

Open Set Action Recognition via Multi-Label Evidential Learning

Chen Zhao, Dawei Du, Anthony Hoogs, Christopher Funk
 Kitware

{chen.zhao, dawei.du, anthony.hoogs, christopher.funk}@kitware.com

Abstract

Existing methods for open set action recognition focus on novelty detection that assumes video clips show a single action, which is unrealistic in the real world. We propose a new method for open set action recognition and novelty detection via **M**ulti-**L**abel **E**vidential learning (MULE), that goes beyond previous novel action detection methods by addressing the more general problems of single or multiple actors in the same scene, with simultaneous action(s) by any actor. Our Beta Evidential Neural Network estimates multi-action uncertainty with Beta densities based on actor-context-object relation representations. An evidence debiasing constraint is added to the objective function for optimization to reduce the static bias of video representations, which can incorrectly correlate predictions and static cues. We develop a primal-dual average scheme update-based learning algorithm to optimize the proposed problem and provide corresponding theoretical analysis. Besides, uncertainty and belief-based novelty estimation mechanisms are formulated to detect novel actions. Extensive experiments on two real-world video datasets show that our proposed approach achieves promising performance in single/multi-actor, single/multi-action settings. Our code and models are released at <https://github.com/charliezhaoyinpeng/mule>.

1. Introduction

Open set human action recognition has been studied in recent years due to its great potential in real-world applications, such as security surveillance [1], autonomous driving [34], and face recognition [26]. It differs from closed set problems that aim to classify human actions into a predefined set of known classes, since open set methods can identify samples with unseen classes with high accuracy [14].

To this end, several recent methods [4, 6, 10] are proposed for open set human action recognition. As shown in the bottom-left of Figure 1, they focus on single-actor, single-action based recognition, assuming that each video contains only one single action. Compared with softmax

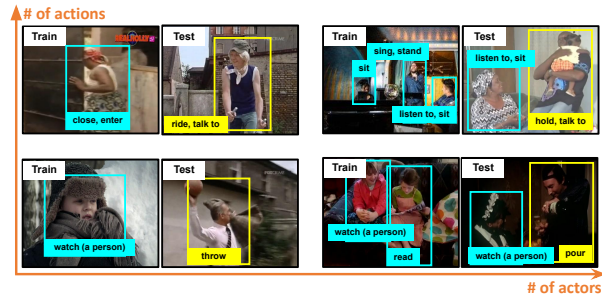


Figure 1. Novelty detection examples of single/multiple actor(s) with single/multiple action(s) in video [16, 38], where an actor is identified as novel (yellow) rather than being from a known category (cyan) in inference. Existing works [4, 6] on open set action recognition focus on single actor associated with single action (bottom-left), while our method can handle different situations.

thresholding [13, 22, 25] for closed set recognition, evidential neural networks (ENNs) [4, 36] can provide a principled way to jointly formulate the multi-class classification and uncertainty modeling to measure novelty of an instance more accurately. It assumes that class probability follows a prior Dirichlet distribution. However, in more realistic situation with multiple actions of actor(s) (see the upper part of Figure 1), the Dirichlet distribution does not hold because the predicted likelihood of each action follows a binomial distribution (*i.e.*, identifying either known or novel action).

In this paper, we introduce a general but understudied problem, namely *novelty detection of actor(s) with multiple actions*. Given real-world use cases [14, 39], the goal is to accurately detect if actor(s) perform novel/unknown action(s) or not. Following [43], an actor is considered unknown if it does not contain any known action(s). Inspired by the belief theory [17, 45], we propose a new framework named **M**ulti-**L**abel **E**vidential learning (MULE), which is composed of three modules: Actor-Context-Object Relation modeling (ACO-R), Beta Evidential Neural Network (Beta-ENN), and Multi-label Evidence Debiasing Constraint (M-EDC). First, we build ACO-R representation to exploit the actors' interactions with the surrounding objects and the context. Then, we use Beta-ENN to estimate the evidence of known actions, and quantify the predictive uncertainty of actions so that unknown actions would incur

high uncertainty, *i.e.*, lack of confidence for known predictions. Here, the evidence indicates actions closest to the predicted one in the feature space and are used to support the decision-making [36]. Instead of relying on Dirichlet distribution [4], the evidence in Beta-ENN is regarded as parameters of a Beta distribution which is a conjugate prior of the Binomial likelihood.

Additionally, in open set recognition, static bias [21] may bring a false correlation between the prediction and static cues, such as scenes, resulting in inferior generalization capability of a model. Therefore, the M-EDC is added to the objective function of our framework to reduce the static bias for video actions. We propose a duality-based learning algorithm to optimize the network. Specifically, we apply an averaging scheme to proximate primal optimal solutions. The primal and dual parameters are updated interactively, where the primal parameters regard model accuracy and dual parameters adjust model debiasing. The theoretical analysis shows the convergence of the primal solution sequence and gives bounds for both the loss function and the violation of the debiasing constraint in MULE. According to the proposed uncertainty and belief based novelty estimation mechanisms, our model outperforms the state-of-the-art on two action recognition datasets (*i.e.*, AVA [16] and Charades [38]) in terms of novelty detection. The main contributions of this work are summarized:

- A new framework MULE is proposed for open set action recognition in videos that contains either a single or multiple actors associated with one or more actions. To the best of our knowledge, this is the first study to detect actors with multiple unknown actions.
- To optimize the Beta-ENN, we develop a primal-dual average scheme update algorithm, with theoretical guarantees on the convergence of the primal solution sequence and bounds for both the loss function and the violation of the debiasing constraint.
- We introduce four novelty estimation mechanisms to calculate novelty score and achieve better performance on novel action detection compared with existing methods.

2. Related Work

Open set action recognition. Despite a great many explorations on video action recognition [11, 12, 47], most existing methods are developed under the assumption that all actions are known a priori (closed set), and limited attention is given to the open set problems. OpenMax [6] is initially proposed for open set recognition, in which it leverages *Extreme Value Theory* [35] to expand the K -class softmax classifier. Roitberg *et al.* [33] develops the voting-based scheme to leverage the estimated uncertainty of action predictions to measure the novelty of the test sample. DEAR [4] formulates the open set action recognition prob-

lem by estimating the uncertainty of single labeled actions to distinguish between the known and unknown samples. However, such methods focus on simple scenarios where each actor has only one action in a video.

Uncertainty estimation. To distinguish unknown samples from known ones, how deep networks identify samples belonging to an unrelated data distribution becomes crucial. To this end, a stream of research on Bayesian Neural Nets (BNNs) [20] is proposed to estimate prediction uncertainty by approximating the moments of the posterior predictive distribution. However, BNNs face several limitations, including the intractability of directly inferring the posterior distribution of the weights given data, the computational expense of sampling during inference, and the question of how to choose a weight prior [2]. Recently, Evidential Neural Networks (ENNs) [36] estimate evidential uncertainty for multi-class classification problems. But they are designed for single-label multi-classification issues by assuming that class probability follows a prior Dirichlet distribution.

Debiasing. Static bias is another challenging issue limiting the generalization capability of a model in an open-set setting [4]. The manifestation of static bias can often be as fraught as the spurious correlation between the prediction and sensitive features like unrelated objects and background [4, 7, 21]. It empirically shows that debiasing the model by input data or learned representation can improve the accuracy. RESOUND [21] indicates that static bias may help to achieve better results in a closed-set setting if an action can overfit it. Choi *et al.* [7] introduce scene-adversarial and human mask confusion losses to mitigate scene bias. For an open-set setting, DEAR [4] introduces the contrastive evidence debiasing by temporally shuffled feature input and 2D convolution. However, it still pushes the sensitive feature to be independent of the non-sensitive one only, in which the dependency of the sensitive feature on model predictions is ignored. In contrast, we propose a new Beta distribution based network with multi-label evidential learning, where static bias is reduced in terms of both indirect and direct dependencies.

3. Method

Given an input video, we aim to determine the predictive uncertainty of the detected actors with unknown actions. For inference, actions with high and low uncertainty can be regarded as unknown and known, respectively. As illustrated in Figure 2, our method extracts Actor-Context-Object Relation (ACO-R) to train the Beta Evidential Neural Network (Beta-ENN) with the Multi-label Evidence Debiasing Constraint (M-EDC).

3.1. Actor-Context-Object Relation Modeling

Previous works usually exploit pairwise relations between actors and objects [31] or actor parts [37]. In this

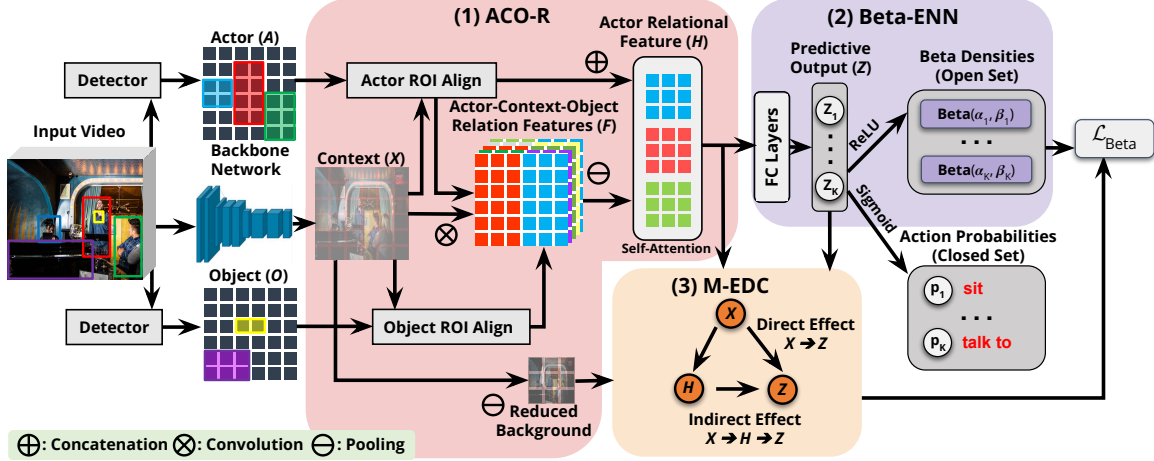


Figure 2. Our MULE framework is composed of three modules. (1) Actor relational features are extracted through the ACO-R module, where they encode information based on interactions between actor/object instances with the context. (2) Positive and negative evidence are estimated through Beta-ENN to quantify the predictive uncertainty of various human actions. (3) M-EDC is a debiasing constraint added to the loss function \mathcal{L}_{beta} in Equation (5), which aims to mitigate static bias.

work, similar to [29, 40], we extract higher-order relations to predict actions in complex scenes. Specifically, each video is split into overlapping clips with the same frame length, where the center frame is regarded as a keyframe. N actors and M objects are detected for each input video clip. To reduce complexity, we duplicate the detections in the keyframe to the remaining frames of the clip as naive temporal linking. Then, actor features $\{A^i\}_{i=1}^N \in \mathbb{R}^C$ and object features $\{O^i\}_{i=1}^M \in \mathbb{R}^C$ are produced by the off-the-shelf detector [32], each of which describes the spatio-temporal appearance and motion of one Region of Interest (RoI). We can also obtain the context feature $X \in \mathbb{R}^{C \times H \times W}$ based on the action backbone network, where C, H, W represent channel, height, and width.

To model actor-context-object relations, we first concatenate all actor features to the context feature followed by object features to form a series of concatenated feature maps $\{F^i\}_{i=1}^{N \times M} \in \mathbb{R}^{3C \times H \times W}$. Following a stack of transformer blocks [29], we can obtain actor relational features $\{H^i\}_{i=1}^{N \times M}$ by calculating the higher-order relations between pairs of $\{F^i\}_{i=1}^{N \times M}$ at the same spatial location, where two actors can be associated via the same spatial context and different objects.

3.2. The Beta Evidential Neural Network

As discussed in Section 1, existing models [13, 22, 25] typically rely on a softmax layer to perform multi-class classification. Since the softmax score is essentially a point estimation of a predictive distribution [4, 13, 36], the models cannot estimate the predictive uncertainty of out-of-distribution. To tackle this limitation, evidential neural networks (ENNs) [36] are developed to jointly formulate the multi-class classification and uncertainty modeling. Specifically, ENNs [2, 4] 1) interpret the standard output of a clas-

sification network as the parameter set of a categorical distribution and 2) assume that class probability follows a prior Dirichlet distribution and replaces this parameter set with the parameters of a Dirichlet density for novelty detection.

However, in our setting, the assumption of using Dirichlet distribution is not a good fit for an actor with multiple actions. It is because the predicted likelihood for each action class follows a binomial distribution whose conjugate prior is a Beta distribution rather than Dirichlet distribution. To this end, we design the Beta-ENN to classify known and novel actions based on Beta distributions.

Beta distribution-based subjective opinions. According to the belief theory [45], it is more reasonable to predict subjective opinions rather than class probabilities in an open-set setting. We hence use the principles of evidential theory to quantify belief masses and uncertainty in the proposed Beta-ENN through Subjective Logic (SL) [17].

In multi-label action recognition, since each action follows a given binomial opinion towards the proposition, the subjective opinion $\omega_i = (b_i, d_i, u_i, a_i)$ for an action $i \in \{1, \dots, K\}$ of an actor is expressed by two belief masses, *i.e.*, belief $b_i \in [0, 1]$ and disbelief $d_i \in [0, 1]$, and one uncertainty mass $u_i \in [0, 1]$, where $b_i + d_i + u_i = 1$. The expected belief probability p_i is defined as $p_i = b_i + a_i \cdot u_i$, where a_i refers to a base rate representing prior knowledge without commitment, such as neither agree nor disagree.

According to [36], a binomial opinion ω_i of action follows a Beta probability density function (pdf) denoted $\text{Beta}(p_i | \alpha_i, \beta_i)$, where $p_i \in [0, 1]$ represents the action assignment probabilities. The Beta pdf is characterized by parameters α_i and β_i , where α_i and β_i are viewed as positive and negative evidence of the observed action i , respectively. The evidence indicates actions closest to the predicted ones in the feature space and is used to support the decision-

making. Action labels are the same for positive evidence, but different for negative evidence. Specifically, for each action, the opinion ω_i is obtained based on its corresponding α_i and β_i using the rule in SL:

$$b_i = \frac{\alpha_i - a_i w}{\alpha_i + \beta_i}, \quad d_i = \frac{\beta_i - a_i w}{\alpha_i + \beta_i}, \quad u_i = \frac{w}{\alpha_i + \beta_i}, \quad (1)$$

where we set the non-informative prior weight $w = 2$ and base rate $a_i = 1$ for each binary action classification (known/unknown) empirically. Therefore, a collection of evidence pairs $\{(\alpha_i, \beta_i)\}_{i=1}^K$ of an actor is estimated to quantify the predictive uncertainty. An actor with multiple unknown actions would incur high uncertainty u with low belief b (see Section 4.2).

Learning opinions through Beta loss. As stated in [36], predictors of an evidence-based learner are presented as a distribution over possible softmax outputs. Therefore, we use fully connected and rectified activation layers (e.g., ELU [8]) as the evidence function $s(\cdot)$ to replace the last softmax layer and keep positive and negative evidences $\alpha_j, \beta_j \succcurlyeq 1$. For an actor j , we have $\alpha_j = s(h(\mathbf{x}_j; \boldsymbol{\theta})) + 1$ and $\beta_j = s(h(\mathbf{x}_j; \boldsymbol{\theta})) + 1$, where $\alpha_j = [\alpha_{1j}, \dots, \alpha_{Kj}]^T$ and $\beta_j = [\beta_{1j}, \dots, \beta_{Kj}]^T$. \mathbf{x}_j denotes the input video, $h(\mathbf{x}_j; \boldsymbol{\theta})$ represents the evidence vector predicted by the network for the classification and, $\boldsymbol{\theta}$ represents parameters for ACO-R modeling.

To learn the above opinions, we define the Beta loss function by computing its Bayes risk for the action predictor. For the binary cross-entropy loss for each action i over a batch of actors, the proposed Beta loss takes

$$\mathcal{L}_{Beta}(\boldsymbol{\theta}) = \sum_{j=1}^N \mathcal{L}_j(\boldsymbol{\theta}), \quad (2)$$

where $j \in \{1, \dots, N\}$ denotes the index of an actor.

$$\begin{aligned} \mathcal{L}_j(\boldsymbol{\theta}) &= \sum_{i=1}^K \int \mathbf{BCE}(y_{ij}, p_{ij}) \mathbf{Beta}(p_{ij}; \alpha_{ij}, \beta_{ij}) dp_{ij} \\ &= \sum_{i=1}^K \left[y_{ij} \left(\psi(\alpha_{ij} + \beta_{ij}) - \psi(\alpha_{ij}) \right) \right. \\ &\quad \left. + (1 - y_{ij}) \left(\psi(\alpha_{ij} + \beta_{ij}) - \psi(\beta_{ij}) \right) \right], \quad (3) \end{aligned}$$

where K is the number of actions, and $\mathbf{BCE}(\cdot)$ is the binary cross-entropy loss, and $\psi(\cdot)$ is the *digamma* function. The log expectation of Beta distribution derives the last equality. $\mathbf{y}_j = [y_{1j}, \dots, y_{Kj}] \in \{0, 1\}^K$ is the K -dimensional ground-truth action(s) label for \mathbf{x}_j . Thus positive and negative evidences α, β can be optimized via the above equations during training.

3.3. Multi-Label Evidence Debiasing Constraint

For open set action recognition, static bias [21] results in a vulnerable model that falsely recognizes an action containing similar static features. For example, the action of “walking” is easily recognized with “road” in the background, but it would be unable to recognize the same ac-

tion with the “treadmill” scene. From the perspective of fairness-aware learning [48–50], as indicated in Figure 2, this is due to the spurious dependency of the prediction Z (e.g., actions) onto sensitive features X (e.g., background scene), and strong dependency indicates strong effects.

To mitigate static bias, we introduce the evidence debiasing constraint in multi-label evidential learning. A fair prediction indicates no direct ($X \rightarrow Z$) or indirect ($X \rightarrow H \rightarrow Z$) dependency of Z on the reduced X . These types of dependencies are supported by frameworks applied to large bodies of cases throughout statistical disparity [5]. Consequently, debiasing through both the direct and indirect effects enforces procedural fairness in decision-making by statistically mitigating the dependency of the sensitive feature X on the prediction Z . It therefore guarantees outcome fairness among sensitive groups [46].

In particular, similar to [3, 4], the Hilbert-Schmidt Independence Criterion (HSIC) function measures the degree of independence between two continuous random variables. With radial basis function kernel k_1 and k_2 , $\text{HSIC}^{k_1, k_2}(Z, \sigma(X)) = 0$ if and only if $Z \perp\!\!\!\perp \sigma(X)$, where $\sigma(\cdot)$ is 2D average pooling operation. As shown in Figure 2, $Z \equiv h(\mathbf{x}; \boldsymbol{\theta})$ represents the evidence vector predicted by the network, and X indicates the context feature from the backbone. The debiasing constraint takes the form

$$g(\boldsymbol{\theta}) \equiv \text{HSIC}(h(\mathbf{x}; \boldsymbol{\theta}), \sigma(X)), \quad (4)$$

which aims to reduce both direct and indirect dependency of predictive outcomes onto background context.

4. Optimization and Inference

In summary, we combine the aforementioned Equations (2) and (4) to formulate the optimization problem as

$$\min_{\boldsymbol{\theta} \in \Theta} \mathcal{L}_{Beta}(\boldsymbol{\theta}) \quad \text{subject to} \quad g(\boldsymbol{\theta}) \leq \gamma, \quad (5)$$

where $\gamma > 0$ is the independence criterion relaxation. Previous methods, such as regularization or projection, can handle the constraint in Equation (5) with near-optimal solutions but do not directly provide primal solutions. It may even fail to produce any useful information for static bias.

To better solve the problem, we develop the primal-dual average scheme update method. Specifically, we apply an averaging scheme to the primal sequence $\{\boldsymbol{\theta}^{(m)}\}_{m=1}^{\infty}$ to approximate primal optimal solutions, where m represents the index of an iteration. In particular, the sequence $\{\tilde{\boldsymbol{\theta}}^{(m)}\}_{m=1}^{\infty}$ is defined as the averages of the previous vectors through $\boldsymbol{\theta}^{(0)}$ to $\boldsymbol{\theta}^{(m-1)}$, i.e.,

$$\tilde{\boldsymbol{\theta}}^{(m)} = \frac{1}{m} \sum_{i=1}^{m-1} \boldsymbol{\theta}^{(i)}, \quad \forall m \geq 1. \quad (6)$$

The primal feasible iterate $\boldsymbol{\theta}^{(m)}$ is given in Equation (7). To simplify, we abuse the notation $\mathcal{L}_{Beta}(\boldsymbol{\theta})$ with $\mathcal{L}(\boldsymbol{\theta})$ in the

rest of the paper.

$$\boldsymbol{\theta}^{(m)} \leftarrow \arg \min_{\boldsymbol{\theta} \in \Theta} \left\{ \mathcal{L}(\boldsymbol{\theta}^{(m-1)}) + \lambda^{(m-1)} \left(g(\boldsymbol{\theta}^{(m-1)}) - \gamma \right) - \frac{\delta}{2} \left(\lambda^{(m-1)} \right)^2 \right\}, \quad (7)$$

where $\delta > 0$ is a constant determined by analysis. Accordingly, the parameter λ in dual solutions is updated as

$$\lambda^{(m)} \leftarrow \max \left\{ \left[\lambda^{(m-1)} + \eta_2 \left(g(\tilde{\boldsymbol{\theta}}^{(m)}) - \gamma - \delta \lambda^{(m-1)} \right) \right], 0 \right\}, \quad (8)$$

where $\eta_2 > 0$ is a constant learning rate of the dual step. By updating the dual parameter, λ , our optimization efficiently approaches the optimal model $\boldsymbol{\theta}^*$ arbitrarily close within a small finite number of steps m . For better understanding, the above algorithm is summarized in Algorithm 1.

4.1. Theoretical Analysis

To derive the bound on the feasibility violation and the primal cost of the running averages, we analyze the statistical guarantees of the solutions in Equations (7) and (8). We first make the following assumption.

Assumption 1. (*Regularity and Feasibility*). *The convex set Θ is compact (i.e., closed and bounded). For any $\boldsymbol{\theta} \in \Theta$, $\mathcal{L}(\boldsymbol{\theta})$ and $g(\boldsymbol{\theta})$ are convex real-valued and bounded functions, where $\inf_{\boldsymbol{\theta} \in \Theta} g(\boldsymbol{\theta}) = 0$ and, for any $\boldsymbol{\theta} \notin \Theta$, $\text{dom}(g(\boldsymbol{\theta})) = \emptyset$.*

Recall that in the proposed Algorithm 1 used to approximate pairs of primal-dual parameters at each iteration m , for the averaged primal sequence $\{\tilde{\boldsymbol{\theta}}^{(m)}\}$, we show that it always converges when Θ is compact.

Proposition 1. (*Convergence of Averaged Primal Sequence*) *Under Assumption 1, when the convex set Θ is compact, let the approximate primal sequence $\{\tilde{\boldsymbol{\theta}}^{(m)}\}_{m=1}^{\infty}$ be the running averages of the primal iterates in Equation (6). Then $\{\tilde{\boldsymbol{\theta}}^{(m)}\}_{m=1}^{\infty}$ can converge to its limit $\tilde{\boldsymbol{\theta}}^*$.*

Next we provide bounds on the feasibility violation $g(\tilde{\boldsymbol{\theta}}^{(m)})$ and the primal cost of the running averages $\mathcal{L}(\tilde{\boldsymbol{\theta}}^{(m)})$, where the bounds are given per iteration m .

Proposition 2. (*Bounds for $\mathcal{L}(\tilde{\boldsymbol{\theta}}^{(m)})$ and the violation of $g(\tilde{\boldsymbol{\theta}}^{(m)})$*) [28] *Let the dual sequence $\{\lambda^{(m)}\}_{m=1}^{\infty}$ be generated through Equation (8) and $\{\tilde{\boldsymbol{\theta}}^{(m)}\}_{m=1}^{\infty}$ be the averages in Equation (6). Under Assumption 1, we have*

1. *An upper bound on the amount of constraint violation of $\tilde{\boldsymbol{\theta}}^{(m)}$ that $\| [g(\tilde{\boldsymbol{\theta}}^{(m)})]_+ \| \leq \frac{\lambda^{(m)}}{m\eta_2}$.*

2. *An upper bound on $\mathcal{L}(\tilde{\boldsymbol{\theta}}^{(m)})$ that $\mathcal{L}(\tilde{\boldsymbol{\theta}}^{(m)}) \leq f^* + \frac{(\lambda^{(0)})^2}{2m\eta_2} + \frac{\eta_2 L^2}{2}$, where $\|g(\tilde{\boldsymbol{\theta}}^{(m)})\| < L$ and $L > 0$.*

3. *A lower bound $\mathcal{L}(\tilde{\boldsymbol{\theta}}^{(m)}) \geq f^* - \lambda^* \cdot \| [g(\tilde{\boldsymbol{\theta}}^{(m)})]_+ \|$.*

where $[u]_+$ denotes the projection of $[u]$ on the nonnegative orthant. f^* is the optimal solution of Equation (5) and λ^*

Algorithm 1 Primal-Dual Average Scheme Update

Input: $\boldsymbol{\theta}^{(0)} \in \Theta, \lambda^{(0)} \in \mathbb{R}_+$: primal and dual parameters

Require: $\eta_1, \eta_2 > 0$: learning rates

- 1: Initialize an empty buffer $B = []$ to store $\boldsymbol{\theta}^{(m)}$
 - 2: **for** $m = 1, 2, \dots$ **do**
 - 3: $L(\boldsymbol{\theta}, \lambda) := \mathcal{L}(\boldsymbol{\theta}) + \lambda(g(\boldsymbol{\theta}) - \gamma) - \frac{\delta}{2}\lambda^2$
 - 4: **Primal Update:**
 $\boldsymbol{\theta}^{(m)} \leftarrow \text{Adam}\left\{L\left(\boldsymbol{\theta}^{(m-1)}, \lambda^{(m-1)}\right), \eta_1, \boldsymbol{\theta}^{(m-1)}\right\}$
 - 5: Add $\boldsymbol{\theta}^{(m)}$ in B
 - 6: **Average Scheme:** $\tilde{\boldsymbol{\theta}}^{(m)} = \frac{1}{|B|} \sum_{i=1}^{|B|-1} \boldsymbol{\theta}^{(i)}$
 - 7: Update $\boldsymbol{\theta}^{(m)} \leftarrow \tilde{\boldsymbol{\theta}}^{(m)}$
 - 8: $L'(\boldsymbol{\theta}, \lambda) := \lambda + \eta_2(g(\boldsymbol{\theta}) - \gamma - \delta\lambda)$
 - 9: **Dual Update:**
 $\lambda^{(m)} \leftarrow \max\left\{L'\left(\tilde{\boldsymbol{\theta}}^{(m)}, \lambda^{(m-1)}\right), 0\right\}$
 - 10: **end for**
-

denotes the optimal value of the dual variable.

Propositions 1 and 2 demonstrate the convergence of the primal solution sequence and give bounds for both the loss function and debiasing constraint in MULE. The detailed proof is given in the Appendix.

4.2. Novelty Score Estimation

During inference, we aim to detect novel actor(s) with single or multiple unknown action(s). According to Equation (1), we develop four novelty quantification scores based on either uncertainty or belief of an actor. To this end, we incorporate the actor's estimated subjective opinions $\{\omega_i\}_{i=1}^K$ for its actions, where $\omega_i = (b_i, d_i, u_i, a_i)$.

Uncertainty-based novelty. As described in Section 3.2, the positive and negative evidence pair $\{(\alpha_i, \beta_i)\}_{i=1}^K$ are used to estimate uncertainty u of an actor with K actions. A value of u close to 1 indicates novelty. Three uncertainty-based novelty score estimation mechanisms are introduced by using positive (PE) or negative (NE) evidence only and aggregating them (PNE), i.e., PE: $u = \frac{2}{1 + \exp(\sum_{i=1}^K \alpha_i - K)}$; NE: $u = \frac{2}{1 + \exp(K - \sum_{i=1}^K \beta_i)} - 1$; PNE: $u = \frac{2K}{\sum_{i=1}^K (\alpha_i + \beta_i)}$.

Belief-based novelty. Another novelty detection scheme is to estimate its belief value b using the binomial multiplication operator (denoted as $*$) [18] for all actions, i.e., $b = b_1 * \dots * b_K$, where $b_i * b_j := b_i + b_j - b_i \cdot b_j, \forall i, j \in \{1, \dots, K\}, i \neq j$. $b_i \in [0, 1]$ is a class-wise belief estimated using its corresponding α_i and β_i in Equation (1). A belief value b close to 0 indicates novelty.

4.3. Relation with Existing Evidential Learning

Although our method shares the basic concept of evidential learning with DEAR [4], it has a significant difference in three aspects.

- *The Beta distribution in Equation (2) is generalized from the Dirichlet distribution in ENNs.* In other words, to

detect an actor with multiple novel actions, we improve original ENNs [4, 36] by using K Beta distributions. If $K = 1$, our loss function is reduced to the counterpart in DEAR [4] (see the proof in the Appendix).

- *The M -EDC in Equation (4) simultaneously mitigates direct and indirect bias.* In contrast, DEAR [4] only considers dependencies of prediction Z on the sensitive feature X through the causal path $X \rightarrow H$. This is viewed as a sub-path of the indirect dependency $X \rightarrow H \rightarrow Z$, resulting in inferior performance.
- *Our optimization method in Algorithm 1 provides optimal hyper-parameter search for robust learning.* DEAR [4] views the debiasing constraint as a regularization term with the empirically set Lagrange multiplier; while an optimal multiplier λ^* is automatically found as a dual variable by Algorithm 1. Iterative update between the primal and dual variable guarantees the achievement of a small duality gap.

In summary, previous ENNs focus on detecting actors with a single action, and assume action probability follows a prior Dirichlet distribution. In contrast, in our MULE, Beta distributions are more general to adapt to different applicable scenarios where a video contains either a single ($N = 1$) or multiple ($N > 1$) actor(s) associated with a single ($K = 1$) or multiple ($K > 1$) action(s).

5. Experiments

Our method is implemented based on PyTorch [30]. All models are trained on 4 NVIDIA Quadro RTX 8000 GPUs.

Datasets. AVA [16] is a video dataset for spatio-temporal localizing atomic visual actions. It contains 430 videos, each with 15 minutes annotated in 1-second intervals. We use version 2.2 of the AVA dataset by default. Charades [38] contains 9,848 videos that average 30s in length. This dataset includes 157 multi-label, daily indoor activities.

Implementation details. Similar to [29], we employ COCO [24] pre-trained Faster R-CNN [32] with a ResNeXt-101-FPN [23] backbone to extract actor and object detections. To extract context features, Kinetics [19] pre-trained¹ SlowFast networks [12] are used as the backbone. For AVA [16], the inputs are 64-frame clips, where we sample $T = 8$ frames with a temporal stride $\tau = 8$ for the slow pathway, and ζT frames, where $\zeta = 4$, for the fast pathway. For Charades [38], the temporal sampling for the slow pathway is changed to 8×4 , and the fast pathway takes as an input 32 continuous frames. We train all models end-to-end using SGD with a batch size of 32. The learning rates in the optimization are set as $\eta_1 = 0.008$, $\eta_2 = 0.001$ empirically. We used both ground-truth boxes and predicted

¹The pretraining dataset contains slight overlap with the unseen actions. However, for a fair comparison with previous works, we use the same Kinetics-based backbone.

human boxes from [44] for training, and scale the shorter side of input frames to 256 pixels. The number of actors and objects N, M is determined by the boxes with scores greater than 0.85 for action detection.

5.1. Open- and Closed-Set Settings

Since the above datasets are used for traditional action recognition, we re-split them to adapt to our problem in this work. For open-set settings, videos are evenly divided into three disjoint sets $\mathcal{Z}_1, \mathcal{Z}_2$, and \mathcal{Z}_3 . We only include actions falling in $\mathcal{Z}_1 \cup \mathcal{Z}_2$ for training and actions in $\mathcal{Z}_2 \cup \mathcal{Z}_3$ for testing. Thus \mathcal{Z}_2 and \mathcal{Z}_3 are a set of known actions and novel actions in inference, respectively. In practice, each subset in AVA [16] contains 20 actions, while actions in Charades [38] are evenly divided into three subsets (52/52/53). An actor is considered as novelty (unknown) if it does not contain any action in the training data. To detect actors with novel actions, in the testing stage, we ensure that each actor contains ground-truth actions in either \mathcal{Z}_2 or \mathcal{Z}_3 exclusively [43]. We hence assign each actor in testing videos with a binary novelty label $\{0, 1\}$, where 0 indicates an actor with all known actions in \mathcal{Z}_2 and correspondingly, 1 indicates an actor with all unknown actions in \mathcal{Z}_3 .

Closed set action recognition refers to classifying actions into pre-defined categories. Following [29], the closed set studies on AVA [16] use 235 videos to train and test on 131 videos with known actions. For Charades [38], following [47], we use the officially provided train-test split (7, 985/1, 863) to evaluate the network where all actions are known. As this work focuses on open set action recognition, closed set accuracy is for reference only.

Evaluation metrics. Similar to [22, 25], we adopt the following four metrics to evaluate the performance on novel action detection, *i.e.*, estimate if the action of an actor is novel or not. 1) **Detection Error** [22] measures the misclassification probability when True Positive Rate (TPR) is 95%. The definition of an error P_e is given by $P_e = 0.5 \cdot (1 - \text{TPR}) + 0.5 \cdot \text{FPR}$, where FPR stands for False Positive Rate. 2) **AUROC** [9] is the Area Under the Receiver Operating Characteristic curve, which depicts the relation between TPR and FPR. A perfect detector corresponds to an AUROC score of 1. 3) **AUPR** [27] is the Area under the Precision-Recall curve. The PR curve is a graph showing the precision and recall against each other. 4) **FPR at 95% TPR** [22] can be interpreted as the probability that a novel example is misclassified as known when TPR is 95%. Additionally, we report the Mean Average Precision (**mAP**) for K -class classification in closed set.

5.2. Ablation Study

To further explore our method, we conduct a detailed ablation study on AVA [16]. In the following tables, evaluation metrics with “ \uparrow ” indicate the larger the better, and “ \downarrow ”

m	PE / NE / PNE / Belief					hours per epoch
	Error↓	AUROC↑		AUPR↑	FPR at 95% TPR↓	
0	31.25 / 35.89 / 25.17 / 40.82	75.23 / 70.01 / 75.04 / 69.45		86.10 / 83.65 / 87.33 / 90.87	7.32 / 9.21 / 8.98 / 8.81	4
1	11.21 / 41.20 / 11.16 / 33.33	86.42 / 61.26 / 85.72 / 73.34		94.23 / 96.47 / 99.43 / 97.07	4.32 / 5.01 / 5.46 / 9.09	5
2	11.22 / 40.14 / 12.18 / 27.91	86.92 / 63.54 / 85.25 / 83.41		96.18 / 89.90 / 93.46 / 90.90	4.98 / 5.09 / 4.51 / 9.15	9
3	11.01 / 45.05 / 11.17 / 27.86	86.81 / 58.23 / 85.25 / 85.18		98.48 / 95.65 / 99.41 / 90.44	4.58 / 5.00 / 3.98 / 9.09	13
4	11.28 / 47.12 / 11.17 / 27.52	86.56 / 59.01 / 85.13 / 84.72		99.43 / 88.98 / 99.40 / 90.42	4.64 / 5.02 / 3.21 / 9.09	18
5	10.22 / 44.32 / 11.15 / 27.49	86.86 / 58.87 / 85.30 / 84.66		99.52 / 96.35 / 99.41 / 90.41	4.32 / 5.00 / 3.09 / 9.09	23

Table 1. Exploration of number of average primal-dual updating step m on AVA [16].

Method	PE / NE / PNE / Belief					Closed Set mAP↑
	Error↓	AUROC↑		AUPR↑	FPR at 95% TPR↓	
MULE, R-50	11.22 / 40.14 / 12.18 / 27.91	86.92 / 63.54 / 85.25 / 83.41		96.18 / 89.90 / 93.46 / 90.90	4.98 / 5.09 / 4.51 / 9.15	27.80
w/o ACO-R	50.12 / 50.04 / 12.45 / 39.84	51.23 / 51.12 / 84.88 / 82.41		86.11 / 86.36 / 88.13 / 92.35	15.23 / 15.56 / 34.44 / 9.57	25.12
w/o Beta-ENN	35.12 / 35.31 / 35.12 / 34.78	57.10 / 57.59 / 57.88 / 57.01		85.71 / 85.65 / 85.66 / 85.66	12.03 / 13.00 / 14.98 / 15.32	28.81
w/o M-EDC	13.16 / 46.32 / 12.16 / 38.23	86.00 / 50.13 / 85.12 / 83.05		95.12 / 90.67 / 92.45 / 89.19	6.31 / 5.05 / 5.01 / 9.59	27.16

Table 2. Exploration of different component in MULE with $m = 2$ on AVA [16].

Methods	Pre-train	Error↓	AUROC↑	AUPR↑	Closed Set mAP↑	
					FPR at 95% TPR	mAP↑
Slowfast, R-101 [12]	K600	60.12	50.15	70.15	20.17	29.00
ACAR, R-50 [29]	K400	35.16	52.89	79.15	14.16	28.84
ACAR, R-101 [29]	K700	32.26	55.18	82.15	10.16	33.30
AFAC, R-101 [47]	K600	53.14	79.69	90.79	7.15	30.20
AIA, R-101 [41]	K700	35.14	54.17	78.49	10.15	32.30
DEAR, R-50 [4]	K400	23.22	82.12	83.15	8.45	18.51
MULE (Ours), R-50	K400	11.22	86.92	96.18	4.98	27.80
MULE (Ours), R-101	K700	10.12	88.75	98.18	4.17	29.87

Table 3. Comparison with state-of-the-art on AVA [16].

indicate the smaller the better.

Effectiveness of optimization algorithm. We investigate the effectiveness of applying an averaging scheme to the primal sequence in Algorithm 1. According to Table 1, the larger the primal-dual updating step m , the better performance and the lower efficiency. If $m = 0$, it indicates that we do not use the proposed optimization method. The dual parameter λ in Equation (7) is then viewed as a Lagrangian multiplier and set empirically. The results demonstrate the effectiveness of our algorithm in reducing static bias. Considering the trade-off between performance and efficiency, we use $m = 2$ in the following experiments.

Contribution of important components. In Table 2, we discuss the contribution of each component of MULE on both open- and closed-set settings as follows.

- **ACO-R.** Without ACO-R representation, we only rely on the video backbone (*i.e.*, SlowFast [12]) and a single layer action classifier to predict actions in the video. The results show a significant performance drop in both closed set mAP (25.12% vs. 27.80%) and open set metrics, demonstrating the importance of ACO-R representation.
- **Beta-ENN.** If we use Beta loss in Equation (2) in the network, the open set recognition performance is improved with a slight sacrifice of the closed set mAP (28.81% vs. 27.80%). This is because each class follows a binomial opinion whose conjugate prior is a Beta distribution to benefit open set recognition, rather than a Dirichlet dis-

tribution aiming to optimize multi-class classification in a closed-set setting. Although using the Beta loss does not outperform using cross-entropy in a closed-set setting, its performance is still competitive.

- **M-EDC.** If we remove the evidence debiasing constraint in network training, our method will suffer from static bias, as explained in Section 3.3. The results show that using our M-EDC module in training brings slight improvement on all four novelty scores, showing the effectiveness to model static bias. This is because the cases with the same action (*e.g.*, “ride”) and different contexts (“bike” or “horse”) cases are only 1% of the dataset. Empirically, we set $\gamma = 0.001$ in Equation (5).

Comparison with novelty estimation mechanisms. As presented in Tables 1 and 2, we compare four novelty estimation mechanisms in Section 4.2. It can be concluded that uncertainty-based scores perform better than belief based one. According to Equation (1), a belief based score is calculated based on negative evidence. We speculate that it is more intuitive and accurate to estimate if an actor performs the same action rather than different actions. In summary, positive evidence is more reliable than negative evidence. In the following, we only report Positive Evidence (PE) scores to detect novel actions for clarity.

5.3. Results Analysis

For a fair comparison with previous single-actor/single-action methods, we enhance them to work within the new multi-actor/multi-action paradigm. Specifically, we calculate positive evidence by the evidence function $s(\cdot)$ in Section 3.2, based on the last layer outputs $h(\mathbf{x}; \theta)$. We apply single-actor based DEAR [4] to a bounding volume around each detected actor for a multi-actor setting.

According to Table 3, our method is compared with existing methods on AVA [16]. The baseline SlowFast [12] achieves a reasonable closed set mAP score but fails in open set recognition. Other methods including ACAR [29],




Figure 3. Visual comparison with our method and state-of-the-art on AVA [16] and Charades [38]. Cyan and (yellow) boxes denote the predictions of actors with known and novel actions, respectively. ✓ marks and X marks indicate correct and false predictions, respectively.

Methods	Pre-train	Error↓	AUROC↑	AUPR↑	FPR at↓ 95% TPR	Closed Set mAP↑
Slowfast, R-101 [12]	K600	25.15	79.12	79.15	50.15	45.20
X3D-XL [11]	K600	8.45	82.15	86.49	39.98	47.20
AFAC, R-101 [47]	K600	7.00	82.20	90.15	35.19	48.10
AFAC, CSN-152 [47]	IG-65M	6.18	80.12	90.79	30.15	50.30
CSN, CSN-152 [42]	IG-65M	6.89	80.15	92.61	35.16	46.40
DEAR, R-50 [4]	K400	12.15	86.15	92.35	29.96	38.12
MULE (Ours), R-50	K400	6.15	85.49	90.78	25.98	45.33
MULE (Ours), R-101	K700	6.23	88.49	91.15	25.15	47.21

Table 4. Comparison with state-of-the-art on Charades [38].

AFAC [47], and AIA [41] obtain better accuracy in the closed set but are still unsatisfying to detect novel actions. Compared with DEAR [4] using ENNs, our method achieves much better performance in both closed set and open set metrics. This is because our network can handle the situation that each actor may contain more than one action better than simple bounding volume strategy. It indicates the effectiveness of the proposed Beta-ENN for multi-label open set action recognition.

We also observe a similar trend on Charades [38]. From Table 4, our MULE achieves comparable closed set mAP and much better open set performance compared with the state-of-the-art. Note that AFAC [47] and CSN [42] achieve the best closed set accuracy by using powerful CSN-152 backbone [47] and IG-65M pre-trained dataset [15]. However, our method using ResNet backbone still obtains a considerable gain in terms of AUROC and FPR at 95% TPR. It is a good fit for the situation that requires lower probability of misclassifying a novel instance as known.

In Figure 3, the visual examples show that other methods like DEAR [4], AFAC [47] and SlowFast [12] output several false predictions of known and novel actions in a multi-actor/multi-action setting. In contrast, our method can handle novelty detection in different situations more accurately.

Results on different splits. Furthermore, we report the results of exchanging the training and testing sets in the dataset split. In Table 5, slightly worse performance indicates that the testing set is more challenging in this split. However, the results are similar to the original split and our

Methods	Pre-train	Error↓	AUROC↑	AUPR↑	FPR at↓ 95% TPR	Closed Set mAP↑
DEAR, R-50 [4]	K400	30.30	79.98	78.59	8.77	17.32
MULE (Ours), R-50	K400	14.92	84.34	90.91	6.02	23.65

Table 5. Comparison with state-of-the-art on AVA [16] on a different split.

method still outperforms the best competitor, DEAR [4]. It indicates that our method is robust on different splits.

6. Conclusion

In this paper, we address the general problem of detecting the novelty of each actor’s action(s) in video for the first time. To this end, we develop a new open set action recognition framework using Multi-Label Evidential Learning (MULE). Based on actor-context-object relation representation, the proposed Beta Evidential Neural Network can formulate Beta distributions for single/multi-actor, single/multi-action settings. To optimize the network, we add the multi-label evidence debiasing constraint and propose the primal-dual average scheme update learning method with theoretical analysis. Thus both direct and indirect dependencies of action predictions are reduced for more robust results. The experiments show that Beta distribution more accurately classifies novel action(s) than the previous Dirichlet distribution.

Acknowledgements. The authors would like to thank Dr. Xujiang Zhao and Ameya Shringi for their helpful discussion and technical advice. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0055 and 140D0422C0043. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA); or its Contracting Agent, the U.S. Department of the Interior, Interior Business Center, Acquisition Services Directorate, Division V.

References

- [1] J. K. Aggarwal and Michael S. Ryoo. Human activity analysis: A review. *ACM Comput. Surv.*, 43(3):16:1–16:43, 2011. 1
- [2] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. In *NeurIPS*, 2020. 2, 3
- [3] Hyojin Bahng, Sanghyuk Chun, Sangdoon Yun, Jaegul Choo, and Seong Joon Oh. Learning de-biased representations with biased representations. In *ICML*, pages 528–539, 2020. 4
- [4] Wentao Bao, Qi Yu, and Yu Kong. