



# Clustering and Recognition of Spatiotemporal Features Through Interpretable Embedding of Sequence to Sequence Recurrent Neural Networks

Kun Su<sup>1</sup> and Eli Shlizerman<sup>1,2\*</sup>

<sup>1</sup> Department of Electrical & Computer Engineering, University of Washington, Seattle, WA, United States, <sup>2</sup> Department of Applied Mathematics, University of Washington, Seattle, WA, United States

## OPEN ACCESS

### Edited by:

Balaraman Ravindran,  
Indian Institute of Technology Madras,  
India

### Reviewed by:

Avinash Sharma,  
International Institute of Information  
Technology, Hyderabad, India  
Sarath Chandar,  
École Polytechnique de Montréal,  
Canada

### \*Correspondence:

Eli Shlizerman  
shlizee@uw.edu

### Specialty section:

This article was submitted to  
Machine Learning and Artificial  
Intelligence,  
a section of the journal  
Frontiers in Artificial Intelligence

**Received:** 17 January 2020

**Accepted:** 29 July 2020

**Published:** 30 September 2020

### Citation:

Su K and Shlizerman E (2020)  
Clustering and Recognition of  
Spatiotemporal Features Through  
Interpretable Embedding of Sequence  
to Sequence Recurrent Neural  
Networks. *Front. Artif. Intell.* 3:70.  
doi: 10.3389/frai.2020.00070

Encoder-decoder recurrent neural network models (RNN Seq2Seq) have achieved success in ubiquitous areas of computation and applications. They were shown to be effective in modeling data with both temporal and spatial dependencies for translation or prediction tasks. In this study, we propose an embedding approach to visualize and interpret the representation of data by these models. Furthermore, we show that the embedding is an effective method for unsupervised learning and can be utilized to estimate the optimality of model training. In particular, we demonstrate that embedding space projections of the decoder states of RNN Seq2Seq model trained on sequences prediction are organized in clusters capturing similarities and differences in the dynamics of these sequences. Such performance corresponds to an unsupervised clustering of any spatio-temporal features and can be employed for time-dependent problems such as temporal segmentation, clustering of dynamic activity, self-supervised classification, action recognition, failure prediction, etc. We test and demonstrate the application of the embedding methodology to time-sequences of 3D human body poses. We show that the methodology provides a high-quality unsupervised categorization of movements. The source code with examples is available in a Github repository<sup>1</sup>.

**Keywords:** spatiotemporal feature, interpretable AI, sequence to sequence (Seq2Seq), clustering, action recognition

## 1. INTRODUCTION

Recurrent Sequence to Sequence (Seq2Seq) network models use internal states to process sequences of inputs (Hochreiter and Schmidhuber, 1997; Cho et al., 2014; Sutskever et al., 2014; Luong et al., 2015). The speciality of Seq2Seq is that these models consist of encoder and decoder components. The encoder typically processes input sequences and constructs a latent representation of the sequences. In addition, the encoder passes the last internal state to the decoder as an “initialization” of the decoder. With this information the decoder transforms or generates novel sequences with a similar distribution. Such an architecture and its variants, e.g., attention-based

<sup>1</sup><https://shlizee.github.io/interpretseq2seq/>

Seq2Seq (Cho et al., 2014), showed compelling performance in applications of machine translation (Sutskever et al., 2014), speech recognition (Graves et al., 2013), and human motion prediction (Gui et al., 2018).

While Seq2Seq and its variants achieve strong performance on various applications, a consistent interpretation of how the encoder-decoder structure is capable to embed the data for general time-series data (i.e., multi-dimensional ordered sequences) and how such interpretation can be used to estimate the performance of the model is an active research topic. In this paper, we propose a dimension reduction approach to visualize and interpret the representation within the Seq2Seq model of general time-series data. The main contribution of this paper is to provide constructive insights on the properties which allow the encoder and the decoder components to operate optimally and to propose a low dimensional visualization method of the representation that the encoder and the decoder construct. With our embedding method, we find a remarkable property of Seq2Seq: the network being trained to predict the future evolution of a sequence self-organizes the hidden units representation into separate identities (clusters or classes). The clusters are embedded in the embedding space as attractors to which embedded encoder trajectories lead. We show on data of poses of the human body, how the organization of the attractors provides an unsupervised easy means for clustering sequential data into distinct identities.

Interpretability has been an important aspect of any artificial neural network (ANN) and is expected to provide generic methodologies to assess the capabilities of different models and evaluate the performance of models for given tasks. Associating interpretability with various types of ANN is often a challenging undertaking since the typical state-of-the-art network models are high dimensional and include many components and processing stages (layers). The problem becomes more challenging when dealing with RNN sequence models and in unsupervised tasks such as synthesis and translation, in which encoder-decoder networks were found to be prevalent. In these tasks, interpretability is aimed to capture the model synthesis procedures and to demonstrate how these are being improved during training and since learning is unsupervised, interpretation methodologies have the potential to provide a framework to assist with enhancing the optimization and learning.

Typically, existing interpretation methods applicable to ANNs are application-oriented. For Convolutional Neural Networks (CNN), models were explained within images (Zeiler and Fergus, 2014; Alain and Bengio, 2016; Kim et al., 2017). For RNN, interpretation and visualization tools focus on natural language processing applications (Karpathy et al., 2015; Collins et al., 2016; Foerster et al., 2017; Strobel et al., 2019). Our proposed work is based and inspired by these works and extends the methodology to generic sequences, in which textual sequences are a sub-class with special time dependence (semantics). Generic multidimensional time series are spatio-temporal sequences including non-trivial correlations in space and time. Beyond text data, there are works inspired by neuronal networks investigating the dynamics of RNN (Recanatani et al., 2018;

Farrell et al., 2019). In our work, we aim to provide an interpretation of encoder-decoder (Seq2Seq) network models for general spatiotemporal data.

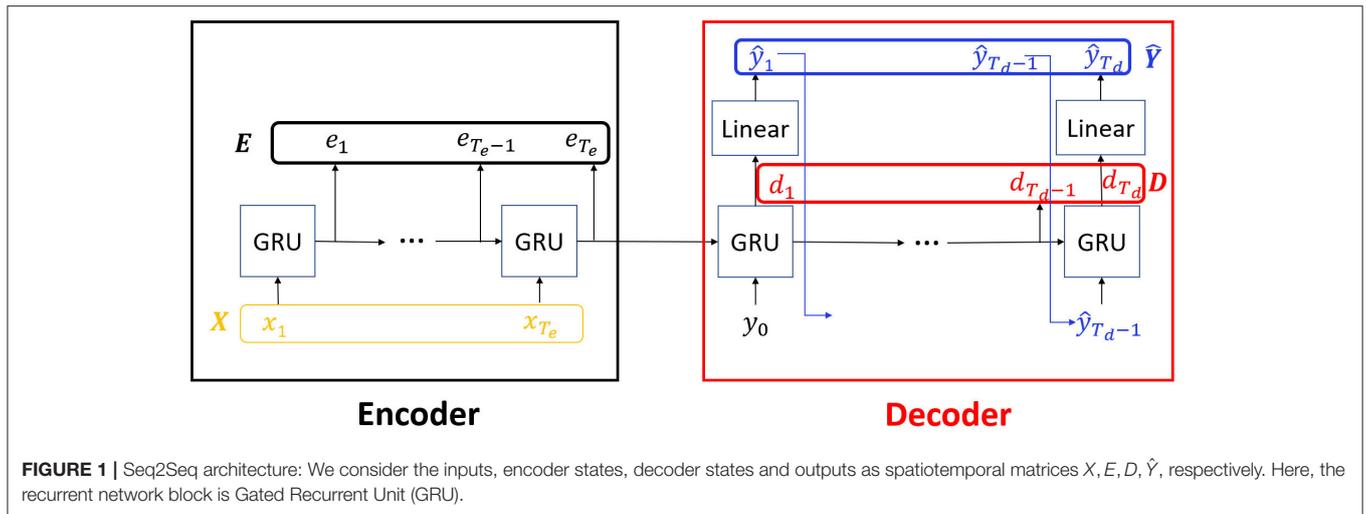
We test our methods on prediction tasks of synthetic data and of typical movements of human body joints. There are several RNN-based Seq2Seq models that achieve success on human motion prediction (Martinez et al., 2017) and outperform the previous non-Seq2Seq based RNN models such as ERD (Fragkiadaki et al., 2015) and S-RNN (Jain et al., 2016). Recently, Generative Adversarial Networks (Gui et al., 2018) have achieved better performance on this task, with the predictor network being RNN Seq2Seq.

