. To this end, a first step is to use compression.

^{*} This work has been co-financed by the European Union (European Regional Development Fund- ERDF) and Greek national funds through the Interreg Greece Albania 2014-2020 Program (project VirtuaLand).

With the evolution of cloud based graphics applications, a compression approach such as Linear Blend Skinning is a necessity for efficiently storing and using animation sequences. Compression is performed by producing an approximation of the animation that consists of an initial pose and a number of transformations that describe each subsequent pose by a deformation of a surface part.

Linear blend skinning (LBS) [24] is a time and space efficient mesh deformation technique where mesh vertices are influenced by a set of bones. In spite of several limitations that have been addressed in the literature [10], LBS-based approaches are significant in the animation industry due to their simplicity and straightforward GPU implementation.

There exists a variety of approaches for compression using clustering techniques, most of which are based on geometric features of vertices over time. We introduce a novel deep learning approach that uses a training set of successfully fully rigged animated models to produce a skinning model. Given a new animated mesh sequence, the trained network derives pseudo-bones and weights. There is no limit on the number of vertices, faces and frames given as input. There is just an upper limit on the number of bones for all animations. A single trained network can be used to predict weights for any animation sequence.

We also improve the efficiency of the least square optimization of transformations and weights that is commonly used to reduce the approximation error by employing conjugate gradient optimization that is suitable for multidimensional systems.

While previous skinning approaches use a predetermined number of bones and several other tuning parameters, our approach is parameter free. An appropriate set of bones is derived based on similar successfully rigged animations of the training dataset. In our method there is no need for preprocessing (scaling, rotation or translation) for the geometry of the input, since we only use vertex trajectories, so only the relative movement of the vertex is taken into account. To evaluate our approach, we use both mesh sequences that are derived from rigged and animated characters and benchmark mesh sequences from available animation sequence datasets. Our experimental evaluation shows that our approach outperforms previous approaches in terms of both compression rate and approximation error for all datasets.

The mesh clustering derived by our method can also be used to create a skeletal rig since it yields segmentations that correspond to the influence of bones on mesh vertices. Therefore, the output of our methods can be easily converted to a fully rigged animated character and used in subsequent phases of animation editing and rendering.

2 Related Work

Although there is a lot of research on skinning of animated models, the use of deep learning techniques for skinning has not been explored thoroughly.

Elastiface [30] indicate that an animated character can be quite complex and that managing and processing needs cumbersome human intervention and a significant amount of computationally intensive tasks. Moreover techniques such as cross-parameterization [17] or procedures that can convert an animated character into an animation sequence have high memory and space requirements. So there is a need for different procedures that can produce animated models in a compressed form without being provided with a skeleton or skin specification [11].

In the context of animation compression, James and Twigg [11] were the first to explore the use of LBS to approximately reproduce articulated characters as a function of their bone movement. Extending this work, Kavan et al. [13,12] presented a dual quaternion skinning scheme that can compute approximations for highly-deformable animations by suggesting that uniformly selected points on the mesh can act as bones. Both of them are enhance their final skinned approximation by exploiting EigenSkin [18]. FESAM [14] introduced an algorithm that optimizes all of the skinning parameters in an iterative manner. While FESAM offers high-quality reproduction results, is limited to download-and-play scenarios, since they do not use information about topology and the location of proxy joints is occluded once the optimization process kicks in.

[28] introduces a pose-to-pose skinning technique that exploits temporal coherence that enables the full spectrum of applications supported by previous approaches in conjunction with a novel pose editing of arbitrary animation poses, which can be smoothly propagated through the subsequent ones generating new deformed mesh sequences.

[19] approaches a set of example poses by defining a constrained optimization problem for deriving weights and transformations which yields better results in terms of error. In our method we ensure convexity by an additional equation for each vertex and a non negative least square solver. Then we employ linear solvers and update the weights and the bone transformations successively.

On the other hand, there are techniques that can create a plausible skeleton for a mesh model either by exploiting the movement of vertices to perform mesh segmentation [5], by exploiting the mesh structure by performing constrained Laplacian smoothing [2], or by analyzing the mesh structure of a set of several sparse example poses [8]. Recently, [21] presented a method that first produces a large number of plausible clusters, then reconstructs mesh topology by removing bones and finally performs an iterative optimization for joints, weights and bone transformations. Such methods are competent and produce a fully animated rigged object, but usually are semi-automated since their effectiveness and efficiency depends on setting a large number of parameters that are associated with the structure of the mesh or the specifics of kinematics.

[22] and [29] predicts a set of vertex weights based on the morphology of a static mesh, by previously training with static meshes and their corresponding weights from animated characters. [7] presents a method for automatically rigging a static mesh by matching the mesh against a set of morphable models. Our method predicts weights based on the vertex trajectories by training with the

motion of the vertices over time and the corresponding weights from animated characters and is not restricted by the morphology of the static mesh.

[4] tries to capture non linear deformations that are used in conjunction with a linear system and an underlying skeleton by employing a deep learning technique to determine the non linear part. [23] captures better nonlinear deformations by including in the animation pipeline a light weight neural network (NNWarp) that is known for its rich expressivity of nonlinear functions.

3 Temporal Deep Skinning

Skinning is based on the core idea that character skin vertices are deformed based on the motion of skeletal bones. One or more weights are assigned to each vertex that represent the percentage of influence vertices receive from each bone. With this approach we can reproduce an animation sequence based on a reference pose, the vertex weights and a set of transformations for every frame and bone.

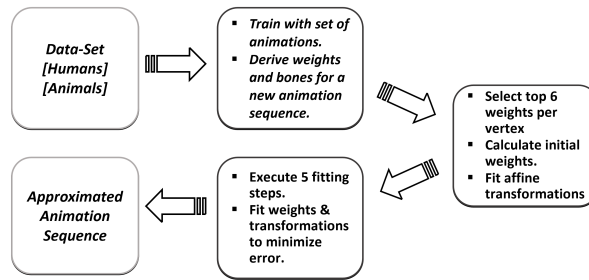


Fig. 1: Temporal Deep Skinning workflow.

Figure 1 illustrates the concept of our *Temporal Deep Skinning* method (or simply *Deep Skinning* for short). First, we build an appropriate neural network model that classifies each vertex by capturing mesh geometry and vertex kinematics. Then we use a set of human and animal animations to train the neural network model. We achieve this by using as input features the trajectories of all vertices and as output the weights that represent how each vertex is influenced by a bone. The output weight is conceived by the network as the probability of a bone to influence the corresponding vertex. Subsequently, we provide as input to our network arbitrary mesh animation sequences and predict their weights. From the per vertex classifier we determine the number of bones and the weights for each vertex.

