

Translating Videos to Commands for Robotic Manipulation with Deep Recurrent Neural Networks

Anh Nguyen¹, Dimitrios Kanoulas¹, Luca Muratore^{1,2}, Darwin G. Caldwell¹, and Nikos G. Tsagarakis¹

Abstract—We present a new method to translate videos to commands for robotic manipulation using Deep Recurrent Neural Networks (RNN). Our framework first extracts deep features from the input video frames with a deep Convolutional Neural Networks (CNN). Two RNN layers with an encoder-decoder architecture are then used to encode the visual features and sequentially generate the output words as the command. We demonstrate that the translation accuracy can be improved by allowing a smooth transaction between two RNN layers and using the state-of-the-art feature extractor. The experimental results on our new challenging dataset show that our approach outperforms recent methods by a fair margin. Furthermore, we combine the proposed translation module with the vision and planning system to let a robot perform various manipulation tasks. Finally, we demonstrate the effectiveness of our framework on a full-size humanoid robot WALK-MAN.

I. INTRODUCTION

The ability to perform actions based on observations of human activities is one of the major challenges to increase the capabilities of robotic systems [1]. Over the past few years, this problem has been of great interest to researchers and remains an active field in robotics [2]. By understanding human actions, robots may be able to acquire new skills, or perform different tasks, without the need for tedious programming. It is expected that the robots with these abilities will play an increasingly more important role in our society in areas such as assisting or replacing humans in disaster scenarios, taking care of the elderly, or helping people with everyday life tasks.

In this paper, we argue that there are two main capabilities that a robot must develop to be able to replicate human activities: *understanding* human actions, and *imitating* them. The imitation step has been widely investigated in robotics within the framework of learning from demonstration (LfD) [3]. In particular, there are two main approaches in LfD that focus on improving the accuracy of the imitation process: kinesthetic teaching [4] and motion capture [5]. While the first approach needs the users to physically move the robot through the desired trajectories, the second approach uses a bodysuit or camera system to capture human motions. Although both approaches successfully allow a robot to imitate a human, the number of actions that the robot can learn is quite limited due to the need of using expensively physical systems (i.e., real robot, bodysuit, etc.) to capture the training data [4] [5].

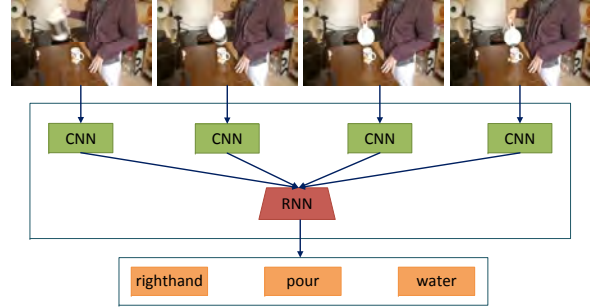


Fig. 1. Translating videos to commands for robotic manipulation.

The *understanding* step, on the other hand, receives more attention from the computer vision community. Two popular problems that receive a great deal of interest are video classification [6] and action recognition [7]. However, the outputs of these problems are discrete (e.g., the action classes used in [7] are “diving”, “biking”, “skiing”, etc.), and do not provide further meaningful clues that can be used in robotic applications. Recently, with the rise of deep learning, the video captioning problem [8] has become more feasible to tackle. Unlike the classification or detection tasks, the output of the video captioning task is a natural language sentence, which is potentially useful in robotic applications.

Inspired by the recent advances in computer vision, this paper describes a deep learning framework that translates an input video to a command that can be used in robotic applications. While the field of LfD [3] focuses mainly on the imitation step, we focus on the understanding step, but our method also allows the robot to perform useful tasks via the output commands. Our goal is to bridge the gap between computer vision and robotics, by developing a system that helps the robot understand human actions, and use this knowledge to complete useful tasks. In particular, we first use CNN to extract deep features from video frames, then two RNN layers are used to learn the relationship between the visual features and the output command. Unlike the video captioning problem [8] which describes the output sentence in a natural language form, we use a grammar-free form to describe the output command. We show that our solely neural architecture further improves the state of the art, while its output can be applied in real robotic applications. Fig. 1 illustrates the concept of our method.

Next, we review the related work in Section II, then describe our network architecture in Section III. In Section IV, we present the experimental results on the new challenging dataset, and on the full-size humanoid robot. Finally, we discuss the future work and conclude the paper in Section V.