We show that Seq2Seq optimization with gradient descent based methods forms a low dimensional embedding of internal states. The embedding can be mapped and visualized through Proper Orthogonal Decomposition (POD) of concatenated encoder and decoder internal states - the interpretable embedding. Within this embedding, the decoder evolution for each distinct sequence (decoder trajectory) is separable from other distinct sequences. Furthermore, each distinct decoder trajectory preserves both the spatial and the temporal properties of the sequence. The encoder trajectory initiated from various starting points connects them in the interpretable embedding space with the appropriate decoder trajectory. Monitoring the interpretable embedding space and projected trajectories in it during training shows the effect of training on data representation and assists to identify an optimal regime between under- and over- fitting. We construct synthetic data examples to demonstrate the construction of the interpretable embedding space. Next, we apply the construction of the interpretable embedding and analyze Seq2Seq performance on human joints movements datasets: Human 3.6 million (H3.6M) that contains 15 different types of real body movement sequences, such as walking, eating, etc. (Ionescu et al., 2014). To show generality and example application for the proposed approach, we apply it to the CMU Motion capture dataset and perform an unsupervised action recognition task.

## 1.1. Setup and Spatiotemporal States of the Seq2Seq Model

RNN Seq2Seq model is utilized for sequence prediction (synthesis of a new sequence based on input sequence) or translation (mapping the input sequence to a new representation) (Figure 1). For general time series prediction, given a sequence of spatiotemporal data as input, Seq2Seq predicts the future sequence. We stack the sequence of input to the encoder as a matrix and call it the **encoder input matrix**  $\mathbf{X} \in \mathbb{R}^{T_e \times M}$ , where each row  $x_t \in \mathbb{R}^{1 \times M}$  is a time step of the input sequence at time  $t$ ,  $T_e$  is the number of time steps and  $M$  is the number of dimensions in the input data. Similarly, we construct the target sequences as **target output matrix**  $\mathbf{Y} \in \mathbb{R}^{T_d \times M}$ , where  $T_d$  is the number of output sequence steps to be predicted.

In addition, forward propagation of the input in RNN Seq2Seq computes the internal states of the encoder at each time step. We concatenate and denote them as the **encoder states matrix**



$E \in \mathbb{R}^{T_e \times N}$ , where each row  $e_t \in \mathbb{R}^{1 \times N}$  represents the states of all  $N$  internal units at time  $t$ . Similarly, we also define the **decoder states matrix**  $D \in \mathbb{R}^{T_d \times N}$ . Typically, there is an additional fully connected linear map transforming the decoder states to the dimensions of the output space. We denote the **decoder output matrix**  $\hat{Y} \in \mathbb{R}^{T_d \times M}$ , where  $\hat{y}_t$  is the predicted  $\hat{y}_t$  output at time  $t$ . **Figure 1** demonstrates the structure of the components of RNN Seq2Seq matrices. In the figure, the encoder and the decoder are single layer GRU networks that can share or not share parameters with each other depending on applications. Our approach is applicable to general types of networks, such as LSTMs/GRUs, and variable number of layers. In our setup, the decoder uses the output of a previous time step as an input to the current time step in both training and testing, except for the initial time step in which it receives its input from the last step of the encoder. This setup can vary as well. In terms of time series prediction task, the cost function is typically defined as the MSE between **target outputs** and **decoder outputs**,  $J = \frac{1}{T_d} \sum_{t=1}^{T_d} (y_t - \hat{y}_t)^2$ , however, other norms or cost functions can be considered.

### 1.2. Proper Orthogonal Decomposition of Spatiotemporal Matrices

Since forward propagation within RNN Seq2Seq can be represented through spatiotemporal matrices, we propose to apply the POD method as a dimensionality reduction algorithm to construct a low dimensional interpretable embedding (Shlizerman et al., 2012). Specifically, we use the Singular Value Decomposition to decompose the matrices into orthogonal spatial modes (PCs) and time-dependent coefficients and singular values (scaling) associated with each mode. Particularly, given a matrix  $A \in \mathbb{R}^{T \times N}$ , we first perform a normalization where we subtract the mean of that column from each entry and obtain  $A_c$ . The normalization ensures that each column has the same mean (zero). We then apply SVD such that  $A_c = U \Sigma V^T$ , where  $U \in \mathbb{R}^{T \times T}$  is an orthogonal matrix of time-dependent coefficients,  $\Sigma \in \mathbb{R}^{T \times N}$  is the matrix of singular values and  $V \in \mathbb{R}^{N \times N}$  is the matrix of spatially dependent components. To determine the number of PCs, we compute the singular value energy (SVE),  $SVE = \sum_{i=1}^n \sigma_i^2$  where  $\sigma_i$  is

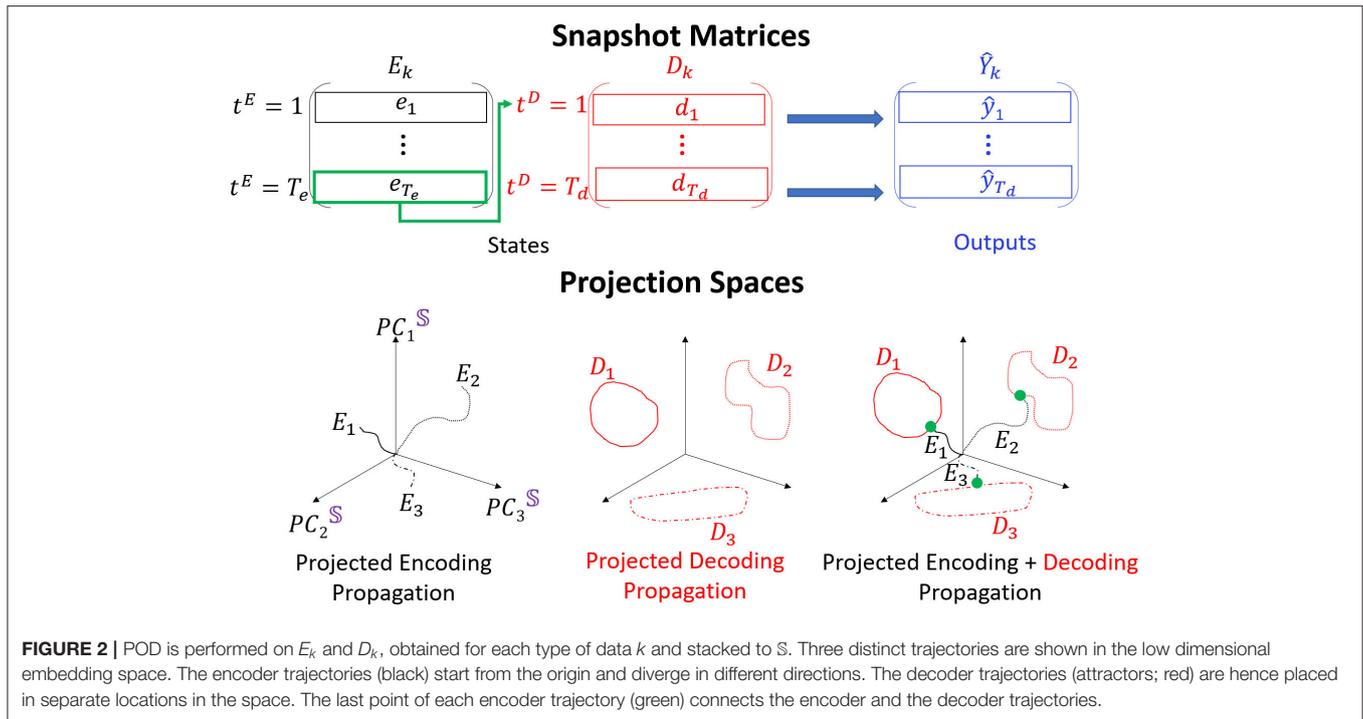
singular value corresponding to PC mode  $i$ . We compute the number of modes such that 90% or 99% of the energy is retained ( $SVE_k = \sum_{i=1}^k \sigma_i^2, \frac{SVE_k}{SVE} \leq p$  where  $k$  is the number of modes and  $p$  is the percentage). With the number of dominant spatial features, we can truncate  $A_c$  by projecting it onto the PC modes to get a low dimensional matrix  $A_{(PC=n)} \in \mathbb{R}^{T \times n}$ , where (PC= $n$ ) denotes  $n$  principal components are retained. If we choose  $n = 2$  or  $3$ , we can visualize the representation in 2D or 3D. The axes are the orthogonal PC modes.