We restrict each vertex to have no more than six weights so as to be compatible with the existing animation pipelines [13]. For simple gaming characters, usually four weights per vertex are enough, but six weights per vertex can be

used to correct artifacts or capture local deformations with pseudo bones. In our comparative evaluation we have implemented all previous methods with six weights as well, so as to conduct an objective comparison. The derived six (or less) weights per vertex correspond to the six highest probability predictions of the network. We then normalize these weights which are already in $[0, 1]$ to sum to 1 (coefficients of a convex combination). Since probability prediction of a vertex towards a specific bone cluster represents similarity to a training example, this is naturally translated to influence of the bone on the vertex.

Subsequently, we perform optimization to minimize the least square error between the original and approximated mesh frames. We do so by optimizing weights and transformations in an iterative manner.

3.1 Training and Test Datasets

The network is trained on a training set using a supervised learning method. Training dataset consists of input vector pairs that represent the motion of each vertex through all frames and the corresponding output vector of labels which determines whether a vertex is influenced by a specific bone. The input vector size is $(3 \cdot F)$, where 3 represents the x, y, z coordinates of a vertex and F the number of frames for the specific animation and the output consists of B_{max} labels, where B_{max} is the maximum number of bones. The current network model is fed with the training dataset and produces a result, which is then compared with the label vector, for each input vector in the training set. Based on the result of the comparison and the specific learning algorithm being used, the parameters of model are adjusted (supervised learning).

We have used two types of training datasets, one that consists of human character animations and one consisting of animal character animations. The animal dataset contains 32 animated animal characters with an average number of 12k vertices each, an average number of 3 animations per character and an average number of 195 frames per animation. The human dataset contains 35 animated human characters with an average number of 10k vertices each, 1 animation per character, and an average number of 158 frames per animation.

Successively, the fitted network model is used to predict the response of observations in a second smaller dataset called the validation set. This set provides an unbiased evaluation of the model and has been used to tune the hyperparameters of the network. Figure 2 indicates the average time that we need for training using LSTM or CNN networks (Section 3.4). After a complete training session, we export the trained network model so that we can use it in our Temporal Deep Skinning method to predict bones and weights for a given animation sequence.

For the test dataset that is used in the experiments we have used a set of human and animal models that are not included in the training datasets. The efficiency of our algorithm has been tested with more than 20 human and animal models to ensure generalization and accuracy. For example the testing dataset includes four animations Spider-man (27,030 Vertices & 28Frames), Man-Walking (15,918 Vertices & 32 Frames), Fox (1,728 Vertices & 400 Frames) and Lizard

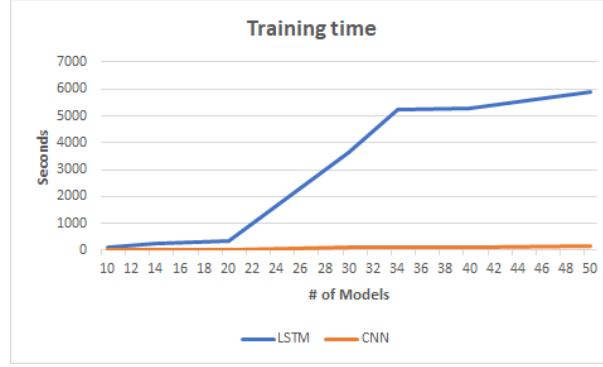


Fig. 2: Training time for LSTM & CNN networks (Section 3.4).

(29,614 Vertices & 75 Frames) . Note that all dataset models are extracted from FBX animations which means that are fully animated with skeletal rigs, skinning information and transformations. The skeletal information is only used for comparison with the outcome of our method.

3.2 Transformation and Weight Optimization

Approximating an animation sequence to produce a more succinct representation is common in the case of articulated models and is carried out through a process called *skinning*.

For every vertex v_i that is influenced by a bone j , a weight w_{ij} is assigned. For skeletal rigs the skeleton and skin of a mesh model is given in a predetermined pose also known as bind or rest pose. The rigging procedure blends the skeleton with skin which is given by the rest pose of the model. Each transformation is the concatenation of a *bind matrix* that takes the vertex into the local space of a given *bone* and a transformation matrix that moves the bone to a new position.

In LBS the new position of vertex v_i^p at pose (frame) p is given by Equation 1. This approach corresponds to using proxy bones instead of the traditional hierarchical bone structure on rigid or even on deformable models [13].

$$v_i^p = \sum_{j=1}^B w_{ij} \cdot T_j^p \cdot v_i \quad (1)$$

In this equation, v_i represents the position of the vertex in rest pose, w_{ij} the weight by which bone j influences vertex v_i and T_j^p is the transformation that is applied to bone j during frame p .

Figure 3 summarizes the successive weight and transformation optimization that aims at reducing the approximation error for all frames.

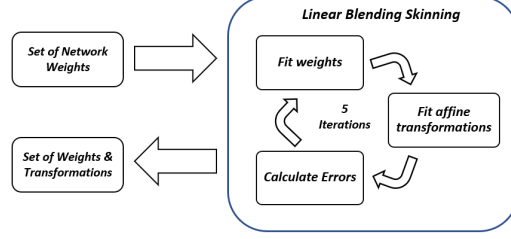


Fig. 3: Deep Skinning optimization workflow for weights and transformations.

Computing a good initial set of weights is a key step for the final result. In temporal deep skinning, a neural network provides the proxy bones and initial weights that are appropriate for an animation sequence. After that, we perform a successive optimization to find weights and proxy bone transformations. Both problems are formulated as least squares optimization problems that minimize the quantity given in Equation 2.

$$\sum_{i=1}^N \|v_i'^p - v_i^p\|^2 \quad (2)$$

where v_i^p denotes the coordinates of the original vertex in pose p , $v_i'^p$ is the approximation based on deep skinning and N is the number of vertices in the model. For the following, the number of vertices is N , the number of frames is P and the number of proxy bones is B . To solve the weight optimization problem, we formulate the system $A\mathbf{x} = b$, where matrix A is a $3PN \times 6N$ (where 6 is 6 the maximum bone number) matrix constructed by combining the rest-pose vertex positions and the corresponding transformations, \mathbf{x} is a $6N$ vector of unknowns that contains the weights and b is a known $3PN$ vector that consists of the original (target) vertex coordinates in all frames. In the case of finding the optimal weights we include the convexity coefficient requirement as an extra equation per vertex (so A becomes $(3P+1)N \times 6N$ and b becomes $(3P+1)N$).