¹Advanced Robotics Department, Istituto Italiano di Tecnologia, Italy.

²School of Electrical and Electronic Engineering, The University of Manchester, UK.

{Anh.Nguyen, Dimitrios.Kanoulas, Luca.Muratore, Darwin.Caldwell, Nikos.Tsagarakis}@iit.it

II. RELATED WORK

In the robotic community, LfD techniques are widely used to teach the robots new skills based on human demonstrations. Koenemann et al. [5] introduced a real-time method to allow a humanoid robot to imitate human whole-body motions. Recently, Welschehold et al. [9] proposed to transform human demonstrations to different hand-object trajectories in order to adapt to robotic manipulation tasks. The advantage of LfD methods is their abilities to let the robots accurately repeat human motions, however, it is difficult to expand LfD techniques to a large number of tasks since the training process is usually designed for a specific task or needs training data from real robotic systems [4].

From a computer vision viewpoint, Aksoy et al. [10] introduced a framework that represents the continuous human actions as “semantic event chains” and solved the problem as an activity detection task. In [11], Yang et al. proposed to learn manipulation actions from unconstrained videos using CNN and grammar based parser. However, this method needs an explicit representation of both the objects and grasping types to generate command sentences. Recently, the authors in [12] introduced an unsupervised method to link visual features to textual descriptions in long manipulation tasks. In this paper, we propose to directly learn the output command sentences from the input videos without any prior knowledge. Our method takes advantage of CNN to learn robust features, and RNN to model the sequences, while being easily adapted to any human activity.

Although commands, or in general natural languages, are widely used to control robotic systems. They are usually carefully programmed for each task. This limitation means programming is tedious if there are many tasks. To automatically understand the commands, the authors in [13] formed this problem as a probabilistic graphical model based on the semantic structure of the input command. Similarly, Guadarrama et al. [14] introduced a semantic parser that used both natural commands and visual concepts to let the robot execute the task. While we retain the concepts of [13] and [14], the main difference in our approach is that we directly use the grammar-free commands from the translation module. This allows us to use a simple similarity measure to map each word in the generated command to the real command on the robot.

In deep learning, Donahue et al. [15] made a first attempt to tackle the video captioning problem. The features were first extracted from video frames with CRF then fed to a LSTM network to produce the output captions. In [8], the authors proposed a sequence-to-sequence model to generate captions for videos from both RGB and optical flow images. Yu et al. [16] used a hierarchical RNN to generate one or multiple sentences to describe a video. In this work, we cast the problem of translating videos to commands as a video captioning task to build on the strong state the art in computer vision. Furthermore, we use the output of the deep network as the input command to control a full-size humanoid robot, allowing it to perform different manipulation tasks.

III. TRANSLATING VIDEOS TO COMMANDS

We start by formulating the problem and briefly describing two popular RNNs use in our method: Long-Short Term Memory (LSTM) [17] and Gated Recurrent Neural network (GRU) [18]. Then we present the network architecture that translates the input videos to robotic commands.

A. Problem Formulation

We cast the problem of translating videos to commands as a video captioning task. In particular, the input video is considered as a list of frames, presented by a sequence of features $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ from each frame. The output command is presented as a sequence of word vectors $\mathbf{Y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$, in which each vector \mathbf{y} represents one word in the dictionary D . The video captioning task is to find for each sequence feature \mathbf{X}_i its most probable command \mathbf{Y}_i . In practice, the number of video frames n is usually greater than the number of words m . To make the problem become more suitable for robotic applications, we use a dataset that contains mainly human’s manipulation actions and assume that the output command \mathbf{Y} is in grammar-free format.

