Learning Robotic Handwriting with Convolutional Image-to-Motion Encoder-Decoder Networks

Barry Ridge Dept. of Automatics, Biocybernetics, and Robotics, Jožef Stefan Institute Ljubljana, Slovenia barry.ridge@ijs.si

ABSTRACT

Learning to recognize and reproduce handwriting is a familiar skill taught to every educated human being but is challenging to teach to a robot given its tight coupling between perception and motion. In this work, we address the specific task of recognizing digits in single images and reproducing the digits in the form of robot end-effector trajectories encoded as dynamic movement primitives (DMPs) used to control the pen strokes. Here we present a convolutional image-to-motion encoder-decoder deep neural network architecture that takes the raw digit images as input and produces the DMP parameters as output, learning a mapping between the two as a latent representation. The architecture is tested on several challenging noisy digit datasets under different training regimes and compared to an architecture without convolutional layers in the image encoder where it is shown to provide robust results for the digit writing task.

KEYWORDS

deep neural networks, dynamic movement primitives

1 INTRODUCTION

Effectively learning to predict action mappings directly from perceptual input is a highly challenging problem in robotics research that has seen a broad variety of approaches attempting to solve it in different settings. The particular setting under consideration in this work is depicted in Fig. 1, in which a robot must learn direct mappings between handwritten characters in input images and the motion trajectories needed to draw them. In previous work we proposed a fully-connected encoder-decoder network architecture [6] that used dynamic movement primitives (DMPs) [5] for movement representation and this proved to be an effective choice both for representation and learning with the neural network and ultimately for control of the robot when drawing the actual digits. The fully-connected architecture, however, was not ideal for image representation.

Here we investigate a different architecture that combines the benefits of convolutional layers for image encoding with those

IS2019, October 7–11, 2019, Ljubljana, Slovenia, Europe

© 2019 Copyright held by the owner/author(s).

Rok Pahič Dept. of Automatics, Biocybernetics, and Robotics, Jožef

> Stefan Institute Ljubljana, Slovenia rok.pahic@ijs.si

of a fully-connected encoder-decoder architecture for DMP parameter prediction and image-to-motion representation in a lowdimensional latent space. This combination allows for relatively robust prediction compared to the previously proposed architecture, even when the input images are heavily corrupted by noise. The use of convolutional layers has the added benefit of significantly reducing the number of network parameters and by pre-training these layers on images from a similar image domain, the learning process is further improved.



Figure 1: Writing digits with a robot using image-to-motion encoder-decoder network prediction.

Autoencoders [3], as well as variational autoencoders [4], have been demonstrated to be quite effective when it comes to calculating DMP-based representations of human motion. Since our focus is on learning direct mappings between images and actions, instead of using such autoencoder networks in which the DMP encoding occurs in the latent space, we use an encoder-decoder architecture in which the image is encoded from the input layer, the DMP parameters are predicted at the output layer and the transformation and generalization of the image-to-motion representation occurs in the low-dimensional latent space. Encoder-decoder networks in combination with convolutional layers have proven to be useful in computer vision. A well-known example is SegNet [1], in which pretrained convolutional layers from a convolutional neural network (CNN) were adapted to form a fully-convolutional encoder-decoder architecture for semantic pixel-wise segmentation.

Usually when CNNs are used for supervised learning of perceptionaction couplings, they are used i combination with another neural network in two separately trainable parts. In [9], Yang et al. first used a deep convolutional autoencoder for finding camera image features and then in combination with recorded robot angles, formed sequences for the learning task dynamics with a time delay neural network. Pervez et al. [7] used a pre-trained CNN for finding task parameters from input images, while using a another fully-connected neural network to learn to generate forcing terms

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

IS2019, October 7-11, 2019, Ljubljana, Slovenia, Europe



Figure 2: The CIMEDNet architecture.

from the clock signal and task parameters, before combining both networks in an end-to-end training scheme. Both of these two examples produce the next step from the image of the current step while working in online loop, whereas our method, by contrast, uses just single images for generating entire trajectories.