Finally, to solve the transformation optimization problem we formulate a linear system that consists of $3N$ equations, the unknowns of which are the (3×4) elements of the transformation matrices T_j^p of each bone j and frame p . This sums to $12BP$ unknowns. The system can be expressed as $A\mathbf{x} = b$, where A is a $3N \times 12BP$ known matrix constructed by combining the rest-pose vertex positions and the corresponding vertex weights. Moreover b is a known $3N$ vector that contains the original (target) vertex coordinates.

To avoid reverting into non linear solvers we alternatively optimize weights and transformations separately. In terms of optimization [28] uses NNLS (non negative least square) optimization for enforcing the convexity condition of the

weights. We express the convexity by a separate equation per vertex which is closer to the approach adopted by [14]. [28] suggests that 5 iterations are enough, whereas [14] employs 15 iterations. We have performed experiments for up to 50 iterations and our conclusion is that after 5 iterations there is no significant error improvement. To perform the optimization problem we have employed conjugate gradient optimization which works better on multidimensional variable spaces and can be carried out efficiently on the GPU.

3.3 Measuring the Error

We used three different types of measures to calculate error of the approximation methods. The first two measures are standard measures used in [13], [11] and [14]. The first error measure is the percentage of deformation known as *distortion percentage* (*DisPer*).

$$DisPer = 100 \cdot \frac{\|A_{orig} - A_{Approx}\|_F}{\|A_{orig} - A_{avg}\|_F}. \quad (3)$$

where $\|\cdot\|_F$ is the Frobenius matrix metric. In Equation 3 A_{orig} is a $3NP$ matrix which consists of the real vertex coordinates in all frames of the model. Similarly, A_{Approx} has all the approximated vertex coordinates and matrix A_{avg} contains in each column, the average of the original coordinates in all frames. [14] replaces 100 by 1000 and divides by the surrounding sphere diameter. Sometimes this measure tends to be sensitive to the translation of the entire character, therefore we use a different measure that is invariant to translation. The *root mean square* (*ERMS*) error measure in Equation 4 is an alternative way to express distortion with the difference that we use $\sqrt{3NP}$ in the denominator so as to obtain the average deformation per vertex and frame during the sequence. $3NP$ is the total number of elements in the A_{orig} matrix. [21] uses as denominator the diameter of the bounding box multiplied by \sqrt{NP} .

$$ERMS = 100 \cdot \frac{\|A_{orig} - A_{Approx}\|_F}{\sqrt{3NP}} \quad (4)$$

We introduce a novel error measure, namely the *max average distance* (*MaxAvgDist*) given by Equation 5) which is a novel quality metric that reflects better the visual quality of the result. Max distance denotes the largest vertex error in every frame. So this measure represents the average of max distances over all frames.

$$MaxAvgDist = \frac{1}{P} \sum_{f=1}^P \max_{i=1, \dots, N} \|v_{orig}^{f,i} - v_{Approx}^{f,i}\| \quad (5)$$

Finally, we introduce an additional measure that characterizes the *normal distortion* - (*NormDistort*) and is used to measure the different behavior of two animation sequences during rendering. We compute the average difference between the original and the approximated face normals by the norm of their

cross product that equals to the sine of the angle between the two normal vectors. Therefore for a model with F faces and P frames, where $NV^{i,j}$ is the normal vector of face j at frame i , Equation 6 computes the normal distortion measure.

$$NormDistort = \sin^{-1} \left(\frac{1}{FP} \sum_{i=1}^P \sum_{j=1}^F \|NV_{orig}^{i,j} \times NV_{Approx}^{i,j}\| \right) \quad (6)$$

3.4 Building and Tuning a Neural Network for Weight Prediction

Network structure Our method adopts a supervised learning approach to leverage the power of neural networks on multiple class classification. Consequently, we utilize a neural network instead of using clustering techniques (unsupervised learning) to obtain better initial weights and bones for skinning. We have experimented with a variety of neural network models that can efficiently be trained to detect vertex motion patterns and mesh geometry characteristics and use similarities among them for clustering vertices into bones and determining weights implicitly through the influence of bones on the mesh surface. We have chosen networks that perform well in sequence learning.

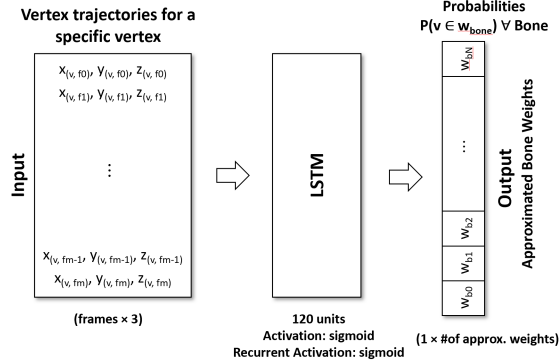


Fig. 4: LSTM-Network.

The first network that we propose as the first step and mean of animation compression is a Recurrent Neural Network (RNN). These networks are created by applying the same set of weights recursively over a differentiable graph-like structure by traversing it in topological order. This makes them suitable for classifying and predicting sequence data.

The type of RNN network used is a Long Short-term Memory network firstly introduced by [9] (LSTM), which consists of units made up of a cell remembering time inconstant data values, a corresponding forget cell, an input and an output

gate being responsible of controlling the flow of data in and out of the remembering component of it Figure 4. The actual difference of an LSTM compared to RNN is that LSTM has the capability of memorizing long-term dependencies regarding time data without resulting in emerging gradient vanishing problems (gradient loss exponentially decay). Not only does this capability make LSTM networks suitable for animation frame learning, but it is also a powerful way of predicting highly accurate weights.

Thus, utilization of many network units for LSTM construction (120 units used) produces a network that is able of predicting weights even for models with a large number of bones. Regarding the activation functions we used (i) an alternative for the activation function (cell and hidden state) by using *sigmoid* instead of *tanh* and (ii) the default for the recurrent activation function (for input, forget and output gate) which is *sigmoid*. The main reason of using the *sigmoid* function instead of the hyperbolic tangent is that our training procedure involves the network deciding per vertex whether it belongs or not to the influence range of a bone. This results in higher efficacy and additionally makes our model learn more effectively.