### 1.3. Clustering

While visualization of projected dynamics could be informative (Maaten and Hinton, 2008), 2D or 3D visualized dynamics do not reveal the intricacies in the representation of different datasets. In particular, here we would like to evaluate the separability of projections of distinct trajectories in the interpretable embedding space. We thereby propose to augment the embedding with clustering approaches such as **K-means++**, an extended version of the standard K-means algorithm (Arthur and Vassilvitskii, 2007) or **agglomerative** clustering, a bottom up approach of hierarchical clustering. Since the number of trajectories and time steps are known, we can use the **Adjusted Rand Index (ARI)** to evaluate the clustering performance.

## 2. INTERPRETABLE EMBEDDING FOR Seq2Seq NETWORKS

The generic goal of RNN Seq2Seq model is to continue (predict) the evolution of each given input sequence chosen from a (test) dataset which includes  $K$  different types of multidimensional time series. Such a goal is challenging as it requires the network to generate a sequence which superimposes the typical dynamics of that particular type and the individual dynamics of the given tested input sequence. We show that the methods described above can be applied to construct an interpretable embedding for the RNN Seq2Seq model. To illustrate the embedding properties we use an example, depicted in **Figure 2**, containing three types of time series sequences. We show further examples



with more sequences in the section of “Human Body Joints Movements Data.”

To construct the embedding space basis we concatenate the matrices  $\mathbf{E}$  and  $\mathbf{D}$  for each single forward propagation into a **states matrix**

$$\mathbf{S} = \begin{bmatrix} \mathbf{E} \\ \mathbf{D} \end{bmatrix} \in \mathbb{R}^{(T_e+T_d) \times N}, \quad \mathbb{S} = \begin{bmatrix} s_1 \\ \vdots \\ s_k \end{bmatrix}.$$

Concatenation of forward propagation evolution for all considered time series in the dataset will result with a **global states matrix** denoted as  $\mathbb{S}$ . POD application on  $\mathbb{S}$  provides the PC modes which are the axes of the interpretable embedding space. Dimension reduction of the embedding space is performed by considering only  $n$  PC modes (denoted as PC $n$ ) which singular values are included in the representation of particular total SVE (e.g., 90% or 99%) and truncating the rest of the modes.

To inspect the propagation of single time series through the network, we can project the matrices  $\mathbf{E}$ ,  $\mathbf{D}$ , or  $\mathbf{S}$  onto the low dimensional space spanned by PC $n$  modes. As we show below, we find that the dimension of the embedding space can be very low even for high dimensional data. We depict the structure of such projections onto PC3 embedding space in **Figure 2**, bottom. Encoder trajectories (black, left) start from initial points projected to the embedding space (we choose initial states as zeros and therefore the trajectories are always initiated at the origin) and evolve in different directions in the space. Decoder trajectories (red, middle) appear as attractors in the space, and clustering approaches are used to determine (e.g., K-means) how separable they are in the space. The encoder and the decoder connect via a single time step which corresponds to a point (green, right). Composition of the three types of projections corresponds to an interpretation of the propagation in RNN

Seq2Seq. In particular, we show that the encoder trajectory takes the sequence from an initial point and evolves it to the corresponding starting point on the decoder trajectory (we call it attractor or cluster). The decoder continues the evolution from there. As we show below, it appears that the gradient descent training succeeds to organize the decoder attractors in the embedding space such that they are easily clustered. Such an arrangement explains the uniqueness of RNN Seq2Seq training in which the cost function minimizes the error between the decoder output and the actual output and there is no minimization on the encoder. Therefore, the decoder is trained to predict different features for different types of inputs (training to optimize clustering of types of data and capture unique features) with the encoder trajectory (and not only the last time step of it) being a sequential constraint that connects the cluster to the initial state. In practice, the length of the encoding sequence should be set such that the encoder can capture enough information but not be sensitive to the vanishing gradient problem. We discuss the encoding sequence length parameter choices and show clustering results in **Supplementary Materials**.

It is possible to monitor the construction and the changes in the embedding space and projected trajectories within it during training. Specifically, at each training iteration  $i$ , for each type of data, we construct the matrices  $E^i, D^i, S^i$  and perform the aforementioned procedures. The pseudo code is shown in the Algorithm 1.

### 3. INTERPRETABLE EMBEDDING SPACE APPLIED TO SYNTHETIC DATA

We first create two types of synthetic 2D trajectories which follow a circumference of (i) a unit circle  $X_{circle}$  (ii) an ellipse

---

**Algorithm 1:** Interpretable Embedding and Clustering on Seq2Seq Networks
 

---

**Input:** RNN Seq2Seq model  $F$   
**Data:** Time series data  $X_1, \dots, X_K$

- 1 Initialize  $\mathbb{S}$  and  $\mathbb{D}$  be empty lists [ ]
- 2 **for**  $k$  in  $(1, \dots, K)$  **do**
- 3   Forward propagate  $X_k$  through  $F$  to obtain  $\mathbf{E}_k, \mathbf{D}_k$
- 4    $\mathbf{S}_k = \begin{bmatrix} \mathbf{E}_k \\ \mathbf{D}_k \end{bmatrix}$
- 5   Append  $\mathbf{D}_k$  to  $\mathbb{D}$
- 6   Append  $\mathbf{S}_k$  to  $\mathbb{S}$
- 7   **if** *Visualize Single Sequence On 3D* **then**
- 8     Normalize  $\mathbf{S}_k$
- 9     Apply SVD to have  $\mathbf{S}_k = \mathbf{U}_k \Sigma_k \mathbf{V}_k^T$
- 10     $\mathbf{S}_k \leftarrow \mathbf{S}_k \mathbf{V}_k[:, :3]$
- 11    Plot  $\mathbf{S}_k$
- 12 **if** *Clustering* **then**
- 13   Normalize  $\mathbb{D}$
- 14   **if** *Clustering On Dimension  $n$*  **then**
- 15     Apply SVD to have  $\mathbb{D} = \mathbf{U}_D \Sigma_D \mathbf{V}_D^T$
- 16      $\mathbb{D} \leftarrow \mathbb{D} \mathbf{V}_D[:, :n]$
- 17   Apply Clustering Algorithm (e.g., K-means) on  $\mathbb{D}$ .
- 18 **if** *Visualize All Sequences on 3D* **then**
- 19   Normalize  $\mathbb{S}$
- 20   Apply SVD to have  $\mathbb{S} = \mathbf{U}_S \Sigma_S \mathbf{V}_S^T$
- 21    $\mathbb{S} \leftarrow \mathbb{S} \mathbf{V}_S[:, :3]$
- 22   Plot  $\mathbb{S}$

---

$X_{\text{ellipse}}$ . The encoder input is binned into  $T_e$  time steps such that  $X_{\text{circle}}, X_{\text{ellipse}} \in \mathbb{R}^{T_e \times 2}$  with  $T_e = 50$  such that it completes a full period. Given full circle or ellipse dynamics as the encoder input, the objective is to predict another full circle or ellipse by the decoder, i.e., the target sequence  $Y$  is expected to be the same as the input sequence ( $Y \equiv X$ ) and the **decoder output matrix**  $\hat{Y}$  has the same dimension as  $Y$  and  $X$ . The cost function  $J = \frac{1}{T_d} \sum_{t=1}^{T_d} (y_t - \hat{y}_t)^2$  is the MSE between the target and the prediction with  $T_d = T_e$ . For both the encoder and the decoder we use a single layer GRU with 16 neurons ( $E, D \in \mathbb{R}^{T_e \times 16}$ ) and the ADAM optimizer to train the model for 5,000 iterations where the cost function almost converges to zero such that the model can predict the trajectory with high accuracy.

We show in **Figure 3** the projections of forward propagation in RNN Seq2Seq trained on (i) circle (ii) ellipse (iii) both circle and ellipse. In each model, we obtain the PC3 embedding space from the matrix  $\mathbb{S}$  and examine the projections of the encoder states matrix  $\mathbf{E}$ , decoder states matrix  $\mathbf{D}$  onto the space and also show the decoder output matrix  $\hat{Y}$  projected onto  $x - y$  space. We observe that the projections of the encoder and the decoder states are deformed and not necessarily preserve the same form as the input and the output, however, linear transformation of the decoder deforms the trajectory to be correctly represented in the  $x - y$  space. Notably, all encoder

projections start near the origin since our initial states are zero by default.