2 CONVOLUTIONAL IMAGE-TO-MOTION ENCODER-DECODER NETWORKS

The structure of the data under consideration in this work is the same as in [6] where the input and output data pairs take the form $\mathbf{D} = \{\mathbf{C}_j, \mathbf{M}_j\}_{j=1}^M$ where M is the number of input and output training pairs, $\mathbf{C}_j \in \mathbb{R}^{H \times W}$ are the input images of width W and height H, and \mathbf{M}_j the corresponding movements associated with each image, i. e. $\mathbf{M}_j = \{y_{i,j}, t_{i,j}\}_{i=1}^{T_j}$. Here $y_{i,j} \in \mathbb{R}^d$ are the vectors describing the movement's degrees of freedom, e.g. Cartesian positions or joint angles, $t_{i,j} \in \mathbb{R}$ the measurement times for the *j*-th movement, and *d* is the number of degrees of freedom. However, it should be noted that in this paper, we convert the movements \mathbf{M}_j to DMPs and construct all of the datasets used to train the network models as follows: $\mathbf{D}' = \{\mathbf{C}_j, \mathbf{k}_j\}_{j=1}^M$, where \mathbf{k}_j are the DMP parameters calculated for each movement \mathbf{M}_j and are represented as

$$\mathbf{k}_{j} = \left\{ \{w_{k}\}_{k=1}^{N}, \ \tau, \ g, \ y_{0} \right\}.$$
(1)

The construction of DMPs and the nature of the parameters $\{w_k\}_{k=1}^N$, τ , g and y_0 are explained in detail in the following subsection.

2.1 Motion Representation with DMPs

Letting a time-dependent movement trajectory be denoted as $y(t) \in \mathbb{R}^d$, a DMP specifying this trajectory is given by the following system of differential equations

$$\tau \dot{z} = \alpha_z (\beta_z (g - y) - z) + \operatorname{diag}(g - y_0) \mathbf{F}(x), \qquad (2)$$

$$\tau \dot{y} = z, \tag{3}$$

where $y_0 \in \mathbb{R}^d$ is the initial position on the trajectory, $g \in \mathbb{R}^d$ the final position on the trajectory, $\operatorname{diag}(g - y_0) \in \mathbb{R}^{d \times d}$ a diagonal matrix with components of vector $g - y_0$ on the diagonal, $\mathbf{F}(x) \in \mathbb{R}^d$ a nonlinear forcing term, $z \in \mathbb{R}^d$ a scaled velocity of motion, and $x \in \mathbb{R}$ the phase defined by the following equation

$$\tau \dot{x} = -\alpha_x x. \tag{4}$$

The phase *x* is used instead of time to avoid explicit time dependency. It is fully defined by setting its initial value to x(0) = 1. Eq. system (2) – (4) constitutes a *dynamic movement primitive* (DMP).

tem (2) – (3) becomes critically damped and y, z monotonically converge to a unique attractor point at y = g, z = 0. The forcing term $\mathbf{F}(x)$ is usually defined as a linear combination of radial basis functions

$$\mathbf{F}(x) = \frac{\sum_{k=1}^{N} w_k \Psi_k(x)}{\sum_{k=1}^{N} \Psi_k(x)} x,$$
(5)

$$\Psi_k(x) = \exp\left(-h_k \left(x - c_k\right)^2\right),\tag{6}$$

where c_k are the centers of Gaussians distributed along the phase of the trajectory, and h_k their widths. The role of **F** is to adapt the dynamics of (2) – (3) to the desired trajectory y(t), thus enabling the system to reproduce any smooth movement from the initial position y_0 to the final configuration g. This can be accomplished by computing the free parameters $w_k \in \mathbb{R}^d$ using regression techniques. See [8] for more details.

 α_z , β_z , and α_x are usually constants that do not change between movements. Thus the neural network needs to learn the other parameters of differential equation system (2) – (4) to fully specify a DMP as defined in Equation (1).