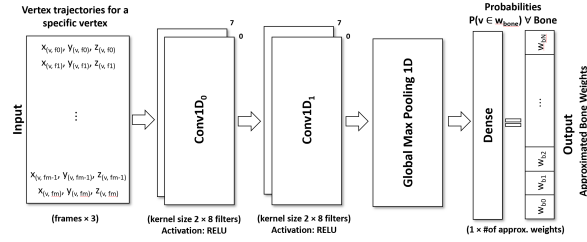


Fig. 5: CNN-Network.

The second network that we have used successfully is a feed-forward network called Convolutional Neural Network (CNN) [15] that uses convolutional operations to capture patterns in order to determine classes mainly in image classification problems. CNNs are additionally able to be used in classification of sequence data with quite impressive results. On top of the two convolutional layers utilized, we have also introduced a global max-pooling layer (down-sampling layer) and a simple dense layer so that we have the desirable number of weights for each proxy bone, as it is illustrated in Figure 5. In the two convolutional layers (Conv1D) used we utilize 8 filters of kernel size 2. The number of filters and kernel size have been determined experimentally.

The last network that we have considered for completeness is a hybrid neural network (Figure 7) that is a combination of the two aforementioned networks with some modifications. Unfortunately, the hybrid network does not perform equally well as its counterparts but it still derives comparable results.

Figures 8 and 9 illustrate the loss and accuracy values that our network models achieved with several batch sizes. For loss we utilized the binary cross-entropy function given by Equation 7, since we have a multi-label problem (a vertex may belong to multiple bones).

$$L(y, y_{pred}) = -\frac{1}{N} \sum_{i=0}^N ((1 - y) \cdot \log(1 - y_{pred}) + y \cdot \log(y_{pred})) \quad (7)$$

Where y are the real values (1: belongs to a bone or 0: does not) and y_{pred} are the predicted values. Binary cross-entropy measures how far in average a prediction is from the real value for every class. To this end, we also used binary accuracy which calculates the percentage of matched prediction-label pairs the 0/1 threshold value set to 0.5. What we have inferred by these plots is that for CNN there is no reason to increase the batch-size higher than 4096 owing to the fact that accuracy and loss values tend to be almost identical after increasing batch-size from 2048 samples to 4096. Likewise, for the LSTM case (see Figure 9) we observe that batch-size 2048 is the best option.

From Figures 8 and 9 we infer that we should use at least 20 epochs for training. After that the improvement of loss and accuracy is negligible but as we observed occasional overfitting is alleviated by increasing further the number of epochs.

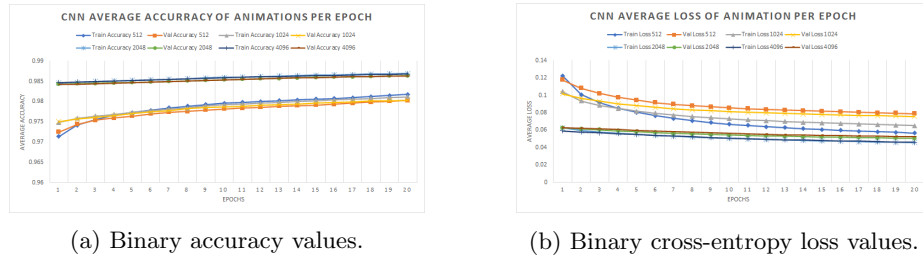
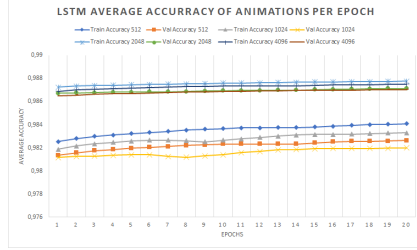


Fig. 8: Average per epoch Accuracy & Loss for CNN.

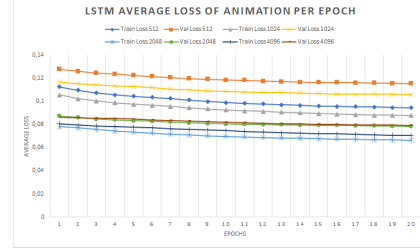
Furthermore, for the LSTM case, from Figure 11 we conclude that the network exhibits similar behavior to the CNN Figure 10. Therefore we select a batch-size of 4096 samples for training our LSTM model as well.

4 Experimental Evaluation of Deep Skinning

One of the main contributions of our work is that it expresses a combinatorial optimization problem with constraints as a classification problem and then proposes a method to solve it using deep learning techniques. For that reason, we have conducted a thorough experimental study to substantiate the effectiveness of our method based on the resulting error.

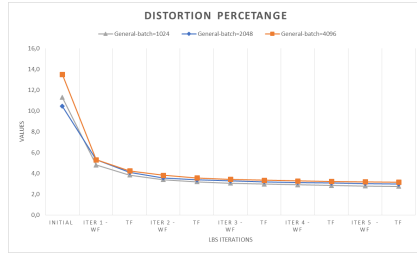


(a) Binary Accuracy values.

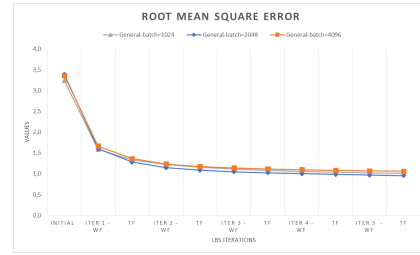


(b) Binary cross-entropy loss values.

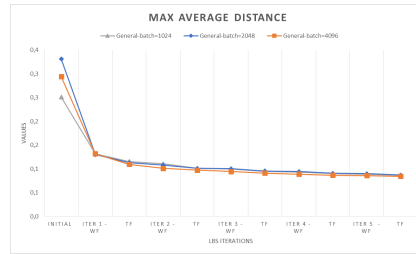
Fig. 9: Average per epoch Accuracy & Loss for LSTM.



(a) Distortion Percentage metric.



(b) Root Mean Square Error metric.



(c) Max Average Distance metric.

Fig. 10: Error Metrics for batch-size tuning in CNN.

