

CERN: Confidence-Energy Recurrent Network for Group Activity Recognition

Tianmin Shu¹, Sinisa Todorovic² and Song-Chun Zhu¹

¹ University of California, Los Angeles ² Oregon State University

tianmin.shu@ucla.edu sinisa@onid.orst.edu sczhu@stat.ucla.edu

Abstract

This work is about recognizing human activities occurring in videos at distinct semantic levels, including individual actions, interactions, and group activities. The recognition is realized using a two-level hierarchy of Long Short-Term Memory (LSTM) networks, forming a feed-forward deep architecture, which can be trained end-to-end. In comparison with existing architectures of LSTMs, we make two key contributions giving the name to our approach as Confidence-Energy Recurrent Network – CERN. First, instead of using the common softmax layer for prediction, we specify a novel energy layer (EL) for estimating the energy of our predictions. Second, rather than finding the common minimum-energy class assignment, which may be numerically unstable under uncertainty, we specify that the EL additionally computes the p -values of the solutions, and in this way estimates the most confident energy minimum. The evaluation on the Collective Activity and Volleyball datasets demonstrates: (i) advantages of our two contributions relative to the common softmax and energy-minimization formulations and (ii) a superior performance relative to the state-of-the-art approaches.

1. Introduction

This paper addresses activity recognition in videos, each showing a group activity or event (e.g., spiking in volleyball) arising as a whole from a number of individual actions (e.g., jumping) and human interactions (e.g., passing the ball). Our goal is to recognize events, interactions, and individual actions, for settings where training examples of all these classes are annotated. When ground truth annotations of interactions are not provided in training data, we only pursue recognition of events and actions.

Recent deep architectures [14, 22], representing a multi-level cascade of Long Short-Term Memory (LSTM) networks [13], have shown great promise in recognizing video events. In these approaches, the LSTMs at the bottom layer are grounded onto individual human trajectories, initially obtained from tracking. These LSTMs are aimed at

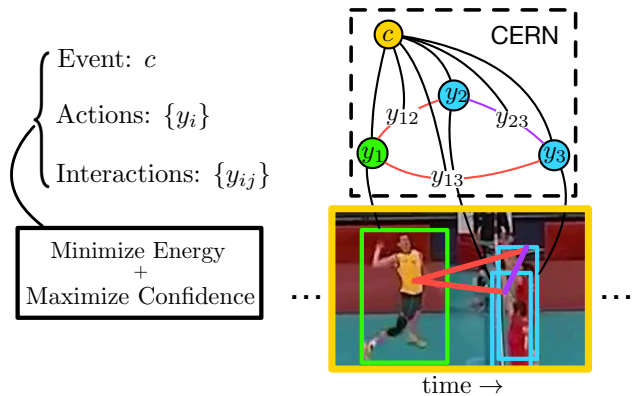


Figure 1: Our CERN represents a two-level hierarchy of LSTMs grounded onto human trajectories, where the LSTMs predict individual actions $\{y_i\}$, human interactions $\{y_{ij}\}$, or the event class c in a given video. CERN outputs an optimal configuration of LSTM predictions which jointly minimizes the energy of the predictions and maximizes their confidence, for addressing the brittleness of cascaded predictions under uncertainty. This is realized by extending the two-level hierarchy with an additional energy layer, which can be trained in an end-to-end fashion.

extracting deep visual representations and predicting individual actions of the respective human trajectories. Outputs of the bottom LSTMs are forwarded to a higher-level LSTM for predicting events. All predictions are made in a feed-forward way using the softmax layer at each LSTM. Such a hierarchy of LSTMs is trained end-to-end using backpropagation-through-time of the cross-entropy loss.

Motivated by the success of these approaches, we start off with a similar two-level hierarchy of LSTMs for recognizing individual actions, interactions, and events. We extend this hierarchy for producing more reliable and accurate predictions in the face of the uncertainty of the visual input.

Ideally, the aforementioned cascade should be learned to overcome uncertainty in a given domain (e.g., occlusion, dynamic background clutter). However, our empirical evaluation suggests that existing benchmark datasets (e.g., the Collective Activity dataset [6] and the Volleyball dataset [14]) are relatively too small for a robust training of

all LSTMs in the cascade. Hence, in cases that have not been seen in the training data, we observe that the feed-forwarding of predictions is typically too brittle, as errors made at the bottom level are directly propagated to the higher level. One way to address this challenge is to augment the training set. But it may not be practical as collecting and annotating group activities is usually difficult.

As shown in Fig. 1, we take another two-pronged strategy toward more robust activity recognition that includes:

1. Minimizing **energy** of all our predictions at the different semantic levels considered, and
2. Maximizing **confidence** (reliability) of the predictions.

Hence the name of our approach – Confidence-Energy Recurrent Network (CERN).