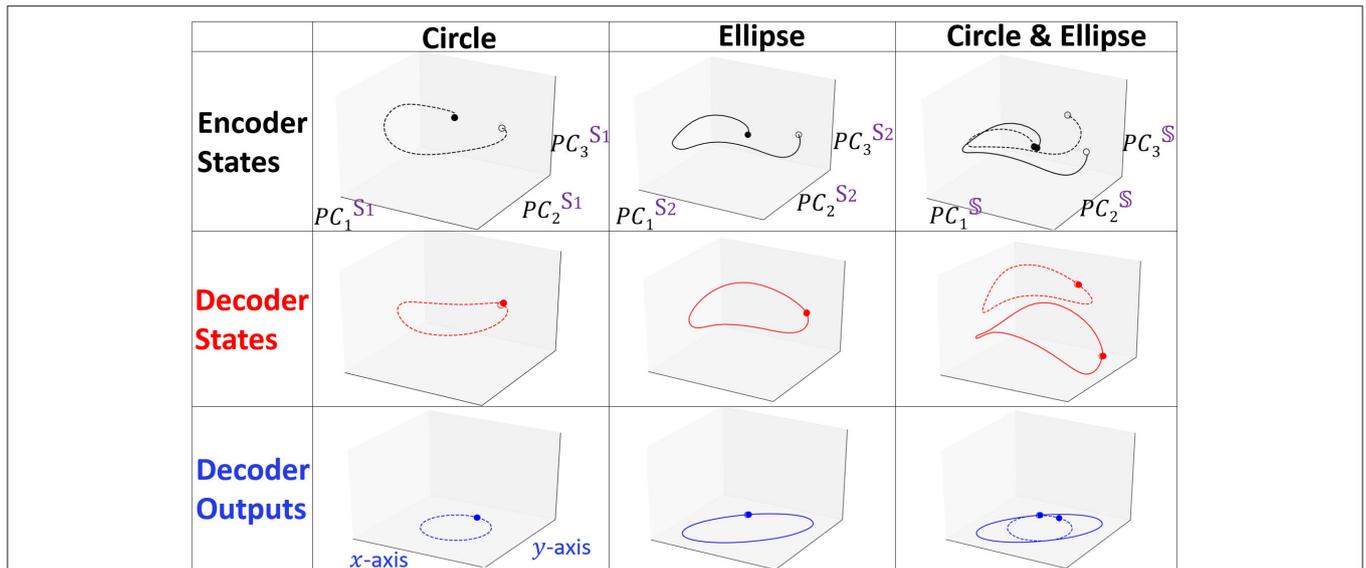
The last state of the encoder in NLP applications is typically considered to contain the information of all previous states, however, our results indicate that the full encoder sequence is important and the last encoder state is simply providing a starting point for the decoder to continue. Such an effect of the encoder is easily observed in considering continuous time series (as we show in the prediction of human movement data) and deviates from the interpretation of the encoder role in textual semantic sequences.

We focus on the model trained on both the circle and the ellipse to better understand how RNN Seq2Seq can predict both spatio-temporal series in an unsupervised way without any guidance. Specifically, we show projected trajectories in the PC embedding space (right column of **Figure 3**). Encoder projections indicate that the two trajectories corresponding to distinct types of output sequences, start from points near the origin, however, diverge as the sequence evolves, and end up at more distant points. The projected trajectories of the decoder states start from these distinct points and continue to perform prediction in completely separable shapes. Effectively we observe that the decoder states projections are clustered in the embedding space. Application of agglomerative clustering and cosine similarity indeed verify these observations.

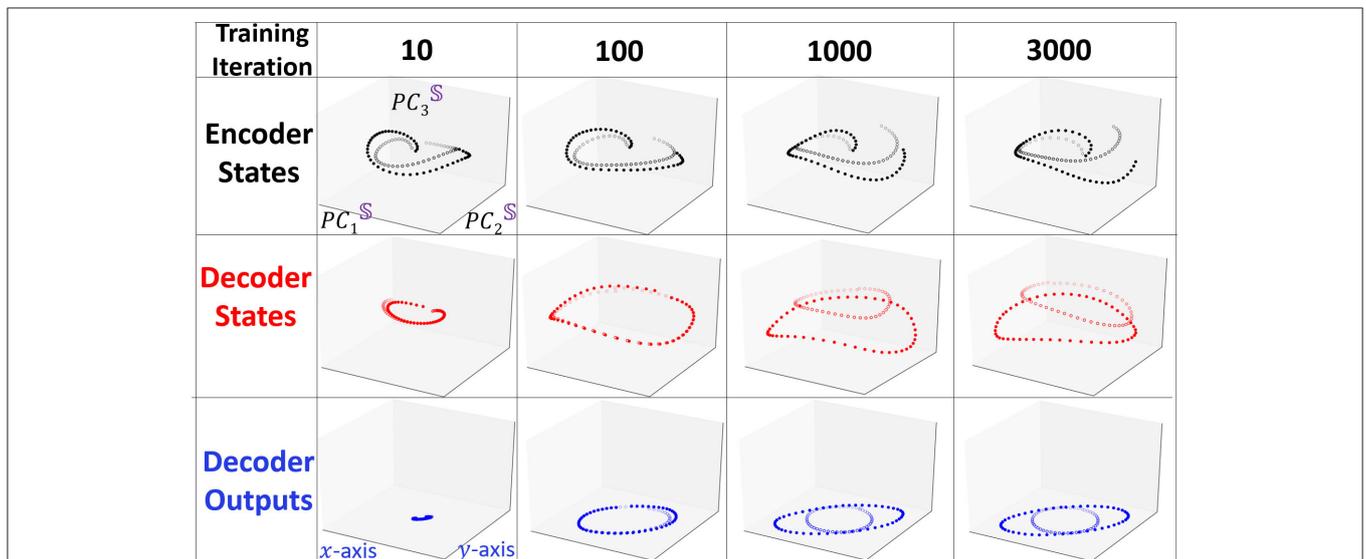
To visualize how Seq2Seq learns to differentiate the two shapes with training, we keep track of the evolution of the representation of the states and depict in **Figure 4** the three projections as in **Figure 3** at iterations  $i = 10, 100, 1,000, 3,000$ .

At the beginning of training ( $i = 10$ ), all projections of the two shapes are similar. Projections of the decoder states appear to be distinct for a few several initial points but when the sequence evolves forward, the trajectories appear to converge to the same point. When the model undergoes additional training, after  $i = 100$ , Seq2Seq appears to have learned a single pattern (ellipse like), however, is unable to generate two distinct predicted shapes. At  $i = 1,000$  the two trajectories separate and obtain accurate predictions at the same time. The evolution during training reveals that the model learns one general pattern first and then gradually evolves into two different separable patterns.

Notably, during training, the loss is inverse proportional to the clustering performance (ARI). This means that with training the RNN Seq2Seq model learns to recognize the type of input sequence. We conjecture that separability is correlated with successful training and test the conjecture by repeating the same prediction task with an additional one-hot encoding label for the circle and the ellipse. Indeed, with the additional labels, the model learns these two sequences much faster and they appear separable in the embedding space (**Figure 5**, top). We also verify that the RNN Seq2seq model can encode both spatial and temporal features in the embedding space. For spatial features, we train the model to learn centered and shifted circles. While both trajectories are circles, RNN Seq2seq and the projections of its decoder states to the embedding space distinguish the trajectories well (**Figure 5**, bottom left). For temporal features, we train the network to predict two centered unit circles, sampled



**FIGURE 3** | Projected trajectories in the embedding space spanned by PC3 for the encoder states, decoder states and decoder outputs matrices. Dashed and solid lines represent  $X_{circle}$  and  $X_{ellipse}$ , respectively. Opaque and transparent points denote the starting and the ending points of the trajectory.



**FIGURE 4** | Snapshots of the evolution of trajectories projected to the interpretable embedding space when they undergo training. We use similar line identifiers and colors as in **Figure 3**. At the beginning of the training, two trajectories are very similar and positioned close to each other. As training proceeds the two type of trajectories emerge and separate from each other.