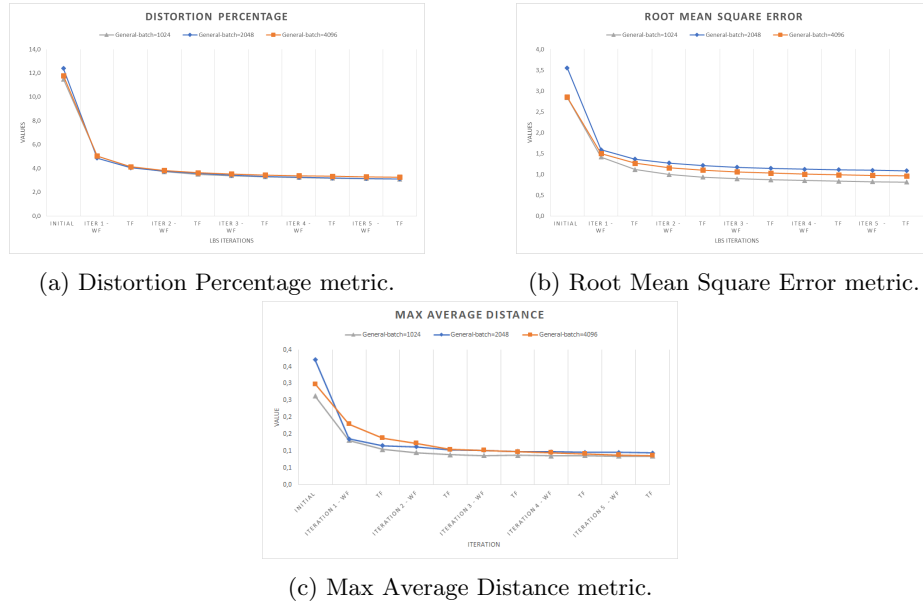


Fig. 11: Error Metrics for batch-size tuning in LSTM.

The entire method was developed¹ using Python and Tensorflow under the Blender 2.79b scripting API. The training part runs on a system with an NVIDIA GeForce RTX 2080Ti GPU with 11GB GDDR6 RAM. We trained our network models with Adam Optimizer [16], *learningRate* = 0.001 for 20 – 100 *epochs* with *batchSize* = 4096 over a training data-set that incorporates 60 animated character models of different size in terms of number of vertices, animations and frames per animation. We have inferred that 20 *epochs* are usually enough to have our method converging in terms of the error metrics and most importantly towards an acceptable visual outcome. However to obtain better RMS and distortion errors without over-fitting 100 *epochs* is a safe choice independently of the training set size. Furthermore, with this choice of batch-size we overcome the over-fitting problem that was apparent by observing the Max Average Distance metric and was manifested by locally distorted meshes.

The rest of our algorithm (prediction and optimization) was developed and ran on a commodity computer equipped with an Intel Core i7-4930K 3.4GHz processor with 48Gb under Windows 10 64-bit operating System. In addition, the FESAM algorithm was developed and ran on the same system.

4.1 Quantitative Results

In this section we present quantitative results for the Temporal Deep Skinning algorithm. We have conducted several experiments with multiple neural network

¹ source code: <https://github.com/AnastasiaMoutafidou/DeepSkinning>

structures to derive the top three choices of classification networks that fit best our training data with generalization capability.

The fitting part of our method optimizes weights and transformations (Weight Fitting-WF & Transformation Fitting-TF in figures 10, 11, 12, 13, 18) alternatively using a linear optimizer for five iterations. After five iterations we have observed that there is practically no improvement of the error for any of the methods. For each iteration the error metrics are computed and registered for every fitting case separately. Additionally, the initial error values in the plots below are the actual errors computed with the weights that each neural network produces.

Figures 12 and 13 provide a comparative evaluation of our method with the most competent skinning approach FESAM [14]. The original FESAM algorithm follows three steps of optimization with the first step being the process of optimizing the initial pose something that is not compatible with the traditional animation pipelines. For that reason this step is not included in our experiments and subsequently we use the FESAM-WT approach with two steps (weight and transformation optimization). More specifically, we evaluate the performance of our approach with distortion and RMS errors for the three top networks as compared to FESAM-WT. Based on Figure 12a and Figure 12b we conclude that LSTM 4096 performs overall better among the three prevalent networks on animal characters. We have performed the same experiments with the same test sets using decimated versions of our animation characters (50% and 20% decimation) and we have obtained the same results with differences only in the fourth decimal place.

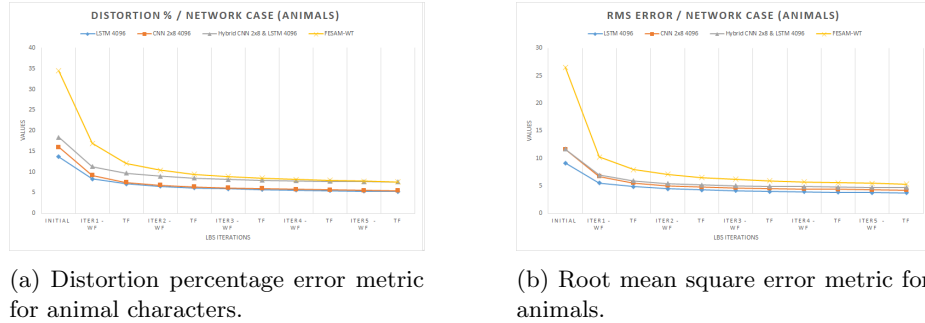
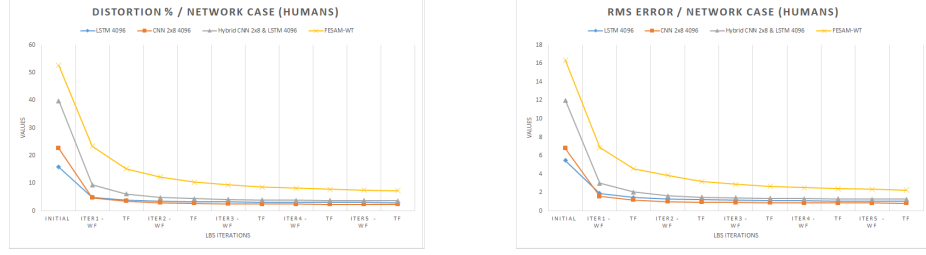


Fig. 12: Quantitative error results for animal characters.

The behavior of our method on human characters is illustrated in Figures 13a and 13b. We conclude that CNN 4096 is the most appropriate network structure in comparison with the other two.