Our first contribution is aimed at mitigating the brittleness of the direct cascading of predictions in previous work. We specify an energy function for capturing dependencies between all LSTM predictions within CERN, and in this way enable recognition by energy minimization. Specifically, we extend the aforementioned two-layer hierarchy of LSTMs with an additional energy layer (EL) for estimating the energy of our predictions. The EL replaces the common softmax layer at the output of LSTMs. Importantly, this extension allows for a robust, energy-based, and end-to-end training of the EL layer on top of all LSTMs in CERN.

Our second contribution is aimed at improving the numerical stability of CERN’s predictions under perturbations in the input, and resolving ambiguous cases with multiple similar-valued local minima. Instead of directly minimizing the energy, we consider more reliable solutions, as illustrated in Fig. 2. The reliability or confidence of solutions is formalized using the classical tool of a statistical hypothesis test [11] – namely, p-values of the corresponding LSTM’s hypotheses (i.e., class predictions). Thus, we seek more confident solutions by regularizing energy minimization with constraints on the p-values. This effectively amounts to a joint maximization of confidence and minimization of energy of CERN outputs. Therefore, we specify the EL to estimate the minimum energy with certain confidence constraints, rather than just the energy.

We also use the energy regularized by p-values for robust deep learning. Specifically, we formulate an energy-based loss which not only accounts for the energy but also the p-values of CERN predictions on the training data.

Our evaluation on the Collective Activity [6] and Volleyball [14] datasets demonstrates: (i) advantages of the above contributions compared with the common softmax and energy-based formulations and (ii) a superior performance relative to the state-of-the-art methods.

In the following, Sec. 2 reviews prior work, Sec. 3 specifies CERN, Sec. 4 and 5 formulate the energy and confidence, Sec. 6 describes the energy layer, Sec. 7 specifies our learning, and finally Sec. 8 presents our results.

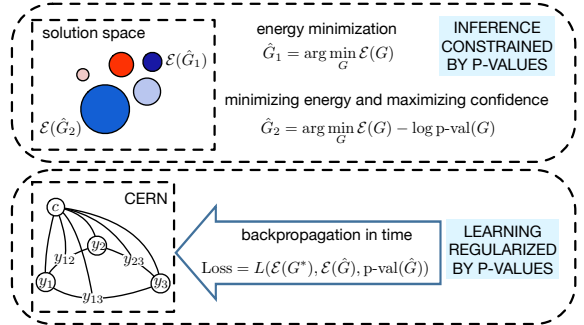


Figure 2: (top) An imaginary illustration of the solution space where each circle represents a candidate solution. The colors and sizes of the circles indicate the energy (red:high, blue:low) and confidence (the larger the radius the higher confidence) computed by the energy layer in CERN. A candidate solution \hat{G}_1 has the minimum energy, but seems numerically unstable for small perturbations in input. A joint maximization of confidence and minimization of energy gives a different, more confident solution \hat{G}_2 . Confidence is specified in terms of p-values of the energy potentials. (bottom) We formulate an energy-based loss for end-to-end learning of CERN. The loss accounts for the energy and p-values.

2. Related Work

Group activity recognition. Group activity recognition often requires the explicit representation of spatiotemporal structures of group activities defined in terms of individual actions and pairwise interactions. Previous work typically used graphical models [18, 17, 23, 1, 5] or AND-OR grammar models [2, 25] to learn the structures grounded on hand-crafted features. Recent methods learn a graphical model, typically MRF [4, 29] or CRF [31, 15, 21], using recurrent neural networks (RNNs). Also, work on group activity recognition [14, 10] has demonstrated many advantages of using deep architectures of RNNs over the mentioned non-deep approaches. Our approach extends this work by replacing the RNN’s softmax layer with a new energy layer, and by specifying a new energy-based model that takes into account p-values of the network’s predictions.

Energy-based learning. While energy-based formulations of inference and learning are common in non-deep group activity recognition [23, 1, 5, 25], they are seldom used for deep architectures. Recently, a few approaches have tried to learn an energy-based model [20, 19] using deep neural networks [3, 30]. They have demonstrated that energy-based objectives have great potential in improving the performance of structured predictions, especially when training data are limited. Our approach extends this work by regularizing the energy-based objective such that it additionally accounts for the confidence of predictions.

Reliability of Recognition. Most energy-based models in computer vision have only focused on the energy minimization for various recognition problems. Our ap-




Figure 3: We specify and evaluate two versions of CERN. CERN is a deep architecture of LSTMs, which are grounded via CNNs to video frames at the bottom. The LSTMs forward their class predictions to the energy layer (EL) at the top. CERN-1 has LSTMs only at the bottom level which compute distributions of individual action classes (colored boxes) or distributions of interaction classes (colored links between green boxes). CERN-2 has an additional LSTM for computing the distribution of event (or group activity) classes. The EL takes the LSTM outputs, and infers an energy minimum with the maximum confidence. The figure shows that CERN-1 and CERN-2 give different results for the group activity *crossing*. CERN-1 wrongly predicts *walking*. CERN-2 typically yields better results for group activities that can not be defined only by individual actions.

proach additionally estimates and regularizes inference with p-values. The p-values are specified within the framework of conformal prediction [24]. This allows the selection of more reliable and numerically stable predictions.

3. Components of the CERN Architecture

For recognizing events, interactions, and individual actions, we use a deep architecture of LSTMs, called CERN, shown in Fig. 3. CERN is similar to the deep networks presented in [14, 15], and can be viewed as a graph $G = \langle V, E, c, Y \rangle$, where $V = \{i\}$ is the set of nodes corresponding to individual human trajectories, and $E = \{(i, j)\}$ is the set of edges corresponding to pairs of human trajectories. These human trajectories are extracted using an off-the-shelf tracker [8]. Also, $c \in \{1, \dots, C\}$ denotes an event class (or group activity), and $Y = Y^V \cup Y^E$ is the union set of individual action classes $Y^V = \{y_i : y_i \in \mathcal{Y}^V\}$ and human interaction classes $Y^E = \{y_{ij} : y_{ij} \in \mathcal{Y}^E\}$ associated with nodes and edges.

In CERN, we assign an LSTM to every node and edge in G . All the node LSTMs share the same weights and all the edge LSTMs also have the same weights. These LSTMs use convolutional neural networks (CNNs) to compute deep features of the corresponding human trajectories, and output softmax distributions of individual action classes, $\psi^V(x_i, y_i)$, or softmax distributions of human in-

teraction classes, $\psi^E(x_{ij}, y_{ij})$. The LSTM outputs are then forwarded to an energy layer (EL) in CERN for computing the energy $\mathcal{E}(G)$. Finally, CERN outputs a structured prediction \hat{G} whose energy has a high confidence:

$$\hat{G} = \arg \min_G \mathcal{E}(G) - \log \text{p-val}(G). \quad (1)$$

As shown in Fig. 3, we specify and evaluate two versions of CERN. CERN-1 uses LSTMs for predicting individual actions and interactions, whereas the event class is predicted by the EL as in (1). CERN-2 has an additional event LSTM which takes features maxpooled from the outputs of the node and edge LSTMs, and then computes the distribution of event classes, $\psi(c)$. The EL in CERN-2 takes all three types of class distributions as input – specifically, $\{\psi^V(x_i, y_i)\}_{i \in V}$, $\{\psi^E(x_{ij}, y_{ij})\}_{(i,j) \in E}$, and $\psi(c)$ – and predicts an optimal class assignment as in (1).

In the following, we specify $\mathcal{E}(G)$ and $\text{p-val}(G)$.

4. Formulation of Energy

For CERN-1, the energy of G is defined as

$$\begin{aligned} \mathcal{E}(G) \propto & \sum_{i \in V} w_{c, y_i}^V \psi^V(x_i, y_i) && \text{node potential} \\ & + \sum_{(i,j) \in E} w_{c, y_{ij}}^E \psi^E(x_{ij}, y_{ij}) && \text{edge potential,} \end{aligned} \quad (2)$$




Figure 4: A simple illustration of the relationship between the non-conformity measure α of individual actions and the p-value, where the ratio of the dashed region to the whole area under the curve indicates the p-value. Clearly, for the given instance, action class 2 has a larger softmax output but action class 1 has a higher confidence. $V_0(c)$ is the training set of videos showing event c .

where w_{c,y_i}^V and $w_{c,y_{ij}}^E$ are parameters, $\psi^V(x_i, y_i)$ denotes the softmax output of the corresponding node LSTM, and $\psi^E(x_{ij}, y_{ij})$ denotes the softmax output of the corresponding edge LSTM (see Sec. 3), and x_i and x_{ij} denote visual cues extracted from respective human trajectories by a CNN as in [10, 14].

For CERN-2, the energy in (2) is augmented by the softmax output of the event LSTM, i.e.,

$$\begin{aligned} \mathcal{E}(G) &\propto \sum_{i \in V} w_{c,y_i}^V \psi^V(x_i, y_i) && \text{node potential} \\ &+ \sum_{(i,j) \in E} w_{c,y_{ij}}^E \psi^E(x_{ij}, y_{ij}) && \text{edge potential} \\ &+ w_c \psi(x, c) && \text{event potential,} \end{aligned} \quad (3)$$

where x in $\psi(x, c)$ is the visual representation of all actions and interactions maxpooled from the outputs of the node LSTMs and edge LSTMs.

5. Formulation of Confidence

There are several well-studied ways to define the p-values [11]. In this paper, we follow the framework of conformal prediction [24]. Conformal prediction uses a non-conformity (dissimilarity) measure to estimate the extent to which a new prediction is different from the system’s predictions made during training. Hence, it provides a formalism to estimate the confidence of new predictions based on the past experience on the training data. Below, we define the nonconformity measure, which is used to compute the p-values for LSTMs’ predictions of individual actions, interactions, and events.

5.1. Nonconformity Measure and P-values

Given the node potential $\psi^V(x_i, y_i)$, we define a non-conformity measure for action predictions:

$$\alpha^V(y_i) = 1 - \frac{\psi^V(x_i, y_i)}{\sum_{y \in \mathcal{Y}^V} \psi^V(x_i, y)} = 1 - \psi^V(x_i, y_i), \quad (4)$$

where the above derivation step holds because $\psi^V(x_i, y_i)$ is the softmax output normalized over action classes. $\alpha^V(y_i)$

is used to estimate the p-value of predicting action class y_i under the context of event class c as

$$p_i^V(c, y_i) = \frac{\sum_{i' \in V_0(c)} \mathbb{1}(y_{i'} = y_i) \mathbb{1}(\alpha^V(y_{i'}) \geq \alpha^V(y_i))}{\sum_{i' \in V_0(c)} \mathbb{1}(y_{i'} = y_i)}. \quad (5)$$

where $\mathbb{1}(\cdot)$ is the indicator, and $V_0(c)$ denotes the set of all human trajectories in training videos with ground truth labels $y_{i'}$ and belonging to the ground truth event class c . From (5), the LSTM prediction $\psi^V(x_i, y_i)$ is reliable – i.e., has a high p-value – when many training examples i' of the same class have larger nonconformity measures.

To better understand the relationship between the non-conformity measure and the p-value, let us consider a simple case illustrated in Fig. 4. The figure plots the two distributions of nonconformity measures of two action classes in the training examples (green: class 1, red: class 2). Suppose that we observe a new instance whose softmax output indicates that action class 2 has a higher probability to be the true label, i.e., $\psi^V(x_i, 1) < \psi^V(x_i, 2)$, and $\alpha^V(1) > \alpha^V(2)$. From the two curves, however, we see that this softmax output is very likely to be wrong. This is because from Fig. 4 we have the p-values $p_i^V(c, 1) > p_i^V(c, 2)$, since a majority of training examples with the class 1 label have larger nonconformity measures than $\alpha^V(1)$, and hence class 1 is a more confident solution.

Similarly, given the softmax output of the edge LSTM, $\psi^E(x_{ij}, y_{ij})$, we specify a nonconformity measure of predicting interaction classes:

$$\alpha_{ij}^E(y_{ij}) = 1 - \frac{\psi^E(x_{ij}, y_{ij})}{\sum_{y \in \mathcal{Y}^E} \psi^E(x_{ij}, y)} = 1 - \psi^E(x_{ij}, y_{ij}), \quad (6)$$

which is then used to estimate the p-value of predicting interaction class y_{ij} under the context of event class c as

$$\begin{aligned} p_{ij}^E(c, y_{ij}) &= \frac{\sum_{(i',j') \in E_0(c)} \mathbb{1}(y_{i'j'} = y_{ij}) \mathbb{1}(\alpha_{i'j'}^E(y_{i'j'}) \geq \alpha_{ij}^E(y_{ij}))}{\sum_{(i',j') \in E_0(c)} \mathbb{1}(y_{i'j'} = y_{ij})}, \end{aligned} \quad (7)$$

where $E_0(c)$ denotes the set of all pairs of human trajectories in training videos with ground truth labels $y_{i'j'}$ and belonging to the ground truth event class c . From (7), the LSTM prediction $\psi^E(x_{ij}, y_{ij})$ has a high p-value when many training examples (i', j') in $E_0(c)$ have larger non-conformity measures.

Finally, in CERN-2, we also have the LSTM softmax output $\psi(x, c)$, which is used to define a nonconformity measure for event predictions:

$$\alpha(c) = 1 - \frac{\psi(x, c)}{\sum_{c \in \mathcal{C}} \psi(x, c)} = 1 - \psi(x, c), \quad (8)$$

and the p-value of predicting event class c as

$$p(c) = \frac{\sum_{v \in V_0} \mathbb{1}(c_v = c) \mathbb{1}(\alpha(c_v) \geq \alpha(c))}{\sum_{v \in V_0} \mathbb{1}(c_v = c)}. \quad (9)$$

where V_0 denotes the set of all training videos.

5.2. Confidence of the Structured Prediction G

To define the statistical significance of the hypothesis G among other hypotheses (i.e., possible solutions), we need to combine the p-values of predictions assigned to nodes, edges and the event of G . More rigorously, for specifying the p-value of a compound statistical test, $p\text{-val}(G)$, consisting of multiple hypotheses, we follow the Fisher's combined hypothesis test [11]. The Fisher's theory states that N independent hypothesis tests, whose p-values are p_1, \dots, p_N , can be characterized by a test statistic χ_{2N}^2 as

$$\chi_{2N}^2 = -2 \sum_{n=1}^N \log p_n, \quad (10)$$

where the statistic χ_{2N}^2 is proved to follow the χ^2 probability distribution with $2N$ degrees of freedom. From (10), it follows that minimization of the statistic χ_{2N}^2 will yield the maximum p-value characterizing the Fisher's combined hypothesis test.

In the following section, we will use this theoretical result to specify the energy layer of our CERN.


6. The Energy Layer of CERN

We extend the deep architecture of LSTMs with an additional energy layer (EL) aimed at jointly minimizing the energy, given by (3), and maximizing a p-value of the Fisher's combined hypothesis test, given by (10). For CERN-2, this optimization problem can be expressed as


$$\begin{aligned} \min_{c, Y} \quad & \mathcal{E}(G) \\ \text{s.t.} \quad & -\sum_{i \in V'} \log p_i^V(c, y_i) \leq \tau^V, \\ & -\sum_{(i,j) \in E'} \log p_{ij}^E(c, y_{ij}) \leq \tau^E, \\ & -\log p(c) < \tau^c, \end{aligned} \quad (11)$$

where τ^V , τ^E , and τ^c are parameters that impose lower-bound constraints on the p-values. Recall that according to the Fisher's theory on a combined hypothesis test, decreasing the constraint parameters τ^V , τ^E , and τ^c will enforce higher p-values of the solution.

From (3) and (11), we derive the following Lagrangian, also referred to as regularized energy $\tilde{\mathcal{E}}(X, Y, c)$, which can



(a) The unit for computing the regularized energy of category c , given by (13).



(b) Diagram of all units in the energy layer.

Figure 5: The EL takes the softmax outputs of all LSTMs along with estimated p-values as input, and outputs a solution that jointly minimizes the energy and maximizes a p-value of the Fisher's combined hypothesis test.

then be readily implemented as the EL:

$$\begin{aligned} \tilde{\mathcal{E}}(X, Y, c) = & \sum_{i \in V} w_{c, y_i}^V \psi^V(x_i, y_i) - \lambda^V \sum_{i \in V} \log p_i^V(c, y_i) \\ & + \sum_{(i,j) \in E} w_{c, y_{ij}}^E \psi^E(x_{ij}, y_{ij}) - \lambda^E \sum_{(i,j) \in E} \log p_{ij}^E(c, y_{ij}) \\ & + w_c \psi(c) - \lambda \log p(c), \end{aligned} \quad (12)$$

Note that for CERN-1, we drop the last two terms in (12), $w_c \psi(c)$ and $\lambda \log p(c)$. $\tilde{\mathcal{E}}(X, Y, c)$ can be expressed in a more compact form as

$$\begin{aligned} \tilde{\mathcal{E}}(X, Y, c) = & \mathbf{w}_c^{V\top} \boldsymbol{\psi}^V - \boldsymbol{\lambda}^{V\top} \log \mathbf{p}_c^V \\ & + \mathbf{w}_c^{E\top} \boldsymbol{\psi}^E - \boldsymbol{\lambda}^{E\top} \log \mathbf{p}_c^E \\ & + w_c \psi_c - \lambda \log p_c, \end{aligned} \quad (13)$$

where all parameters, potentials, and p-values are grouped into corresponding vectors. For brevity, we defer the specification of these vectors to the supplemental material.

Fig. 5a shows a unit in the EL which computes (13). After stacking these units, as shown in Fig. 5b, we select the solution \hat{G} with the minimum $\tilde{\mathcal{E}}(\hat{G})$.

In the following, we explain our energy-based end-to-end training of the EL.

7. Learning Regularized By Confidence

Following [19, 3], we use an energy-based loss for a training instance X^i and its ground truth labels (Y^i, c^i)

to learn parameters of the EL, i.e., the regularized energy, specified in (12):

$$L(X^i, Y^i, c^i) = \max \left(0, \tilde{\mathcal{E}}(X^i, Y^i, c^i) - \tilde{\mathcal{E}}(X^i, \bar{Y}, \bar{c}) + \mathbb{1}(c^i \neq \bar{c}) \right), \quad (14)$$

where $\bar{Y}, \bar{c} = \operatorname{argmin}_{Y, c \neq c^i} \tilde{\mathcal{E}}(X^i, Y, c) - \mathbb{1}(c^i \neq c)$ is the most violated case. Alternatively, this loss can be replaced by other energy-based loss functions also considered in [19]. Here we treat Y as latent variables for simplicity and thus only consider accuracy of c . However, one can include a comparison between Y and its corresponding ground truth label Y^i into the loss function. It is usually difficult to find the most violated case. However, as [20] points out, the inference of the most violated case does not require a global minimum solution since the normalization term is not modeled in our energy-based model, so we can simply set \bar{Y} to be the output of the node and edge LSTMs.

In practice, one can first train a network using common losses such as cross-entropy to learn the representation excluding the EL, namely from the input layer to softmax layers. Then the p-value of a training instance can be computed by removing itself from the training sets V_0 and E_0 . Finally we train the weights in (12) by minimizing the loss.

8. Results

Implementation details. We stack the node LSTMs and edge LSTMs on top of a VGG-16 model [26] without the FC-1000 layer. The VGG-16 is pre-trained on ImageNet [9], and fine-tuned with LSTMs jointly. We train the top layer of CERN by fixing the weights of the CNNs and the bottom layer LSTMs. The batch size for the joint training of the bottom LSTMs and VGG-16 is 6. The training converges within 20000 iterations. The event LSTM and the EL are trained using 10000 iterations with a batch size of 2000. For the mini-batch gradient descent, we use RM-Sprop [28] with a learning rate ranging from 0.000001 to 0.001. We use Keras [7] with Theano [27] as the backend to implement CERN, and run training and testing with a single NVIDIA Titan X (Pascal) GPU. For a fair comparison with [14], we use the same tracker and its implementation as in [14]. Specifically, we use the tracker of [8] from the Dlib library [16]. The cropped image sequences of persons and pairs of persons are used as the inputs to node LSTMs and edge LSTMs, respectively.

We compare our approach with the state-of-the-art methods [12, 14]. In addition, we evaluate the following reasonable baselines.

Baselines:

- 2-layer LSTMs (B1). We test a network of 2-layer LSTMs similar to [14]. All other baselines below and our full models use B1 to compute their potentials and p-values. B1 does not have the energy layer, but only a

feed-forward network. The event class is predicted by the softmax output of the event LSTM.

- CERN-1 w/o p-values (B2). This baseline represents the CERN-1 network with the EL, however, the p-values are not computed and not used for regularizing energy minimization. Hence, the event class prediction of B2 comes from the standard energy minimization.
- CERN-2 w/o p-values (B3). Similar to B2, in this B3, we do not estimate and do not use the p-values in the EL of CERN-2.

Datasets. We evaluate our method in two domains: collective activities and sport events using the Collective Activity dataset [6] and the Volleyball dataset [14] respectively.

8.1. Collective Activity Dataset

The Collective Activity dataset consists of 44 videos, annotated with 5 activity categories (*crossing, walking, waiting, talking, and queuing*), 6 individual action labels (*NA, crossing, walking, waiting, talking, and queuing*), and 8 pairwise interaction labels (*NA, approaching, leaving, passing-by, facing-each-other, walking-side-by-side, standing-in-a-row, standing-side-by-side*). The interaction labels are provided by the extended annotation in [5].

For this dataset, we first train the node LSTMs and edge LSTMs with 10 time steps and 3000 nodes. Then, we concatenate the outputs of these two types of LSTMs at the bottom layer of CERN, along with their VGG-16 features, and pass the concatenation to the bidirectional event LSTM with 500 nodes and 10 time steps at the top layer of CERN. The concatenation is passed through a max pooling layer and a fully-connected layer with a output dimension of 4500.

For comparison with [12, 14] and baselines B1-B3, we use the following performance metrics: (i) multi-class classification accuracy (MCA), and (ii) mean per-class accuracy (MPCA). Our split of training and testing sets is the same as in [12, 14]. Tab. 1 summarizes the performance of all methods on recognizing group activities. Note that in Tab. 1 only [12] does not use deep neural nets. As can be seen, our energy layer significantly boosts the accuracy, outperforming the state-of-the-art by a large margin. Even when we only have the bottom layer of LSTMs, CERN-1 still outperforms the 2-layer LSTMs in [14] thanks to the EL. Without the EL, the baseline B1 yields lower accuracy than [14] even with additional LSTMs for the interactions.

Our accuracies of recognizing individual actions and interactions on the Collective Activity dataset are 72.7% and 59.9%, using the node LSTMs and edge LSTMs respectively. Note that B1, CERN-1 and CERN-2 share the same node and edge LSTMs.

For evaluating numerical stability of predicting group activity classes by CERN-2, we corrupt all human trajectories in the testing data, and control the amount of corruption with the corruption probability. For instance, for the cor-

Method	MCA	MPCA
Cardinality kernel [12]	83.4	81.9
2-layer LSTMs [14]	81.5	80.9
B1: 2-layer LSTMs	79.7	80.3
B2: CERN-1 w/o p-values	83.8	84.3
B3: CERN-2 w/o p-values	83.8	83.7
CERN-1	84.8	85.5
CERN-2	87.2	88.3

Table 1: Comparison of different methods for group activity recognition on the Collective Activity dataset.




Figure 6: Performance decrease of group activity recognition for a varying percentage of corruption of human trajectories in the Collective Activity dataset. We compare 2-layer LSTMs (B1), CERN-2 w/o p-values (B3) and CERN-2 using the same corrupted trajectories as input.

ruption probability of 0.5, we corrupt one bounding box of a person in every video frame with a 0.5 chance. When the bounding box is selected, we randomly shift it with a horizontal and a vertical displacement ranging from 20% to 80% of the original bounding box’s width and height respectively. As Fig. 6 shows, CERN-2 consistently experiences a lower degradation in performance compared to the baselines without p-values. This indicates that incorporating the p-values into the energy model indeed benefits the inference stability. Such benefit becomes more significant as the amount of corruption in input data increases.

Fig. 8 shows an example of the *crossing* activity. As can be seen, although B1 and CERN-2 share the same individual action labels, where a majority of the people are assigned incorrect action labels, CERN-2 can still correctly recognize the activity.

8.2. Volleyball Dataset

The Volleyball dataset consists of 55 videos with 4830 annotated frames. The actions labels are *waiting*, *setting*, *digging*, *failing*, *spiking*, *blocking*, *jumping*, *moving*, and *standing*; and the group activity classes include *right set*, *right spike*, *right pass*, *right winpoint*, *left winpoint*, *left pass*, *left spike*, and *left set*. Interactions are not annotated in this dataset, so we do not recognize interactions and remove the edge LSTMs.

The node LSTMs have 3000 nodes and 10 time steps

Method	MCA	MPCA
2-layer LSTMs [14] (1 group)	70.3	65.9
B1: 2-layer LSTMs (1 group)	71.3	69.5
B2: CERN-1 w/o p-values (1 group)	33.3	34.3
B3: CERN-2 w/o p-values (1 group)	71.7	69.8
CERN-1 (1 group)	34.4	34.9
CERN-2 (1 group)	73.5	72.2
2-layer LSTMs [14] (2 groups)	81.9	82.9
B1: 2-layer LSTMs (2 group)	80.3	80.5
B3: CERN-2 w/o p-values (2 groups)	82.2	82.3
CERN-2 (2 groups)	83.3	83.6

Table 2: Comparison of different methods for group activity recognition on the Volleyball dataset. The first block is for the methods with 1 group and the second one is for those with 2 groups.




Figure 7: The decrease of group activity recognition accuracy over different input distortion percentages on the Volleyball dataset (all use the 2 groups style). CERN-2 is compared with 2-layer LSTMs (B1) and CERN-2 w/o p-values (B3).

(including 5 preceding and 4 succeeding frames). The event LSTM in CERN-2 is a bidirectional LSTM with 1000 nodes and 10 time steps. In [14], the max pooling has two types: 1) pooling over the output of all node LSTMs, or 2) dividing the players into two groups (the left team and the right team) first and pooling over each group separately. We test both types of max pooling for our approach to rule out the effect of pooling type in the comparison. CERN-1 does not have the pooling layer, thus is categorized as 1 group style.

Recognition accuracy of individual actions is 69.1% using node LSTMs, and the accuracies of recognizing group activities are summarized in Tab. 2. Clearly, the regularized energy minimization increases the accuracy compared to the conventional energy minimization (B2 and B3), and CERN-2 outperforms the state-of-the-art when using either of the pooling types. CERN-1 does not achieve accuracy that is comparable to that of CERN-2 on the Volleyball dataset. This is mainly because CERN-1 reasons the group activity based on individual actions, which may not provide sufficient information for recognizing complex group activities in sports videos. CERN-2 overcomes this problem by adding the event LSTM.

We also evaluate the stability of recognizing group activities by CERN-2 under corruption of input human trajec-




Figure 8: The qualitative results on the Collective Activity dataset. From left to right, we show the inference results from B1, CERN-2 and the ground truth (GT) labels respectively. The colors of the bounding boxes indicate the individual action labels (green: *crossing*, red: *waiting*, magenta: *walking*). The interaction labels are not shown here for simplicity.



Figure 9: The qualitative results on the Volleyball dataset: results of B1 (top), results of CERN-2 (middle) and the ground truth (GT) labels (bottom). The colors of the bounding boxes indicate the individual action labels (green: *waiting*, yellow: *digging*, red: *falling*, magenta: *standing*), and the numbers are the frame IDs.

tories. As Fig. 7 indicates, the p-values in the EL indeed increase the inference reliability on the Volleyball dataset.

The qualitative results (2 groups) of a *right pass* activity is depicted in Fig. 9, which demonstrates the advantage of the inference based on our regularized energy compared to the softmax output of the deep recurrent networks when the action predictions are not accurate.

9. Conclusion

We have addressed the problem of recognizing group activities, human interactions, and individual actions with a novel deep architecture, called Confidence-Energy Recurrent Network (CERN). CERN extends an existing two-level hierarchy of LSTMs by additionally incorporating a confidence measure and an energy-based model toward improving reliability and numerical stability of inference. Inference is formulated as a joint minimization of the energy and maximization of the confidence measure of predictions made by the LSTMs. This is realized through a new differentiable energy layer (EL) that computes the energy reg-

ularized by a p-value of the Fisher’s combined statistical test. We have defined an energy-based loss in terms of the regularized energy for learning the EL end-to-end. CERN has been evaluated on the Collective Activity dataset and Volleyball dataset. In comparison with previous approaches that predict group activities in a feed-forward manner using deep recurrent networks, CERN gives a superior performance, and also gives more numerically stable solutions under uncertainty. For collective activities, our simpler variant CERN-1 gives more accurate predictions than a strong baseline representing a two-level hierarchy of LSTMs with softmax outputs taken as predictions. Our variant CERN-2 increases complexity but yields better accuracy on challenging group activities which are not merely a sum of individual actions but a complex whole.

Acknowledgements

This research was supported by grants DARPA MSEE project FA 8650-11-1-7149, ONR MURI project N00014-16-1-2007, and NSF IIS-1423305.

References

- [1] M. R. Amer, P. Lei, and S. Todorovic. Hirf: Hierarchical random field for collective activity recognition in videos. In *European Conference on Computer Vision (ECCV)*, pages 572–585, 2014. 2
- [2] M. R. Amer, D. Xie, M. Zhao, S. Todorovic, and S.-C. Zhu. Cost-sensitive top-down/bottom-up inference for multiscale activity recognition. In *European Conference on Computer Vision (ECCV)*, pages 187–200, 2012. 2
- [3] D. Belanger and A. McCallum. Structured prediction energy networks. In *International Conference on Machine Learning (ICML)*, 2016. 2, 5
- [4] L.-C. Chen, A. G. Schwing, A. L. Yuille, and R. Urtasun. Learning deep structured models. In *International Conference on Machine Learning (ICML)*, 2015. 2
- [5] W. Choi and S. Savarese. Understanding collective activities of people from videos. *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, 36(6):1242–1257, 2014. 2, 6
- [6] W. Choi, K. Shahid, and S. Savarese. What are they doing? : Collective activity classification using spatio-temporal relationship among people. In *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1282–1289, 2009. 1, 2, 6
- [7] F. Chollet. Keras. <https://github.com/fchollet/keras>, 2015. 6
- [8] M. Danelljan, G. Häger, F. Khan, , and M. Felsberg. Accurate scale estimation for robust visual tracking. In *British Machine Vision Conference (BMVC)*, 2014. 3, 6
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 6
- [10] Z. Deng, A. Vahdat, H. Hu, and G. Mori. Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4772–4781, 2016. 2, 4
- [11] R. A. Fisher. *Statistical Methods for Research Workers*. London: Oliver and Boyd, 11 edition, 1950. 2, 4, 5
- [12] H. Hajimirsadeghi, W. Yan, A. Vahdat, and G. Mori. Visual recognition by counting instances: A multi-instance cardinality potential kernel. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2596–2605, 2015. 6, 7
- [13] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 1
- [14] M. S. Ibrahim, S. Muralidharan, Z. Deng, A. Vahdat, and G. Mori. Hierarchical deep temporal models for group activity recognition. *arXiv preprint*, arXiv:1607.02643, 2016. 1, 2, 3, 4, 6, 7
- [15] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5308–5317, 2016. 2, 3
- [16] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009. 6
- [17] T. Lan, L. Sigal, and G. Mori. Social roles in hierarchical models for human activity recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4321–4328, 2012. 2
- [18] T. Lan, Y. Wang, W. Yang, S. Robinovitch, and G. Mori. Discriminative latent models for recognizing contextual group activities. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 34(8):1549–1562, 2012. 2
- [19] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, , and F. J. Huang. A tutorial on energy-based learning. In G. Bakir, T. Hofman, B. Scholkopf, A. Smola, and B. Taskar, editors, *Predicting Structured Data*, pages 191–246. MIT Press, 2006. 2, 5, 6
- [20] Y. LeCun and F. J. Huang. Loss functions for discriminative training of energy-based models. In *Artificial Intelligence and Statistics Conference (AISTATS)*, 2005. 2, 6
- [21] X. Liang, X. Shen, J. Feng, L. Lin, and S. Yan. Semantic object parsing with graph lstm. In *European Conference on Computer Vision (ECCV)*, pages 125–143, 2016. 2
- [22] V. Ramanathan, J. Huang, S. Abu-El-Haija, A. Gorban, K. Murphy, and L. Fei-Fei. Detecting events and key actors in multi-person videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3043–3053, 2016. 1
- [23] V. Ramanathan, B. Yao, and L. Fei-Fei. Social role discovery in human events. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2475–2482, 2013. 2
- [24] G. Shafer and V. Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9:371–421, 2008. 3, 4
- [25] T. Shu, D. Xie, B. Rothrock, S. Todorovic, and S.-C. Zhu. Joint inference of groups, events and human roles in aerial videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4576–4584, 2015. 2
- [26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint*, arXiv:1409.1556, 2014. 6
- [27] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv preprint*, abs:1605.02688, May 2016. 6
- [28] T. Tieleman and G. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. Coursera: Neural networks for machine learning. 6
- [29] Z. Wu, D. Lin, and X. Tang. Deep markov random field for image modeling. In *European Conference on Computer Vision (ECCV)*, pages 295–312, 2016. 2
- [30] J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial network. *arXiv preprint*, arXiv:1609.03126, 2016. 2
- [31] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. Torr. Conditional random fields as recurrent neural networks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1529–1537, 2015. 2