B. Recurrent Neural Networks

1) *LSTM*: LSTM is a well-known RNN for effectively modelling the long-term dependencies from the input data. The core of an LSTM network is a memory cell \mathbf{c} which has the gate mechanism to encode the knowledge of the previous inputs at every time step. In particular, the LSTM takes an input \mathbf{x}_t at each time step t , and computes the hidden state \mathbf{h}_t and the memory cell state \mathbf{c}_t as follows:

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}_{xi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{b}_i) \\ \mathbf{f}_t &= \sigma(\mathbf{W}_{xf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1} + \mathbf{b}_f) \\ \mathbf{o}_t &= \sigma(\mathbf{W}_{xo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{b}_o) \\ \mathbf{g}_t &= \phi(\mathbf{W}_{xg}\mathbf{x}_t + \mathbf{W}_{hg}\mathbf{h}_{t-1} + \mathbf{b}_g) \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t \\ \mathbf{h}_t &= \mathbf{o}_t \odot \phi(\mathbf{c}_t) \end{aligned} \quad (1)$$

where \odot represents element-wise multiplication, the function $\sigma : \mathbb{R} \mapsto [0, 1], \sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid non-linearity, and $\phi : \mathbb{R} \mapsto [-1, 1], \phi(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$ is the hyperbolic tangent non-linearity. The weight W and bias b are trained parameters. With this gate mechanism, the LSTM network can remember or forget information for long periods of time, while is still robust against vanishing or exploding gradient problems. In practice, the LSTM network is straightforward to train end-to-end and can handle inputs with different lengths using the padding techniques.

2) *GRU*: A popular variation of the LSTM network is the GRU proposed by Cho et al. [18]. The main advantage of the GRU network is that it requires fewer computations in comparison with the standard LSTM, while the accuracy between these two networks are competitive. Unlike the LSTM network, in a GRU, the update gate controls both the input and forget gates, and the reset gate is applied before the nonlinear transformation as follows:

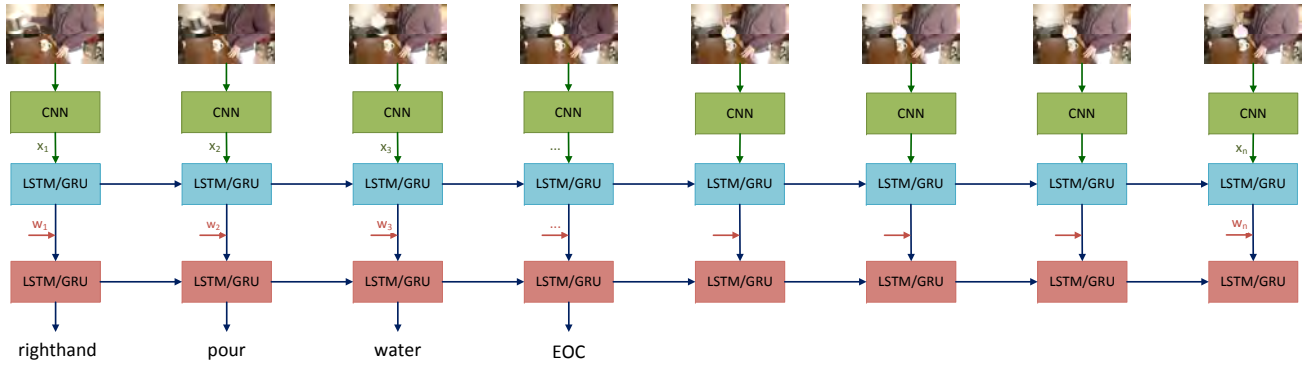


Fig. 2. An overview of our approach. We first extract the deep features from the input frames using CNN. Then the first LSTM/GRU layer is used to encode the visual features. The input words are fed to the second LSTM/GRU layer and this layer sequentially generates the output words.

$$\begin{aligned}
 \mathbf{r}_t &= \sigma(\mathbf{W}_{xr}\mathbf{x}_t + \mathbf{W}_{hr}\mathbf{h}_{t-1} + \mathbf{b}_r) \\
 \mathbf{z}_t &= \sigma(\mathbf{W}_{xz}\mathbf{x}_t + \mathbf{W}_{hz}\mathbf{h}_{t-1} + \mathbf{b}_z) \\
 \tilde{\mathbf{h}}_t &= \phi(\mathbf{W}_{xh}\mathbf{x}_t + \mathbf{W}_{hh}(\mathbf{r}_t \odot \mathbf{h}_{t-1}) + \mathbf{b}_h) \\
 \mathbf{h}_t &= \mathbf{z}_t \odot \mathbf{h}_{t-1} + (1 - \mathbf{z}_t) \odot \tilde{\mathbf{h}}_t
 \end{aligned} \quad (2)$$

where \mathbf{r}_t , \mathbf{z}_t , \mathbf{h}_t represent the reset, update, and hidden gate respectively.