Table 1 is a comparison of our method with other similar methods producing LBS schemes with pseudo bones when presented with several benchmark animation sequences from literature. More specifically, presents a comparison of



(a) Distortion percentage error metric on human characters.

(b) Root mean square error metric on human characters.

Fig. 13: Quantitative error results for human characters.

our method on four benchmark animation sequences, that were not produced by fully animated rigs, with all previous combinations of LBS, quaternion-based and SVD methods. N is the number of Vertices, F is the number of frames and the number in round brackets is the result of the method combined with SVD. Our method derives better results in terms of both error and compression rate as compared to methods I-IV. Method V is only cited for reference since it only obtains compression and is not compatible with any of the standard animation pipelines.

Approximation Error ERMS															
Input Data				Our Method		Method I		Method II		Method III		Method IV		Method V	
Dataset	N	F		Bones	ERMS	Bones	ERMS	Bones	ERMS	Bones	ERMS	Bones	ERMS	Bones	ERMS
Horse-gallop	8,431	48		26	0.15	30	2.3(0.3)	30	4.9(2.9)	30	1.3	30	2.4	-	2E-5
Elephant-gallop	42,321	48		18	0.35	25	2.6(0.5)	25	15(6.5)	25	1.4	25	2.3	-	6E-5
Camel-gallop	21,887	48		16	0.22	23	3.1(0.5)	23	4.7(2.2)	23	1.4	23	2.8	-	2E-4
Samba	9,971	175		17	0.60	30	8.6(3.6)	30	11.4(6)	30	1.5	30	4	-	0.2

Table 1: Comparative evaluation of our method versus Method I [11], Method II [13], Method III [14], Method IV [26], Method V [1].

Finally, Figure 14 shows the speed up that we have achieved in the fitting time by using the conjugate gradient method which is more efficient in multi dimensional problems such as the ones that we are solving ($12BP$ variables for transformations, and $6N$ variables for weights).

4.2 Visual Quality Evaluation Results

In computer graphics qualitative results (visual and otherwise) is an important means of assessing a novel method. In this section we present three processes for assessing the visual quality of temporal deep skinning.

Quality measure evaluation We use the MaxAvgDist quality assessment measure that indicates how far in terms of visual quality the generated frames

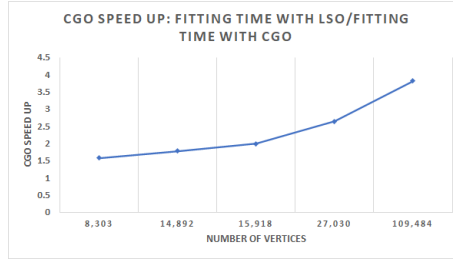
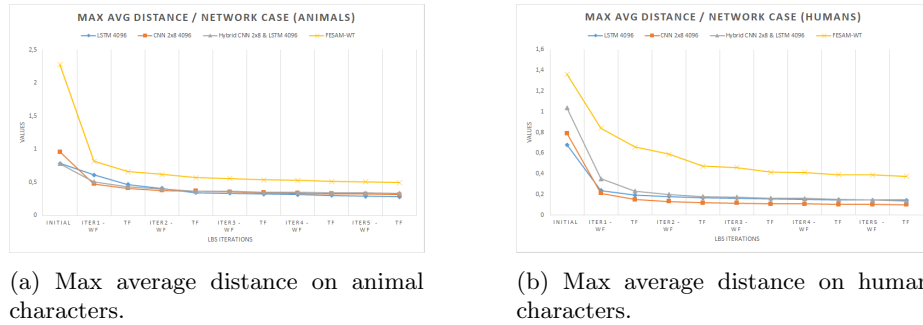


Fig. 14: Speed up of fitting time by using the conjugate gradient optimization method.

are from the original frame sequence (see Section 3.3). Low measure values correspond to high quality animation.



(a) Max average distance on animal characters.

(b) Max average distance on human characters.

Fig. 15: Qualitative error metric *MaxAvgDistr* results for humans and animals.

Figure 15 suggests that Temporal Deep Skinning yields results with better quality measure as compared to FESAM-WT. The results of Figures 15a and 15b confirm the quantitative results. Specifically, the LSTM network for animals and the CNN for humans are the most appropriate choices quantitatively and qualitatively.

Visualization-based Evaluation As an additional assessment criterion for our method we provide an illustration of the visual outcome. By using the term visual outcome, we refer to the approximated output frames as compared to the frames of original 3D model. After conducting several experiments we have observed that our approach seems to approximate better the original model. In every case there is a noticeable difference between temporal deep skinning and FESAM-WT. To this end, a demonstration video is also provided as supplementary material.

Figure 16 illustrates the differences of the two approximation methods as compared to the original model animation. Several frames have been selected with noticeable structural flaws.

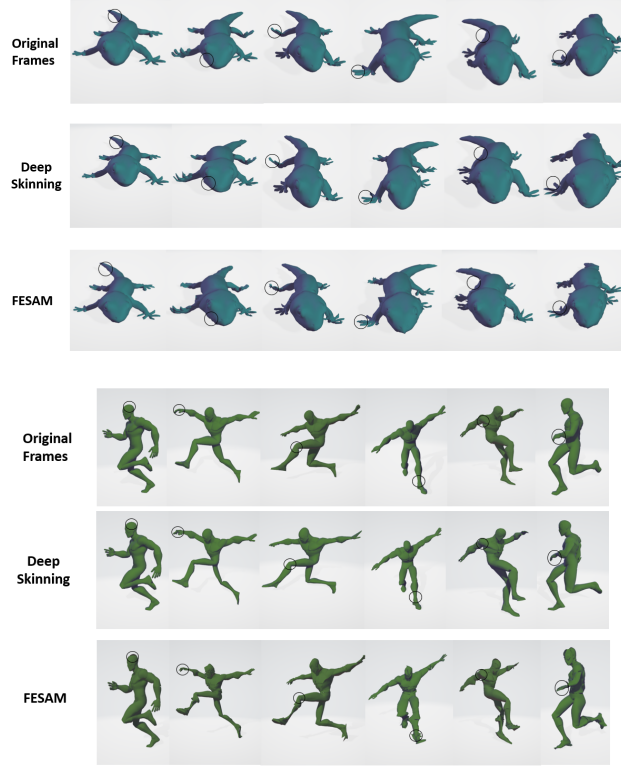


Fig.16: Visual comparison of Deep Skinning, FESAM-WT and the original frames for two models. Six frames have been selected in which structural flaws are marked by small circles.