2.2 Network Architecture

In our improved architecture, images are encoded via convolutional layers that are pre-trained as part of a basic CNN classifier that was trained on the original MNIST dataset. The input is a $40 \times 40 \times 1$ grayscale pixel image, followed by a convolutional layer with 5×5 kernel size and 10 feature maps, a convolutional layer, a fully-connected layer of size 320, a fully-connected layer of size 50 and the output layer of size 10 matching the number of digits. After training the classifier, the fully-connected layers are removed and the convolutional layers are retained and are used to form the first layers of the encoder in our proposed architecture. These two convolutional layers are followed by two added fully-connected layers with sizes of 600 neurons and 200 neurons, illustrated on the left side of Fig. 2.

Following the bottleneck of the network that forms the latent space representation, a decoder is formed via a number of fullyconnected layers that gradually expand the number of units in each layer until the final output layer which has a size set to 55 units in order to match the DMP parameters $\{w_k\}_{k=1}^N$, τ , g and y_0 . The layers of the decoder are illustrated on the right side of 2 starting with the bottleneck of size 20, followed by a layer of size 35 and finishing with the output layer. This is the same decoder structure as used [6] and we retain it here as-is, having found it to be effective throughout our experiments for this particular use case. The cost function used to evaluate the output of the network is the same as that of Equation (9) in [6], which is defined for the *j*-th DMP as follows:

$$E_{p}(j) = \frac{1}{2} \left(\sum_{k=1}^{N} \|w_{k} - w_{k,j}\|^{2} + (\tau - \tau_{j})^{2} + \|g - g_{j}\|^{2} + \|y_{0} - y_{0,j}\|^{2} \right),$$
(7)

where $\{\{w_k\}_{k=1}^N, \tau, g, y_0\}$ denotes the output of the neural network and $\{\{w_{k,j}\}_{k=1}^N, \tau_j, g_j, y_{0,j}\}$ the DMP parameters from the training data $\mathbf{k}_j \in \mathbf{D}'$. For further details on the gradient calculations required for minimizing the cost function via backpropagation we refer the reader to [6].

3 EXPERIMENTS

In our experiments, we trained both the fully-connected image-tomotion encoder-decoder architecture (IMEDNet) and the convolutional architecture (CIMEDNet) on various digit image and motion trajectory datasets. The IMEDNet architecture was the same as described in [6] with fully-connected hidden layer sizes of 1500, 1300, 1000, 600, 200, 20, and 35 neurons, respectively. The CIMEDNet architecture was as described in Section 2.2 and as illustrated in Fig. 2.

In the case of CIMEDNet, we also experimented with either freezing the convolutional layer weights or training the entire network end-to-end. The results for these different training regimes are cataloged in Table 1.

3.1 Datasets

In order to construct D, we employed the same scheme described in [6] to generate 40×40 images of synthetically written digits and associated two-dimensional artificial writing trajectory movements. Briefly, the synthetic trajectory data was generated using a combination of straight lines and elliptic arcs. These geometric elements were used to generate grayscale digit images and their paramaters were varied according to a uniform distribution. The resulting images were processed with a Gaussian filter and some moderate salt-and-pepper noise was added to the foreground pixels. Finally, both the generated trajectories and the resulting images were transformed using affine transformations composed of translation, rotation, scaling, and shearing. These parameters were again taken from a uniform distribution. For the DMP representation of the trajectories, 25 radial-basis functions were selected for every dimension. The weights of these basis functions form together with the common time constant (1 parameter) and the start and the goal values of a planar movement $(2 \times 2 \text{ parameters})$, the full set of 55 DMP parameters that represent the motion. Using this procedure, several datasets were generated both with and without similar noise as used in the noisy MNIST (n-MNIST) datasets [2] as follows:

- s-MNIST: 2000 pairs of images and trajectories without any added noise were generated for each digit, for a total of 20000 samples that were split in a 70%/15%/15% ratio between training/validation/test data,
- s-MNIST-AWGN-19.0: 300 samples per digit/3000 total samples, using additive white gaussian noise with a signal-to-noise ratio of 19.0,
- s-MNIST-AWGN-9.5: 300 samples per digit/3000 total samples, using additive white gaussian noise with a signal-to-noise ratio of 9.5,
- s-MNIST-MB: 300 samples per digit/3000 total samples, using a motion blur filter emulating a linear motion of the camera of 5 pixels and a 15 degree motion in the counter-clockwise direction,

s-MNIST-RC-AWGN: 300 samples per digit/3000 total samples, using a contrast range scaled down to half as well as additive white gaussian noise with a signal-to-noise ratio of 9.5.

It should be emphasized that in the results that follow, only the s-MNIST dataset was used for training the presented models.

3.2 Results

The main quantitative results are presented in Table 1 while qualitative results for selected samples are presented in Fig. 3. After training on the noiseless s-MNIST dataset each of the models were tested on all five of the noiseless and noisy s-MNIST datasets described in the previous section. The CIMEDNet architecture was trained with two separate training regimes in which the convolutional layer weights were frozen and the models were trained end-to-end respectively. For the quantitative evaluation, dynamic time warping was used to measure the mean pointwise pixel distance between the trajectories generated by the DMPs predicted by the networks from the digit images and the actual digit trajectories.

Table 1: DMP reconstruction statistics. The results are in pixels. The best result for each dataset is highlighted in boldface.

IMEDNet (End-to-End)	CIMEDNet (Frozen Conv.)	CIMEDNet (End-to-End)
0.22 ± 0.08	0.26 ± 0.10	0.19 ± 0.08
0.56 ± 0.20	0.54 ± 0.20	0.36 ± 0.14
1.66 ± 0.60	1.48 ± 0.55	1.02 ± 0.45
0.35 ± 0.15	0.47 ± 0.25	0.36 ± 0.12
2.32 ± 0.77	2.19 ± 0.76	1.93 ± 0.66
	IMEDNet (End-to-End) 0.22 ± 0.08 0.56 ± 0.20 1.66 ± 0.60 0.35 ± 0.15 2.32 ± 0.77	IMEDNet (End-to-End)CIMEDNet (Frozen Conv.) 0.22 ± 0.08 0.26 ± 0.10 0.56 ± 0.20 0.54 ± 0.20 1.66 ± 0.60 1.48 ± 0.55 0.35 ± 0.15 0.47 ± 0.25 2.32 ± 0.77 2.19 ± 0.76

As can be seen in Table 1, the CIMEDNet model that is trained end-to-end significantly out-performs the IMEDNet model on both the noiseless s-MNIST dataset and on most of the noisy s-MNIST datasets, apart from the dataset featuring motion blur noise. We reason that this may be due to the fact that motion blur can significantly distort overall object shape and edge profiles and given that convolutional neural networks function the basis of exploiting hierarchies of image filters often heavily represented by edge detectors, this may impact on their effectiveness in such circumstances. The CIMEDNet that was trained with frozen convolutional layers also fared well, beating the IMEDNet model on the same noisy datasets despite not scoring as well on the noiseless dataset. This indicates that the feature detectors in the convolutional layers allow for more robust generalization whereas fully-connected layers are more inclined to overfit.

The qualitative result samples in Fig. 3, are also interesting. Here, original trajectories are shown in blue whereas trajectories calculated by the neural networks are shown in red and samples in matching dataset rows are identical for a fair comparison between each network. Results using the s-MNIST-RC-AWGN dataset are omitted as the noise levels are so pathologically difficult that the qualitative results are comparatively worthless. However, the CIMEDNet



Figure 3: Example results for IMEDNet (rows 1, 2, 5, 6, 9, 10, 13 & 14) & CIMEDNet trained end-to-end (rows 3, 4, 7, 8, 11, 12, 15 & 16). Rows 1-4: s-MNIST, rows 5-8: s-MNIST-AWGN-19.0, rows 9-12: s-MNIST-AWGN-9.5 and rows 13-16: s-MNIST-MB.

model often performs surprisingly well given that it was not trained or fine-tuned on the noisy data. Both models appear to produce highly legible writing trajectories that closely match the actual trajectories in the case of the s-MNIST-MB dataset, but the CIMEDNet model is demonstrably superior to IMEDNet in many cases with the s-MNIST-AWGN-19.0 and s-MNIST-AWGN-9.5 data, producing much more legible results and demonstrating the robustness of the convolutional layers in dealing with even high noise levels.

4 CONCLUSIONS AND FUTURE WORK

We have presented an extended form of an encoder-decoder neural network for image-to-motion prediction that employs convolutional layers in the encoder in order to make the image recognition component more robust to noisy input. We have demonstrated that this architecture outperforms its predecessor on a variety of different kinds of noise. Regarding future work, we intend to further expand the capabilities of this model by incorporating layers from more powerful pre-trained CNN models into the encoder and training the network on more challenging image sets. One challenge here lies in either finding suitable image datasets that include trajectory information in their target outputs or in finding other means of producing images with corresponding motion trajectories, e.g. by gathering both in a robot simulation environment.

ACKNOWLEDGMENTS

This work has received funding from the EU's Horizon 2020 RIA AUTOWARE (GA no. 723909); the Slovenian Research Agency under GA no. J2-7360; JSPS KAKENHI JP16H06565; NEDO; the Commissioned Research of NICT; the NICT Japan Trust (International research cooperation program); and JST-Mirai Program Grant Number JPMJMI18B8, Japan.

REFERENCES

- V. Badrinarayanan, A. Kendall, and R. Cipolla. 2017. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Transactions* on Pattern Analysis and Machine Intelligence 39, 12 (Dec. 2017), 2481–2495.
- [2] Saikat Basu, Manohar Karki, Sangram Ganguly, Robert DiBiano, Supratik Mukhopadhyay, Shreekant Gayaka, Rajgopal Kannan, and Ramakrishna Nemani. 2017. Learning Sparse Feature Representations Using Probabilistic Quadtrees and Deep Belief Nets. *Neural Processing Letters* 45, 3 (June 2017), 855–867.
- [3] N. Chen, J. Bayer, S. Urban, and P. van der Smagt. 2015. Efficient Movement Representation by Embedding Dynamic Movement Primitives in Deep Autoencoders. In 2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids). 434–440.
- [4] N. Chen, M. Karl, and P. van der Smagt. 2016. Dynamic Movement Primitives in Latent Space of Time-Dependent Variational Autoencoders. In 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). 629–636.
- [5] Auke Jan Ijspeert, Jun Nakanishi, Heiko Hoffmann, Peter Pastor, and Stefan Schaal. 2013. Dynamical Movement Primitives: Learning Attractor Models for Motor Behaviors. *Neural computation* 25, 2 (2013), 328–373.
- [6] Rok Pahič, Andrej Gams, Aleš Ude, and Jun Morimoto. 2018. Deep Encoder-Decoder Networks for Mapping Raw Images to Dynamic Movement Primitives. In 2018 IEEE International Conference on Robotics and Automation (ICRA). Brisbane, Australia, 5863–5868.
- [7] A. Pervez, Y. Mao, and D. Lee. 2017. Learning Deep Movement Primitives Using Convolutional Neural Networks. In 2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids). 191–197.
- [8] A. Ude, A. Gams, T. Asfour, and J. Morimoto. 2010. Task-Specific Generalization of Discrete and Periodic Dynamic Movement Primitives. *IEEE Transactions on Robotics* 26, 5 (Oct. 2010), 800–815.
- [9] P. C. Yang, K. Sasaki, K. Suzuki, K. Kase, S. Sugano, and T. Ogata. 2017. Repeatable Folding Task by Humanoid Robot Worker Using Deep Learning. *IEEE Robotics* and Automation Letters 2, 2 (April 2017), 397–403.