C. Videos to Commands

1) *Command Embedding*: Since a command is a list of words, we have to represent each word as a vector for computation. There are two popular techniques for word representation: *one-hot* encoding and *word2vec* [19] embedding. Although the one-hot vector is high dimensional and sparse since its dimensionality grows linearly with the number of words in the vocabulary, it is straightforward to use this embedding in the video captioning task. In this work, we choose the one-hot encoding technique as our word representation since the number of words in our dictionary is relatively small (i.e., $|D| = 128$). The one-hot vector $\mathbf{y} \in \mathbb{R}^{|D|}$ is a binary vector with only one non-zero entry indicating the index of the current word in the vocabulary. Formally, each value in the one-hot vector \mathbf{y} is defined by:

$$\mathbf{y}^j = \begin{cases} 1, & \text{if } j = \text{ind}(\mathbf{y}) \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where $\text{ind}(\mathbf{y})$ is the index of the current word in the dictionary D . In practice, we add an extra word EOC to the dictionary to denote the end of command sentences.

2) *Visual Features*: We first sample n frames from each input video in order to extract deep features from the images. The frames are selected uniformly with the same interval if the video is too long. In case the video is too short and there are not enough n frames, we create an artificial frame from the mean pixel values of the ImageNet dataset [20] and pad this frame at the end of the list until it reaches n frames. We then use the state-of-the-art CNN to extract deep features from these input frames. Since the visual features provide the key information for the learning process, three popular CNN are used in our experiments: VGG16 [21], Inception_v3 [22], and ResNet50 [23].

Specifically, for the VGG16 network, the features are extracted from its last fully connected `fc2` layer. For the Inception_v3 network, we extract the features from its `pool_3:0` tensor. Finally, we use the features from `pool5` layer of the ResNet50 network. The dimension of the extracted features is 4096, 2048, 2048, for the VGG16, Inception_v3, and ResNet50 network, respectively. All these CNN are pretrained on ImageNet dataset for image classifications. We notice that the names of the layers we mention here are based on the Tensorflow [24] implementation of these networks.

3) *Architecture*: Our architecture is based on the encoder-decoder scheme [8] [25] [26], which is adapted from the popular sequence to sequence model [27] in machine translation. Although recent approaches to video captioning problem use attention mechanism [26] or hierarchical RNN [16], our proposal solely relies on the neural architecture. Based on the input data characteristics, our network smoothly encodes the input visual features and generates the output commands, achieving a fair improvement over the state of the art without using any additional modules.

In particular, given an input video, we first extract visual features from the video frames using the pretrained CNN network. These features are encoded in the first RNN layer to create the encoder hidden state. The input words are then fed to the second RNN layer, and this layer will decode sequentially to generate a list of words as the output command. Fig. 2 shows an overview of our approach. More formally, given an input sequence of features $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, we want to estimate the conditional probability for an output command $\mathbf{Y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$ as follows:

$$P(\mathbf{y}_1, \dots, \mathbf{y}_m | \mathbf{x}_1, \dots, \mathbf{x}_n) = \prod_{i=1}^m P(\mathbf{y}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_1, \mathbf{X}) \quad (4)$$

Since we want a generative model that encodes a sequence of features and produces a sequence of words in order as a command, the LSTM/GRU is well suitable for this task. Another advantage of LSTM/GRU is that they can model the long-term dependencies in the input features and the output words. In practice, we conduct experiments with the LSTM and GRU network as our RNN, while the input visual features are extracted from the VGG16, Inception_v3, and ResNet50 network, respectively.

In the encoding stage, the first LSTM/GRU layer converts the visual features $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ to a list of hidden state vectors $\mathbf{H}^e = (\mathbf{h}_1^e, \mathbf{h}_2^e, \dots, \mathbf{h}_n^e)$ (using Equation 1 for LSTM or Equation 2 for GRU). Unlike [25] which takes the average of all n hidden state vectors to create a fixed-length vector, we directly use each hidden vector \mathbf{h}_i^e as the input \mathbf{x}_i^d for the second decoder layer. This allows the smooth transaction from the visual features to the output commands without worrying about the harsh average pooling operation, which can lead to the loss of temporal structure underlying the input video.