Error visualization techniques can provide an insight for the parts where errors occur. We use the turbo colormap[25] to represent the error per vertex which is color blind friendly. This error is the distance in a particular frame of the approximated vertex from the original one. Figure 17 illustrates the per vertex error in a particular frame for deep skinning and FESAM_WT.

Lighting Quality Evaluation Finally, we offer the results of evaluating the average distortion of normal vectors. The normal distortion measure (see section 3.3) shows how close the normal vectors of the approximated sequence are to

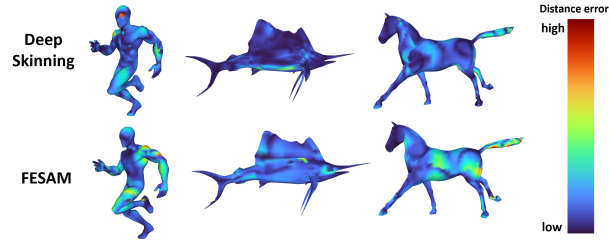
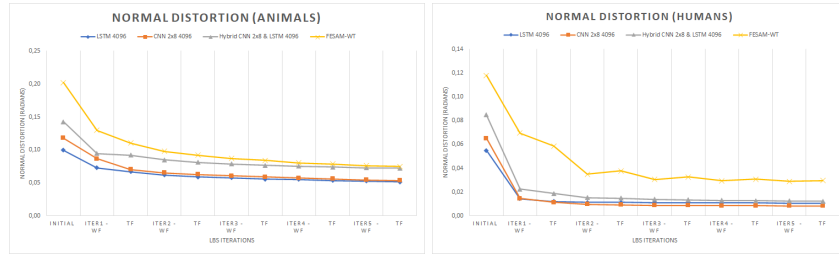


Fig. 17: Distance error comparison in a particular frame between Deep Skinning and FESAM-WT.

the normal vectors of the original animation sequence. This determines how the approximated character will behave in an lighting model as compared to the original animated character. The results of Figure 18 exhibit an average error of 0.01 radians for human characters and an average error of 0.05 radians for animal characters.



(a) Normal distortion on animal characters. (b) Normal distortion on human characters.

Fig. 18: Qualitative normal distortion metric results for humans and animals.

4.3 Discussion and Applications

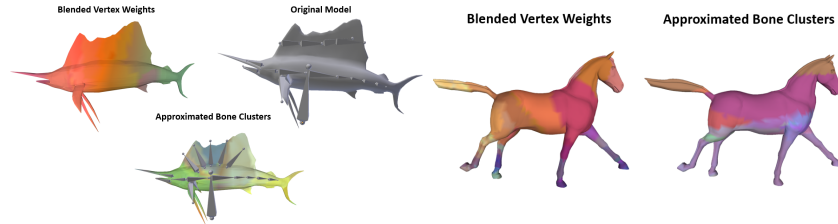
We have presented a method called *Temporal Deep Skinning* that feeds an animation sequence with no underlying skeleton or rigging/skinning information to a pre-trained neural network to generate an approximated compressed skinning model with pseudo-bones.

Moreover, we have developed a post processing tool that using the compressed skinning model with pseudo bones and per frame transformations obtained by temporal deep skinning produces the corresponding hierarchical skeleton, skinning data and transformations. More specifically, using the mesh clustering derived by our method, the pseudo bones and transformations we produce a fully

Input Data			Approximation Error <i>ERMS</i>														
Dataset	N	F	Our Method			Method A			Method B			Method C			Method D		
			Bones	ERMS	CRP	Bones	ERMS	CRP	Bones	ERMS	CRP	Bones	ERMS	CRP	Bones	ERMS	CRP
<i>Horse-gallop</i>	8,431	48	26	0.15	92.5	27	0.19	92.4	27	0.44	92.4	27	1.10	92.4	27	0.88	92.4
<i>Samba</i>	9,971	175	17	0.60	97.6	22	0.63	97.4	22	1.29	97.4	22	1.57	97.4	22	1.79	97.4

Table 2: Comparison between Temporal Deep Skinning and four methods. Specifically Method A [21], Method B [27], Method C [6], Method D [8].

animated character model. This is accomplished by the following steps: (i) perform weight regularization and derive disjoint vertex clusters that are influenced by each bone, (ii) based on the neighboring clusters export the joints of the entire model (the structure of the skeleton) [20] and (iii) finally perform joint location adjustment by geometric constraints and a simple recalculation of orientation and rotation for each of the joints that yields their final position [3]. Figure 19a illustrates the original and approximated representation of a 3D model. This animated model consisting of 14,007 vertices and 4,669 faces was approximated by the deep skinning algorithm with 24 bone clusters and up to six weights per vertex.



(a) Original and approximate representation for a Sailfish. (b) Approximate bone clusters and vertex weights for an animation sequence.

Fig. 19: Original and approximate representations.

Figure 19b presents the computed bones and weights for an animation sequence. This animation sequence consists of 48 different frames from the horse-gallop sequence. After the Deep Skinning algorithm we were able to extract 19 bone clusters and up to six weights per vertex. Subsequently, we have produced a fully animated character.

Table 2 provides a comparison of our method with four methods that produce actual skeletal rigs. In this case of 2 we cite the results from the papers since such methods are difficult to reproduce and this goes beyond the scope of this paper. For two models (*horse gallop* and *samba*) we have measured the *ERMS* error and the compression rate percentage (CRP). Note that the results of [21] were converted to our *ERMS* metric by multiplying by $\frac{D}{\sqrt{3}}$, where D is the diagonal of the bounding box of the rest pose.

For the horse gallop model our method approximates the sequence by using 26 bones and achieves a smaller *ERMS* error as compared to all previous methods. For the samba model our method uses 17 bones and outperforms all previous methods.

5 Conclusions

We have introduced a novel approach that derives pseudo-bones and weights for an animated sequence using deep learning on a training set of fully rigged animated characters. We have experimented with a variety of neural network models that can efficiently be trained to detect vertex motion patterns and mesh geometry characteristics and exploit similarities among them for clustering vertices into bones and determining weights implicitly through the influence of bones on the mesh surface. Our method does not require setting or tuning any parameters regarding the mesh structure or the kinematics of the animation.

Based on a comparative evaluation, we conclude that the approximation error of our method is always smaller than the error of previous approaches that are compatible with existing animation pipelines.

6 Conclusions

References

1. Alexa, M., Müller, W.: Representing animations by principal components. *Comput. Graph. Forum* **19** (2000)
2. Au, O.K.C., Tai, C.L., Chu, H.K., Cohen-Or, D., Lee, T.Y.: Skeleton extraction by mesh contraction. *ACM Trans. Graph.* **27**(3), 44:1–44:10 (2008)
3. Avril, Q., Ghafourzadeh, D., Ramachandran, S., Fallahdoust, S., Ribet, S., Dionne, O., de Lasa, M., Paquette, E.: Animation setup transfer for 3d characters. In: *Proceedings of the 37th Annual Conference of the European Association for Computer Graphics, EG '16*, p. 115–126. Eurographics Association, Goslar, DEU (2016)
4. Bailey, S.W., Otte, D., Dilorenzo, P., O'Brien, J.F.: Fast and deep deformation approximations. *ACM Trans. Graph.* **37**(4) (2018)
5. De Aguiar, E., Theobalt, C., Thrun, S., Seidel, H.P.: Automatic conversion of mesh animations into skeleton-based animations. *Computer Graphics Forum* **27**(2), 389–397 (2008)
6. De Aguiar, E., Theobalt, C., Thrun, S., Seidel, H.P.: Automatic Conversion of Mesh Animations into Skeleton-based Animations. *Computer Graphics Forum* (2008)
7. Feng, A., Casas, D., Shapiro, A.: Avatar reshaping and automatic rigging using a deformable model. In: *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games, MIG '15*, p. 57–64. Association for Computing Machinery, New York, NY, USA (2015). DOI 10.1145/2822013.2822017
8. Hasler, N., Thormählen, T., Rosenhahn, B., Seidel, H.P.: Learning skeletons for shape and pose. In: *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '10*, p. 23–30. Association for Computing Machinery, New York, NY, USA (2010). DOI 10.1145/1730804.1730809
9. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9**(8), 1735–1780 (1997)
10. Jacobson, A., Deng, Z., Kavan, L., Lewis, J.P.: Skinning: Real-time shape deformation. In: *ACM SIGGRAPH 2014 Courses, SIGGRAPH '14*. Association for Computing Machinery, New York, NY, USA (2014). DOI 10.1145/2614028.2615427
11. James, D.L., Twigg, C.D.: Skinning mesh animations. In: *ACM SIGGRAPH 2005 Papers, SIGGRAPH '05*, p. 399–407. Association for Computing Machinery, New York, NY, USA (2005)
12. Kavan, L., Collins, S., Žára, J., O'Sullivan, C.: Skinning with dual quaternions. In: *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games, I3D '07*, p. 39–46. Association for Computing Machinery, New York, NY, USA (2007). DOI 10.1145/1230100.1230107
13. Kavan, L., McDonnell, R., Dobbyn, S., Žára, J., O'Sullivan, C.: Skinning arbitrary deformations. In: *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games, I3D '07*, p. 53–60. Association for Computing Machinery, New York, NY, USA (2007)
14. Kavan, L., Sloan, P.P., O Sullivan, C.: Fast and Efficient Skinning of Animated Meshes. *Computer Graphics Forum* (2010)
15. Khan, A., Sohail, A., Zahoor, U., Qureshi, A.S.: A survey of the recent architectures of deep convolutional neural networks (2019)
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2014)
17. Kraevoy, V., Sheffer, A.: Cross-parameterization and compatible remeshing of 3d models. In: *ACM SIGGRAPH 2004 Papers, SIGGRAPH '04*, pp. 861–869. ACM, New York, NY, USA (2004)

18. Kry, P.G., James, D.L., Pai, D.K.: Eigenskin: Real time large deformation character skinning in hardware. In: Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '02, p. 153–159. Association for Computing Machinery, New York, NY, USA (2002)
19. Le, B.H., Deng, Z.: Smooth skinning decomposition with rigid bones. *ACM Trans. Graph.* **31**(6), 199:1–199:10 (2012)
20. Le, B.H., Deng, Z.: Smooth skinning decomposition with rigid bones. *ACM Trans. Graph.* **31**(6) (2012). DOI 10.1145/2366145.2366218
21. Le, B.H., Deng, Z.: Robust and accurate skeletal rigging from mesh sequences. *ACM Trans. Graph.* **33**(4) (2014). DOI 10.1145/2601097.2601161
22. Liu, L., Zheng, Y., Tang, D., Yuan, Y., Fan, C., Zhou, K.: Neuroskinning: Automatic skin binding for production characters with deep graph networks. *ACM Trans. Graph.* **38**(4) (2019)
23. Luo, R., Shao, T., Wang, H., Xu, W., Chen, X., Zhou, K., Yang, Y.: Nnwarp: Neural network-based nonlinear deformation. *IEEE Transactions on Visualization and Computer Graphics* **26**(4), 1745–1759 (2020)
24. Magnenat-Thalmann, N., Laperrière, R., Thalmann, D.: Joint-dependent local deformations for hand animation and object grasping. In: Proceedings on Graphics Interface '88, p. 26–33. Canadian Information Processing Society, CAN (1989)
25. Mikhailov, A.: Turbo, An Improved Rainbow Colormap for Visualization, Google AI Blog, August 20, 2019
26. Sattler, M., Sarlette, R., Klein, R.: Simple and efficient compression of animation sequences. In: Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '05, p. 209–217. Association for Computing Machinery, New York, NY, USA (2005)
27. Schaefer, S., Yuksel, C.: Example-based skeleton extraction. In: Proceedings of the Fifth Eurographics Symposium on Geometry Processing, SGP '07, p. 153–162. Eurographics Association, Goslar, DEU (2007)
28. Vasilakis, A.A., Fudos, I., Antonopoulos, G.: Pps: Pose-to-pose skinning of animated meshes. In: Proceedings of the 33rd Computer Graphics International, CGI '16, pp. 53–56. ACM, New York, NY, USA (2016)
29. Xu, Z., Zhou, Y., Kalogerakis, E., Landreth, C., Singh, K.: Rignet: Neural rigging for articulated characters. *ACM Trans. on Graphics* **39** (2020)
30. Zell, E., Botsch, M.: Elastiface: Matching and blending textured faces. In: Proceedings of the Symposium on Non-Photorealistic Animation and Rendering, NPAR '13, pp. 15–24. ACM, New York, NY, USA (2013)