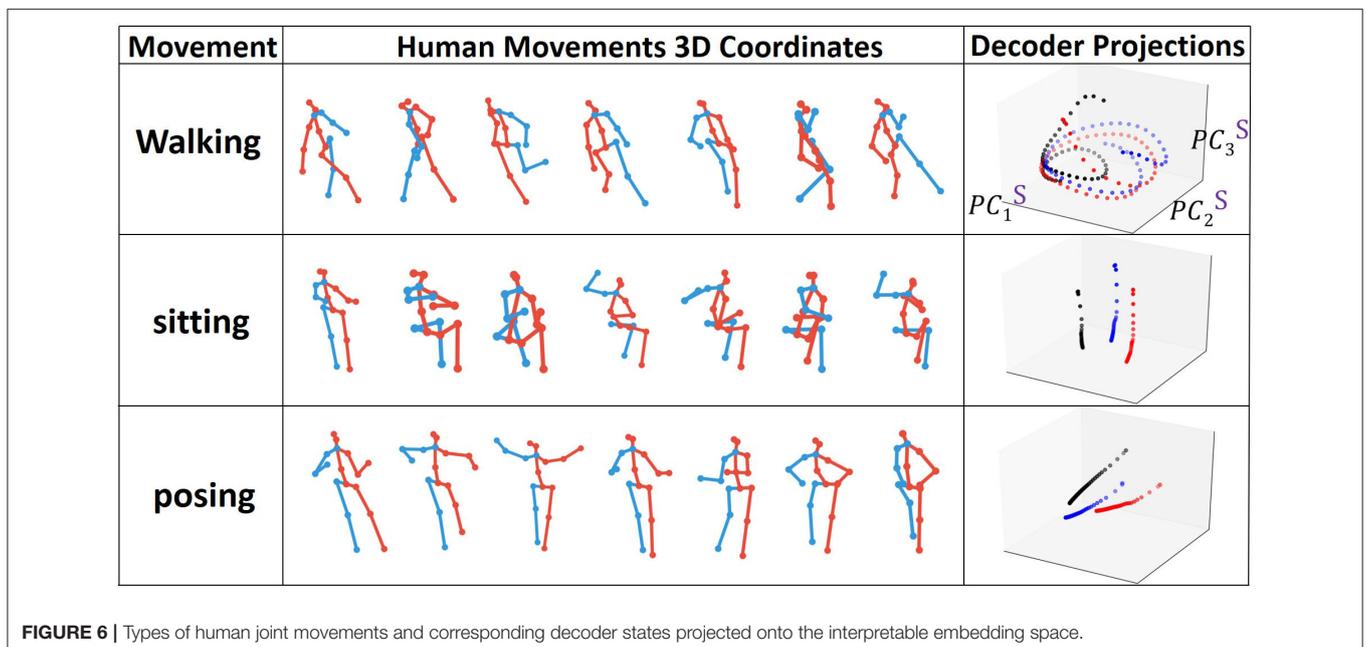
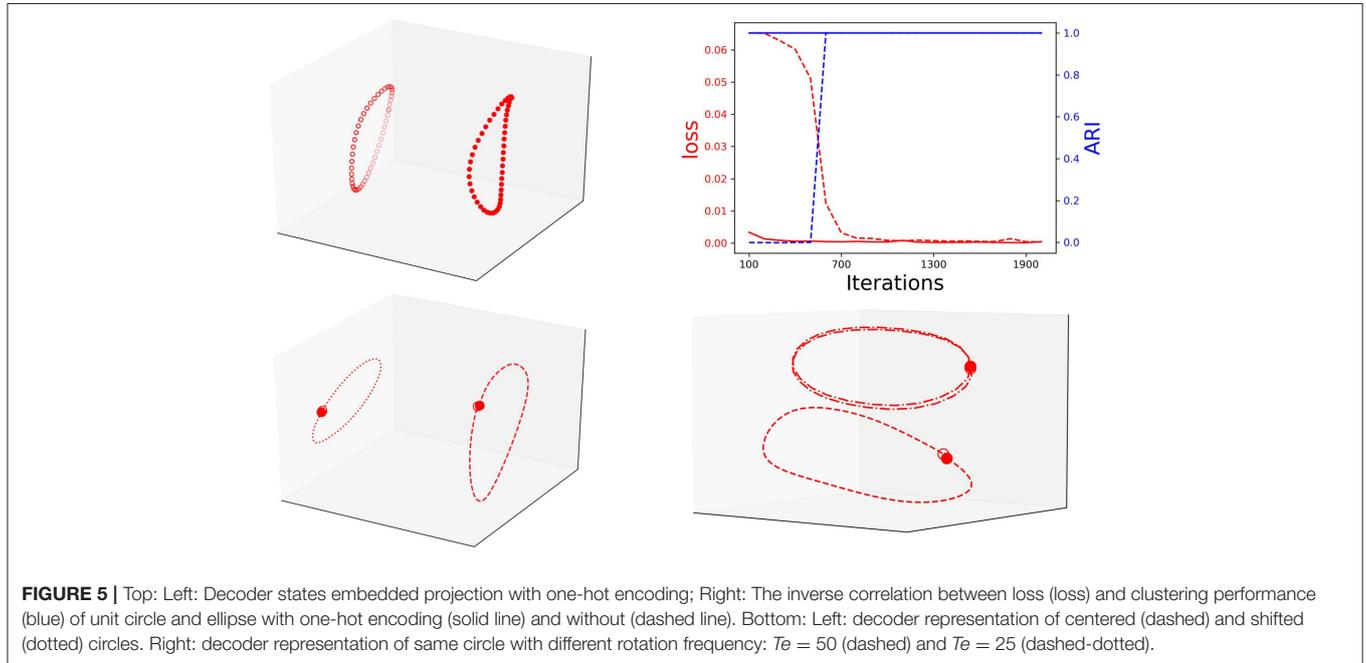
with different rates of  $T_e = 50$  and  $T_e = 25$  (different rotation speeds on the circle). While both circle trajectories coincide in  $x - y$  space, RNN Seq2seq and the projections of its decoder states to the embedding space result with separable attractors with apparent temporal feature difference between the first and second cycles (**Figure 5**, bottom right). These investigations help us to conclude that the embedding space can be effectively used with clustering to represent and assist in the evaluation of different types of learned spatial and temporal features.

### 4. HUMAN BODY JOINTS MOVEMENTS DATA

To test the proposed interpretable embedding methodology on realistic data, we use the Human 3.6 million (H3.6M), which is currently one of the largest publicly available data sets of motion capture data (Ionescu et al., 2014). It includes 7 actors performing 15 various activities such as walking, sitting, and posing. Each movement is repeated in 2 different trials. We use different people for training and testing; 6 of the actors'

motion as the training set and the other actors' motion as a testing set. The human pose is represented as an exponential map representation of each joint, with a special pre-processing of global translation and rotation. The number of features of body joints dynamics is  $M = 54$ . We find that choosing  $T_e = 50$  frames of input data to predict the next  $T_d = 50$  frames corresponds to the best performance. We show results for other lengths of sequences in the **Supplementary Materials**. As a result, the **actual input matrix**, **actual output matrix** and **decoder output matrix** will be  $X, Y, \hat{Y} \in \mathbb{R}^{T_d \times 54}$ .

Seq2Seq model was shown to be successful in such a prediction task (Martinez et al., 2017). We use a similar setup with a single layer GRU of  $N = 1,024$  units sharing the weights between the encoder and the decoder. We find that the setup which shares the weights between the two components converges faster than the setup which does not. However, clustering performance evaluated after the training process converged is similar in both setups. The **encoder and the decoder states** are  $E, D \in \mathbb{R}^{T_e \times N}$ . The cost function is the mean square error between the ground truth and the predicted output



$J = \frac{1}{T_d} \sum_{t=1}^{T_d} (y_t - \hat{y}_t)^2$ . We use a batch size of  $B = 15$  such that every training iteration can contain all actions. We train the model with various gradient descent based optimizers and show our results for ADAM (the fastest converging optimizer for this model). Here we show results with a similar network setup as in previous work, however, our investigations indicate that our method is applicable to other variants of RNN Seq2seq model (e.g., non-parameter sharing RNNs, multi-layer RNNs, etc.).

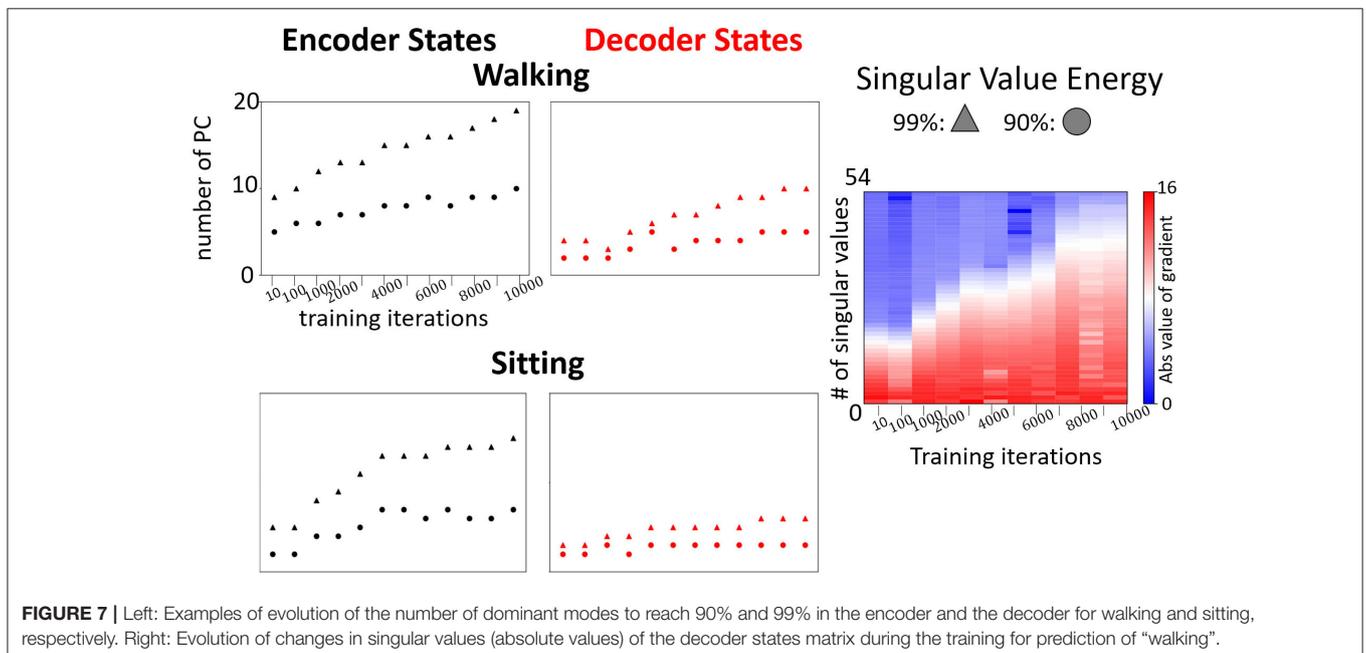
We visualize the projections onto PC3 embedding space for every type of action. In order to see the pattern for the full action, we train the model for prediction and then continuously perform forward propagation on encoder and decoder sequences in a sliding window starting from the very beginning. In **Figure 6**, we show examples of the joints evolution (connected with lines) alongside with the decoder states projection onto PC3 embedding space. We clearly observe that each individual action has its own trajectory evolution (attractor) pattern even in 3D. For example, as expected, walking data corresponds to a periodic circular pattern with a rather fixed period and actions such as sitting correspond to non-periodic trajectories in the embedding space. To get a better understanding of what should be the appropriate dimensions of the embedding space, we monitor the number of dominant PC modes to reach 90 and 99% SVE for encoder and decoder matrices POD for every type of action during training (**Figure 7**).

We observe a low number of dominant modes ( $< 10$  and  $< 20$  on average for 90% and 99% respectively) needed to span most of the energy. The number of modes increases with training and the requirement of the energy threshold (accuracy of the embedding). The rate of the increase depends on the type of movement. For example, as shown in **Figure 7**, the number of dominant modes for “sitting” is much smaller than for

walking. Furthermore, the number of dominant modes required to reach 90% SVE increases more slowly than the number of modes to reach 99%. These results indicate that the model learns additional details with training. However, most of the main features captured in (90%) are learned in the first iterations. Such an observation is supported by analyzing the gradients

**TABLE 1** | Comparison of clustering results on Encoding states, Decoding States, and Joints (RAW) data marked by black, red, and blue colors respectively.

	Training iterations	Dimension	ARI (%)	
			Mean	Std
Encoder states	10	dim = 3	53.7	5.0
		dim = 10	63.0	5.0
		dim = 1,024	64.5	4.2
	4,000	dim = 3	54.9	9.2
		dim = 10	75.5	5.0
		dim = 1024	76.8	2.8
	10,000	dim = 3	60.1	6.6
		dim = 10	81.9	3.5
		dim = 1,024	88.6	8.8
Decoder states	10	dim = 3	26.5	7.4
		dim = 10	29.9	2.5
		dim = 1024	34.6	3.2
	4,000	dim = 3	90.8	6.6
		dim = 10	96.3	3.7
		dim = 1,024	99.1	2.4
	10,000	dim = 3	86.6	9.5
		dim = 10	94.8	5.4
		dim = 1,024	96.2	3.7
Joints data		dim = 3	58.1	4.1
		dim = 54	75.3	4.6



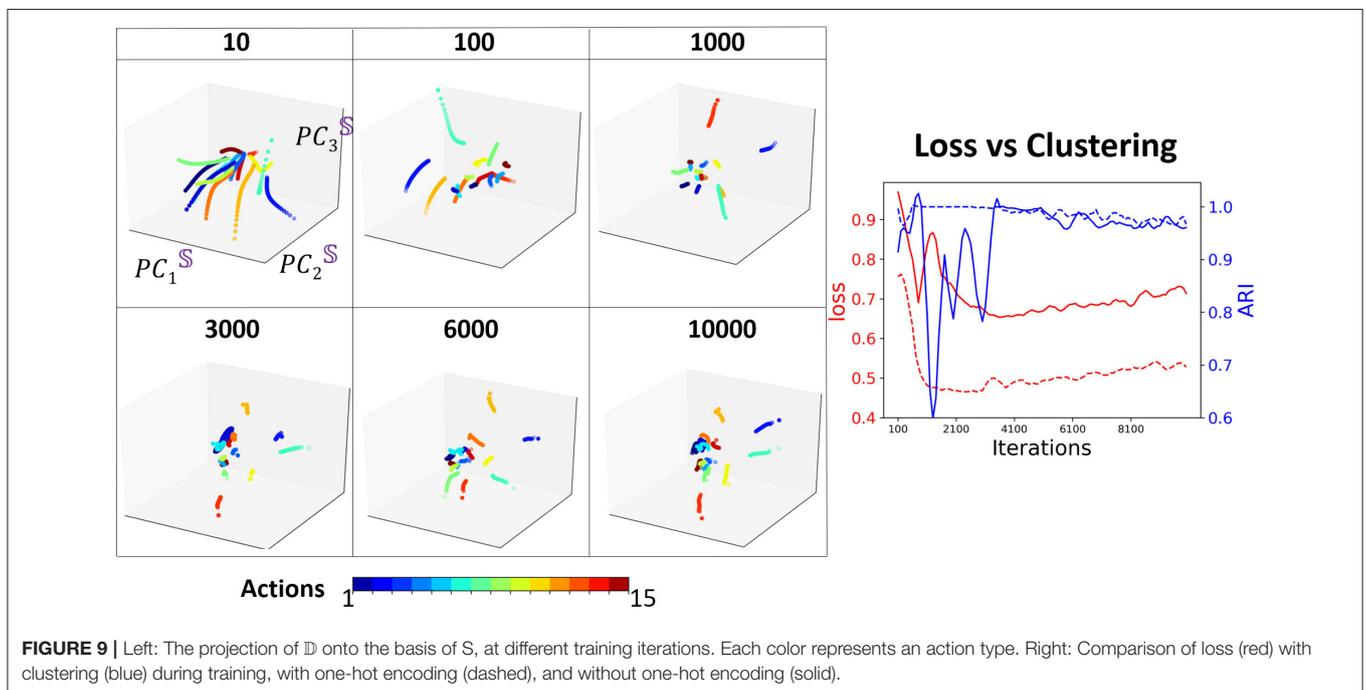
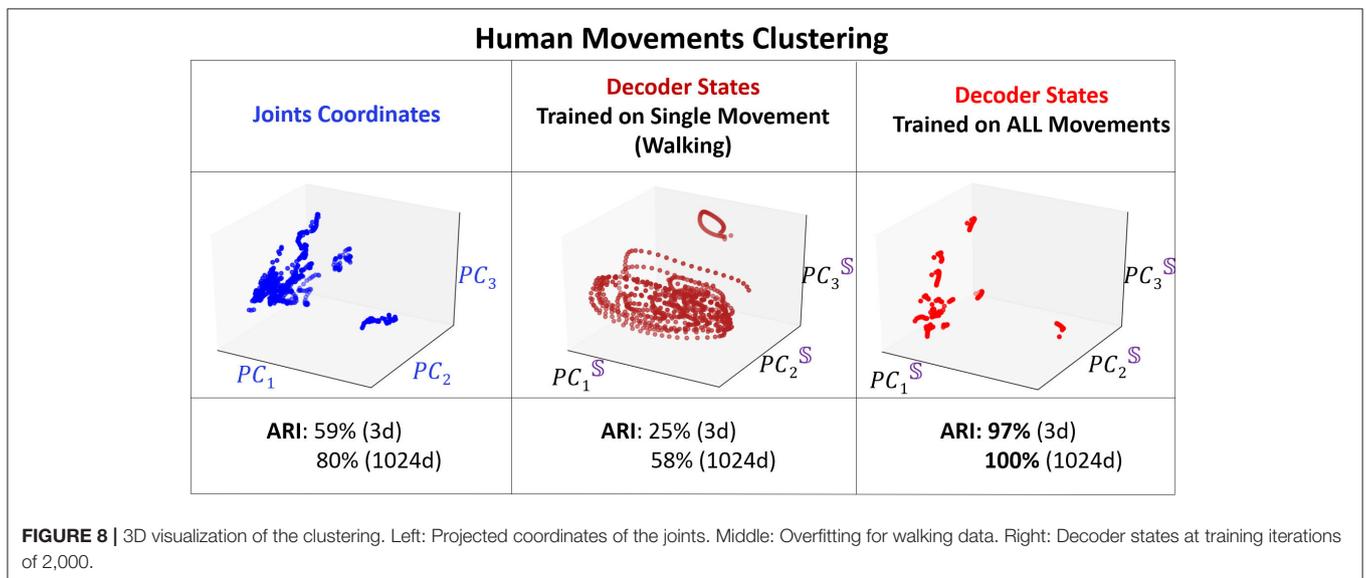
**FIGURE 7** | Left: Examples of evolution of the number of dominant modes to reach 90% and 99% in the encoder and the decoder for walking and sitting, respectively. Right: Evolution of changes in singular values (absolute values) of the decoder states matrix during the training for prediction of “walking”.

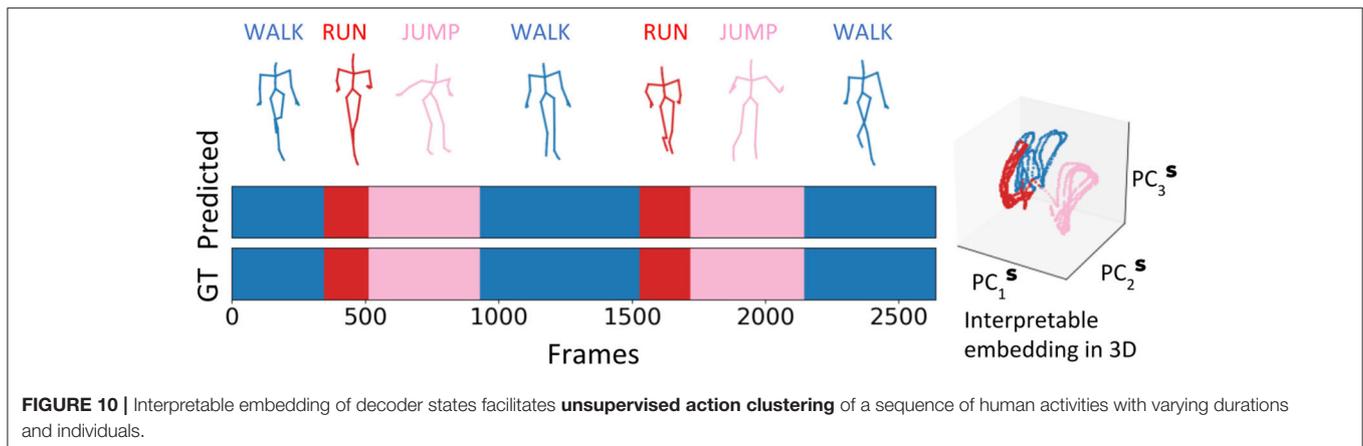
of the singular values during training (Figure 7, right). As training proceeds, additional singular values are gradually being optimized. Notably, we observe that encoder states would have more modes than decoder states. The reason stems from the encoder trajectory starting from the origin and connecting to the decoder attractors in various parts of the embedding space. Such trajectories are hence including mixed characteristics of the attractor and of the path to it resulting in more irregular trajectories requiring additional modes to be represented.

To understand how Seq2Seq forms distinct attractors and differentiates various actions, we construct matrices  $E, D, \hat{Y}$  for all 15 actions for which the model was trained. We apply the K-means++ clustering to  $\mathbb{D}$  to evaluate the separability

of Seq2Seq in the interpretable embedding space and the dimension of the space which provides efficient clustering property (Table 1).

We compare the clustering of the encoder and the decoder trajectories with respect to ARI at different training iterations and for different dimensions (dim = 3 and dim = 1,024) of the embedding space. In addition, we cluster the body joints data (in dim = 3 and dim = 54). Only the clustering of the decoder attractors is able to reach 100% clustering in both measures for dim = 1,024 of the embedding space. The next best clustering performance is for the decoder attractors in dim = 3 (97% ARI which is significantly higher than joints data for dim = 3 and dim = 54) at iterations = 4,000. After this number of iterations





the attractors start to approach each other instead of diverging. The encoder trajectories are not clustered well in any dimension and their clustering does not significantly change with training. In **Figure 8**, we visualize the clusters (for joints coordinates, overfitted trained model, trained model on all movements) in 3D. We show that even visually, the decoder attractors are scattered in the space indicating almost perfect clustering property of PC3 space. If we keep adding the number of PCs to 10, it would be enough to achieve the perfect clustering result.

To understand how training shapes the decoder attractors and their clustering, we mark each distinct movement attractor by a different color and monitor their representation in PC3 embedding space for various training iterations (**Figure 9**). One of our key observations is that the **clustering performance** reaches a **peak** after considerable training (4,000 iterations) which is at the same time that the **validation loss** reaches its **minimum** point. After that, the validation loss increases as clustering becomes inferior. Such performance is known as over-fitting and eventually will lead to a model performing on certain actions only. We demonstrate such a case by training Seq2Seq with more “walking” action data and then test the model with all actions. Indeed, no matter what type of action is given as an input, the decoder states will always be similar to a circular trajectory that represents the “walking” action. Hence, we propose that the **clustering property** of the embedding space could be used as an indication for **sufficient and optimal training** of the model. To show the importance of the separability of clusters, similar to the synthetic case, we compare the results with one-hot encoding data. We find that the model converges much faster and achieves stable high clustering performance very early. In both cases, clustering performance is strongly correlated with clustering property, i.e., when overfitting starts to occur, clustering performance starts to deteriorate.

## 5. UNSUPERVISED ACTION RECOGNITION

To demonstrate the generality and a practical application for our methodology, we utilize the embedding space of RNN Seq2Seq and its clustering property to perform **unsupervised action**

**TABLE 2** | Comparison of clustering performances for continuous multiple actions sequences.

Method	Accuracy (%)
LRR (Liu et al., 2010)	65.08
SSC (Elhamifar and Vidal, 2009)	76.65
ACA (Zhou et al., 2008)	84.50
TSSC-LC (Clopton et al., 2017)	86.08
Ours (Mean, Std)	95.25, 4.34

**recognition** on CMU human body motion capture dataset. We evaluate our methods in two cases: (i) we randomly choose sequences from up to eight different actions (walking, running, jumping, soccer, basketball, washing a window, directing traffic, have a discontinuity signal) and concatenate them together (up to 14,400 frames, 2 mins), following the same rule from (Li et al., 2018), see **Figure 10**. (ii) Since manually concatenated sequences have discontinuity in the prediction, in the second case we choose trials 1 to 14 from subject 86<sup>2</sup>, where each sequence contains continuous multiple actions with variable duration. On average each sequence contains 8,000 frames and time segments of each action sequence are annotated manually (Barbič et al., 2004). For each of the two cases, we train various RNN Seq2Seq predictive models to predict a future sequence task with no usage of labels in training. Our results indicate that the Adversarial Geometry-Aware Encoder-Decoder (AGED) model (Gui et al., 2018), shown to be one of the state-of-art motion prediction model based on RNN Seq2Seq as the predictor, obtains the best prediction results. We, therefore, apply our embedding approach in conjunction with the trained AGED model to perform action recognition. Specifically, all sequences which include a variety of movements, composed using approach (i) or (ii), are scanned with AGED forward propagation performed to collect the decoder states. The decoder states are then projected to the embedding space. Agglomerative clustering with single linkage and cosine similarity is then used to

<sup>2</sup><http://mocap.cs.cmu.edu/search.php?subjectnumber=86>

generate clustering (labeling similar segments to belong to a cluster). With the embedding space and decoder attractors within it, **unsupervised activity recognition** of sequences can be performed. The only information required is the number of unique actions (clusters) that we would like to obtain (see **Figure 10** and **Supplementary Materials** for examples). Testing the algorithm on composed actions sequences produces the following performance: in case (i), the approach achieves above 98% frame-level accuracy on average; in case (ii), we evaluate our accuracy on by frame-level as well and compare with previously reported methods (Clopton et al., 2017). Our method reaches an accuracy of 95.25% on average and outperforms previous methods (**Table 2**).

## 6. CONCLUSION

We propose a novel construction of interpretable embedding for the hidden states of the Seq2Seq model. The embedding clarifies the role of the encoder and the decoder components of Seq2Seq such that encoder embedded trajectories direct the evolution from the origin to the decoder trajectories represented as attractors. Our findings indicate a remarkable property of networks: the network trained to predict the future evolution of a sequence self-organizes the hidden units representation into separate identities. The identities are revealed through our proposed embedding and clustering. We demonstrate the construction and the utilization of the embedding space on both synthetic and human body joints datasets. We show that the embedding can inspect training and determine the goodness of fit. Furthermore, we present an algorithm for unsupervised clustering of any spatio-temporal features. It utilizes training of Seq2Seq to predict future actions and analyzes the learned representation with the interpretable embedding to generate clustering. We show that such an approach allows performing action recognition on human body pose data, i.e., to generate

an unsupervised time-segmentation and clustering of human movements. The algorithm achieves significantly more robust performance and accuracy than previously proposed approaches.

## DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: <http://mocap.cs.cmu.edu/>.

## AUTHOR CONTRIBUTIONS

ES and KS conceptualized and designed the study. ES acquired the funding for the study. KS organized the data. KS implemented the algorithms and performed the experiments. KS and ES validated the results. KS and ES wrote the original draft of the manuscript. KS and ES edited and revised the manuscript, read and approved the submitted version.

## ACKNOWLEDGMENTS

We acknowledge the support of the National Science Foundation under Grant No. DMS-1361145, and Washington Research Foundation Innovation Fund. The authors would also like to acknowledge the partial support by the Departments of Applied Mathematics, Electrical & Computer Engineering, the Center of Computational Neuroscience (CNC), and the eScience Center at the University of Washington in conducting this research. This manuscript has been released as a pre-print Su and Shlizerman (2019).

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2020.00070/full#supplementary-material>

## REFERENCES

- Alain, G., and Bengio, Y. (2016). Understanding intermediate layers using linear classifier probes. *arXiv [Preprint]*. *arXiv:1610.01644*.
- Arthur, D., and Vassilvitskii, S. (2007). “k-means++: the advantages of careful seeding,” in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, LA: Society for Industrial and Applied Mathematics), 1027–1035.
- Barbič, J., Safonova, A., Pan, J.-Y., Faloutsos, C., Hodgins, J. K., and Pollard, N. S. (2004). “Segmenting motion capture data into distinct behaviors,” in *Proceedings of Graphics Interface 2004* (Waterloo, ON: Canadian Human-Computer Communications Society), 185–194.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., et al. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv [Preprint]*. *arXiv:1406.1078*. doi: 10.3115/v1/D14-1179
- Clopton, L., Mavroudi, E., Tsakiris, M., Ali, H., and Vidal, R. (2017). Temporal subspace clustering for unsupervised action segmentation. *CSMR REU*. 1–7.
- Collins, J., Sohl-Dickstein, J., and Sussillo, D. (2016). Capacity and trainability in recurrent neural networks. *arXiv [Preprint]*. *arXiv:1611.09913*.
- Elhamifar, E., and Vidal, R. (2009). “Sparse subspace clustering,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (Miami, FL: IEEE), 2790–2797. doi: 10.1109/CVPR.2009.5206547
- Farrell, M. S., Recanatesi, S., Lajoie, G., and Shea-Brown, E. (2019). Dynamic compression and expansion in a classifying recurrent network. *bioRxiv [Preprint]*. 564476. doi: 10.1101/564476
- Foerster, J. N., Gilmer, J., Sohl-Dickstein, J., Chorowski, J., and Sussillo, D. (2017). “Input switched affine networks: an RNN architecture designed for interpretability,” in *Proceedings of the 34th International Conference on Machine Learning-Volume 70* (Sydney, NSW), 1136–1145.
- Fragkiadaki, K., Levine, S., Felsen, P., and Malik, J. (2015). “Recurrent network models for human dynamics,” in *Proceedings of the IEEE International Conference on Computer Vision* (Las Condes), 4346–4354. doi: 10.1109/ICCV.2015.494
- Graves, A., Mohamed, A.-R., and Hinton, G. (2013). “Speech recognition with deep recurrent neural networks,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (Vancouver, BC: IEEE), 6645–6649. doi: 10.1109/ICASSP.2013.6638947
- Gui, L.-Y., Wang, Y.-X., Liang, X., and Moura, J. M. (2018). “Adversarial geometry-aware human motion prediction,” in *Proceedings of the European Conference on Computer Vision (ECCV)* (Munich), 786–803. doi: 10.1007/978-3-030-01225-0\_48

- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780. doi: 10.1162/neco.1997.9.8.1735
- Ionescu, C., Papava, D., Olaru, V., and Sminchisescu, C. (2014). Human3.6M: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. Pattern Anal. Mach. Intell.* 36, 1325–1339. doi: 10.1109/TPAMI.2013.248
- Jain, A., Zamir, A. R., Savarese, S., and Saxena, A. (2016). “Structural-RNN: deep learning on spatio-temporal graphs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5308–5317. doi: 10.1109/CVPR.2016.573
- Karpathy, A., Johnson, J., and Fei-Fei, L. (2015). Visualizing and understanding recurrent networks. *arXiv [Preprint]*. arXiv:1506.02078.
- Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., et al. (2017). Interpretability beyond feature attribution: quantitative testing with concept activation vectors (TCAV). *arXiv [Preprint]*. arXiv:1711.11279.
- Li, C., Zhang, Z., Sun Lee, W., and Hee Lee, G. (2018). “Convolutional sequence to sequence model for human dynamics,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Salt Lake City, UT). doi: 10.1109/CVPR.2018.00548
- Liu, G., Lin, Z., and Yu, Y. (2010). “Robust subspace segmentation by low-rank representation,” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)* (Haifa), 663–670.
- Luong, M.-T., Pham, H., and Manning, C. D. (2015). Effective approaches to attention-based neural machine translation. *arXiv [Preprint]*. arXiv:1508.04025. doi: 10.18653/v1/D15-1166
- Maaten, L. V. D., and Hinton, G. (2008). Visualizing data using T-SNE. *J. Mach. Learn. Res.* 9, 2579–2605.
- Martinez, J., Black, M. J., and Romero, J. (2017). “On human motion prediction using recurrent neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Honolulu), 2891–2900. doi: 10.1109/CVPR.2017.497
- Recanatani, S., Farrell, M., Lajoie, G., Deneve, S., Rigotti, M., and Shea-Brown, E. (2018). Signatures and mechanisms of low-dimensional neural predictive manifolds. *bioRxiv [Preprint]*. 471987.
- Shlizerman, E., Schroder, K., and Kutz, J. N. (2012). Neural activity measures and their dynamics. *SIAM J. Appl. Math.* 72, 1260–1291. doi: 10.1137/110843630
- Strobelt, H., Gehrmann, S., Behrisch, M., Perer, A., Pfister, H., and Rush, A. M. (2019). Seq2Seq-Vis: a visual debugging tool for sequence-to-sequence models. *IEEE Trans. Visual. Comput. Graph.* 25, 353–363. doi: 10.1109/TVCG.2018.2865044
- Su, K., and Shlizerman, E. (2019). Dimension reduction approach for interpretability of sequence to sequence recurrent neural networks. *arXiv [Preprint]*. arXiv:1905.12176.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). “Sequence to sequence learning with neural networks,” in *The Paper Appears in Advances in Neural Information Processing Systems 27*, eds Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Montreal, QC: Neural Information Processing Systems), 3104–3112.
- Zeiler, M. D., and Fergus, R. (2014). “Visualizing and understanding convolutional networks,” in *European Conference on Computer Vision* (Zurich: Springer), 818–833. doi: 10.1007/978-3-319-10590-1\_53
- Zhou, F., De la Torre, F., and Hodgins, J. K. (2008). “Aligned cluster analysis for temporal segmentation of human motion,” in *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition* (Amsterdam: IEEE), 1–7. doi: 10.1109/AFGR.2008.4813468

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Su and Shlizerman. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.