In the decoding stage, the second LSTM/GRU layer converts the list of hidden encoder vectors \mathbf{H}^e into the sequence of hidden decoder vectors \mathbf{H}^d . The final list of predicted words $\hat{\mathbf{Y}}$ is achieved by applying a softmax layer on the output \mathbf{H}^d of the LSTM/GRU decoder layer. In particular, at each time step t , the output \mathbf{z}_t of each LSTM/GRU cell in the decoder layer is passed through a linear prediction layer $\hat{\mathbf{y}} = \mathbf{W}_z \mathbf{z}_t + \mathbf{b}_z$, and the predicted distribution $P(\mathbf{y}_t)$ is computed by taking the softmax of $\hat{\mathbf{y}}_t$ as follows:

$$P(\mathbf{y}_t = \mathbf{w} | \mathbf{z}_t) = \frac{\exp(\hat{\mathbf{y}}_{t,w})}{\sum_{w' \in D} \exp(\hat{\mathbf{y}}_{t,w'})} \quad (5)$$

where \mathbf{W}_z and \mathbf{b}_z are learned parameters, \mathbf{w} is a word in the dictionary D .

In this way, the LSTM/GRU decoder layer sequentially generates a conditional probability distribution for each word of the output command given the encoded features representation and all the previously generated words. In practice, we preprocess the data so that the number of input words m is equal to the number of input frames n . For the input video, this is done by uniformly sampling n frames in the long video, or padding the extra frame if the video is too short. Since the number of words m in the input commands is always smaller than n , we pad a special empty word to the list until we have n words.

4) *Training*: The network is trained end-to-end with Adam optimizer [28] using the following objective function:

$$\arg \max_{\theta} \sum_{i=1}^m \log P(\mathbf{y}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_1; \theta) \quad (6)$$

where θ represents the parameters of the network.

During the training phase, at each time step t , the input feature \mathbf{x}_t is fed to an LSTM/GRU cell in the encoder layer along with the previous hidden state \mathbf{h}_{t-1}^e to produce the current hidden state \mathbf{h}_t^e . After all the input features are exhausted, the word embedding and the hidden states of the first LSTM/GRU encoder layer are fed to the second LSTM/GRU decoder layer. This decoder layer converts the inputs into a sequence of words by maximizing the log-likelihood of the predicted word (Equation 6). This decoding process is performed sequentially for each word until the network generates the end-of-command (EOC) token.

IV. EXPERIMENTS

A. Dataset

Recently, the task of describing video using natural language has gradually received more interest in the computer vision community. Eventually, many video description datasets have been released [29]. However, these datasets only provide general descriptions of the video and there is no detailed understanding of the action. The captions are also written using natural language sentences which can not be used directly in robotic applications. Motivated by these limitations, we introduce a new *video to command* (IIT-V2C) dataset which focuses on *fine-grained* action understanding [30]. Our goal is to create a new large scale dataset that provides fine-grained understanding of human actions in a grammar-free format. This is more suitable for robotic applications and can be used with deep learning methods.

Video annotation Since our main purpose is to develop a framework that can be used by real robots for manipulation tasks, we use only videos that contain human actions. To this end, the raw videos in the Breakfast dataset [31] are best suited to our purpose since they were originally designed for activity recognition. We only reuse the raw videos from this dataset and manually segment each video into short clips in a fine granularity level. Each short clip is then annotated with a *command sentence* that describes the current human action.

Dataset statistics In particular, we reuse 419 videos from the Breakfast dataset. The dataset contains 52 unique participants performing cooking tasks in different kitchens. We segment each video (approximately 2 – 3 minutes long) into around 10 – 50 short clips (approximately 1 – 15 seconds long), resulting in 11,000 unique short videos. Each short video has a single command sentence that describes human actions. We use 70% of the dataset for training and the remaining 30% for testing. Although our new-form dataset is characterized by its grammar-free property for the convenience in robotic applications, it can easily be adapted to classical video captioning task by adding the full natural sentences as the new groundtruth for each video.

B. Evaluation Metric, Baseline, and Implementation

Evaluation Metric We report the experimental results using the standard metrics in the captioning task [29]: BLEU, METEOR, ROUGE-L, and CIDEr. This makes our results directly comparable with the recent state-of-the-art methods in the video captioning field.

Baseline We compare our results with two recent methods in the video captioning field: S2VT [8] and SGC [26]. The authors of S2VT used LSTM in the encoder-decoder architecture, while the inputs are from the features of RGB images (extracted by VGG16) and optical flow images (extracted by AlexNet). SGC also used LSTM with encoder-decoder architecture, however, this work integrated a saliency guided method as the attention mechanism, while the features are from Inception_v3. We use the code provided by the authors of the associated papers for the fair comparison.



Fig. 3. Example of translation results of the S2VT, SGC and our LSTM_Inception_v3 network on the IIT-V2C dataset.

Implementation We use 512 hidden units in both LSTM and GRU in our implementation. The first hidden state of LSTM/GRU is initialized uniformly in $[-0.1, 0.1]$. We set the number of frames for each input video at 30. Sequentially, we consider each command has maximum 30 words. If there are not enough 30 frames/words in the input video/command, we pad the mean frame (from ImageNet dataset)/empty word at the end of the list until it reaches 30. During training, we only accumulate the softmax losses of the real words to the total loss, while the losses from the empty words are ignored. We train all the networks for 150 epochs using Adam optimizer with a learning rate of 0.0001. The batch size is empirically set to 16. The training time for each network is around 3 hours on a NVIDIA Titan X GPU.

C. Results

TABLE I
PERFORMANCE ON IIT-V2C DATASET

	Bleu.1	Bleu.2	Bleu.3	Bleu.4	METEOR	ROUGE.L	CIDEr
S2VT [8]	0.383	0.265	0.201	0.159	0.183	0.382	1.431
SGC [26]	0.370	0.256	0.198	0.161	0.179	0.371	1.422
LSTM_VGG16	0.372	0.255	0.193	0.159	0.180	0.375	1.395
GRU_VGG16	0.350	0.233	0.173	0.137	0.168	0.351	1.255
LSTM_Inception_v3	0.400	0.286	0.221	0.178	0.194	0.402	1.594
GRU_Inception_v3	0.391	0.281	0.222	0.188	0.190	0.398	1.588
LSTM_ResNet50	0.398	0.279	0.215	0.174	0.193	0.398	1.550
GRU_ResNet50	0.398	0.284	0.220	0.183	0.193	0.399	1.567


Table I summarizes the captioning results on the IIT-V2C dataset. Overall, the LSTM network that uses visual features from Inception_v3 (LSTM_Inception_v3) achieves the highest performance, winning on the Blue_1, Blue_2, METEOR, ROUGE.L, and CIDEr metrics. Our LSTM_Inception_v3 also outperforms S2VT and SGC in all metrics by a fair margin. We also notice that both the LSTM_ResNet50 and GRU_ResNet50 networks give competitive results in comparison with the LSTM_Inception_v3 network. Overall, we observe that the architectures that use LSTM give slightly better results than those using GRU. However, this difference is not significant when the ResNet50 features are used to train the models (LSTM_ResNet50 and GRU_ResNet50 results are a tie).

From the experiments, we notice that there are two main factors that affect the results of this problem: the network architecture and the input visual features. Since the IIT-V2C dataset contains mainly the fine-grained human actions in a limited environment (i.e., the kitchen), the SGC architecture that used saliency guide as the attention mechanism does not perform well as in the normal video captioning task. On the other hand, the visual features strongly affect the final results. Our experiments show that the ResNet50 and Inception_v3 features significantly outperform the VGG16 features in both LSTM and GRU networks. Since the visual features are not re-trained in the sequence to sequence model, in practice it is crucial to choose the state-of-the-art CNN as the feature extractor for the best performance.

Fig. 3 shows some examples of the generated commands by our LSTM_Inception_v3, S2VT, and SGC models on the test videos of the IIT-V2C dataset. These qualitative results show that our LSTM_Inception_v3 gives good predictions in many cases, while S2VT and SGC results are more variable. In addition to the good predictions that are identical with the groundtruth, we note that many other generated commands are relevant. Due to the nature of the IIT-V2C dataset, most of the videos are short and contain fine-grained human manipulation actions, while the groundtruth commands are also very short. This makes the problem of translating videos to commands is more challenging than the normal video captioning task since the network has to rely on the minimal information to predict the output command.

D. Robotic Applications

Given the proposed translation module, we build a robotic framework that allows the robot to perform various manipulation tasks by just “watching” the input video. Our goal in this work is similar to [11], however, we propose to keep the video understanding separately from the vision system. In this way, the robot can learn to understand the task and execute it independently. This makes the proposed approach more practical since it does not require a dataset that has both the caption and the object (or grasping) location. It is also important to note that our goals differ from LfD since we only focus on finding a general way to let the robot execute different manipulation actions, while the trajectory in each action is assumed to be known.



(a) Pick and place task



(b) Pouring task

Fig. 4. Example of manipulation tasks performed by WALK-MAN using our proposed framework. (a) Pick and place task. (b) Pouring task. The frames from human instruction videos are on the left side, while the robot performs actions on the right side. We notice that there are two sub-tasks (i.e., two commands) in these tasks: grasping the object and manipulating it. More illustrations can be found in the supplemental video.

In particular, for each task presented by a video, the translation module will generate an output command sentence. Based on this command, the robot uses its vision system to find relevant objects and plan the actions. Experiments are conducted using the humanoid WALK-MAN [32]. The robot is controlled using the XBotCore software architecture [33], while the OpenSoT library [34] is used to plan full-body motion. The relevant objects and their affordances are detected using AffordanceNet framework [35]. For simplicity, we only use objects in the IIT-Aff dataset [36] in the demonstration videos so the robot can recognize them. Using this setup, the robot can successfully perform various manipulation tasks by closing the loop: understanding the human demonstration from the video using the proposed method, finding the relevant objects and grasping poses [36], and planning for each action [34].

Fig. 4 shows some manipulation tasks performed by WALK-MAN using our proposed framework. For a simple task such as “righthand grasp bottle”, the robot can effectively repeat the human action through the command. Since the output of our translation module is in grammar-free format, we can directly map each word in the command sentence to the real robot command. In this way, we avoid using other modules as in [13] to parse the natural command into the one that uses in the real robot. The visual system also plays an important role in our framework since it provides the object location and the target frames (e.g., grasping frame, ending frame) for the robot to plan the actions. Using our approach, the robot can also complete long manipulation tasks by stacking a list of demonstration videos in order for the translation module. Note that, for the long manipulation tasks, we assume that the ending state of one task will be the starting state of the next task. Overall, WALK-MAN successfully performs various manipulation tasks such as grasping, pick and place, or pouring. The experimental video and our IIT-V2C dataset can be found at the following link:

<https://sites.google.com/site/video2command/>

V. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a new method to translate human demonstration videos to commands using deep recurrent neural networks. We conducted experiments with the LSTM and GRU network using different visual feature representations. The experimental results showed that our purely neural sequence to sequence architecture outperformed current state-of-the-art methods by a fair margin. We also introduced a new large-scale videos to commands dataset that is suitable for deep learning methods. Finally, we combined our proposed method with the vision and planning module, and performed various manipulation tests on a real full-size humanoid robot.

Our robotic experiments so far are qualitative. We have focused on demonstrating how our approach can be used in a real robotic system to reduce the (tedious) programming when there are many manipulation tasks. Although using the learning approach to translate the demonstration videos to commands could help the robot understand human actions in a meaningful way, the imitation step is still challenging since it requires a robust vision, planning (and LfD) system. Currently, our framework relies solely on the vision system to plan the actions. This does not allow the robot to perform accurate tasks such as “hammering” or “cutting” which require precise skills. Therefore, an interesting problem is to combine our approach with LfD techniques to improve the robot manipulation capabilities.

ACKNOWLEDGMENT

This work is supported by the European Union Seventh Framework Programme [FP7-ICT-2013-10] under grant agreement no 611832 (WALK-MAN).

REFERENCES

- [1] K. D. Christopher L. Nehaniv, *Imitation and Social Learning in Robots, Humans and Animals Behavioural, Social and Communicative Dimensions*. Cambridge University Press, 2009.

- [2] K. Ramirez-Amaro, M. Beetz, and G. Cheng, "Transferring skills to humanoid robots by extracting semantic representations from observations of human activities," *Artificial Intelligence*, 2015.
- [3] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, 2009.
- [4] B. Akgun, M. Cakmak, J. W. Yoo, and A. L. Thomaz, "Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective," in *International Conference on Human-Robot Interaction (HRI)*, 2012.
- [5] J. Koenemann, F. Burget, and M. Bennewitz, "Real-time imitation of human whole-body motions by humanoids," in *International Conference on Robotics and Automation (ICRA)*, 2014.
- [6] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [7] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Conference on Neural Information Processing Systems (NIPS)*, 2014.
- [8] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko, "Sequence to sequence - Video to text," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] T. Welschhold, C. Dornhege, and W. Burgard, "Learning manipulation actions from human demonstrations," in *International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [10] E. E. Aksoy, A. Orhan, and F. Wörgötter, "Semantic decomposition and recognition of long and complex manipulation action sequences," in *International Journal of Computer Vision*, 2016.
- [11] Y. Yang, Y. Li, C. Fermuller, and Y. Aloimonos, "Robot learning manipulation action plans by "watching" unconstrained videos from the world wide web," in *AAAI Conference on Artificial Intelligence*, 2015.
- [12] E. E. Aksoy, E. Ovchinnikova, A. Orhan, Y. Yang, and T. Asfour, "Unsupervised linking of visual features to textual descriptions in long manipulation activities," *Robotics and Automation Letters*, 2017.
- [13] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. J. Teller, and N. Roy, "Understanding natural language commands for robotic navigation and mobile manipulation." in *AAAI Conference on Artificial Intelligence*, 2011.
- [14] S. Guadarrama, L. Riano, D. Golland, D. Go, Y. Jia, D. Klein, P. Abbeel, T. Darrell, *et al.*, "Grounding spatial relations for human-robot interaction," in *International Conference on Intelligent Robots and Systems (IROS)*, 2013.
- [15] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, "Long-term recurrent convolutional networks for visual recognition and description," in *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [16] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu, "Video paragraph captioning using hierarchical recurrent neural networks," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [17] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computing*, 1997.
- [18] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," *arXiv:1406.1078*, 2014.
- [19] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- [20] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, 2015.
- [21] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [22] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [24] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org.
- [25] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko, "Translating videos to natural language using deep recurrent neural networks," *arXiv preprint arXiv:1412.4729*, 2014.
- [26] V. Ramanishka, A. Das, J. Zhang, and K. Saenko, "Top-down visual saliency guided by captions," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [27] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [28] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference for Learning Representations (ICLR)*, 2014.
- [29] J. Xu, T. Mei, T. Yao, and Y. Rui, "Msr-vtt: A large video description dataset for bridging video and language," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] C. Lea, R. Vidal, and G. D. Hager, "Learning convolutional action primitives for fine-grained action recognition," in *International Conference Robotics and Automation (ICRA)*, 2016.
- [31] H. Kuehne, A. Arslan, and T. Serre, "The language of actions: Recovering the syntax and semantics of goal-directed human activities," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [32] N. G. Tsagarakis, D. G. Caldwell, F. Negrello, W. Choi, L. Baccelliere, V. Loc, J. Noorden, L. Muratore, A. Margan, A. Cardellino, L. Natale, E. Mingo Hoffman, H. Dallali, N. Kashiri, J. Malzahn, J. Lee, P. Kryczka, D. Kanoulas, M. Garabini, M. Catalano, M. Ferrati, V. Varricchio, L. Pallottino, C. Pavan, A. Bicchi, A. Settini, A. Rocchi, and A. Ajoudani, "Walk-man: A high-performance humanoid platform for realistic environments," *Journal of Field Robotics*, 2017.
- [33] L. Muratore, A. Laurenzi, E. Mingo Hoffman, A. Rocchi, D. G. Caldwell, and N. G. Tsagarakis, "Xbotcore: A real-time cross-robot software platform," in *International Conference on Robotic Computing*, 2017.
- [34] A. Rocchi, E. Mingo Hoffman, D. Caldwell, and N. Tsagarakis, "OpenSoT: A Whole-Body Control Library for the Compliant Humanoid Robot COMAN," in *International Conference on Robotics and Automation (ICRA)*, 2015.
- [35] T.-T. Do, A. Nguyen, I. Reid, D. G. Caldwell, and N. G. Tsagarakis, "Affordancenet: An end-to-end deep learning approach for object affordance detection," *arXiv:1709.07326*, 2017.
- [36] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis, "Object-Based Affordances Detection with Convolutional Neural Networks and Dense Conditional Random Fields," in *International Conference on Intelligent Robots and Systems (IROS)*, 2017.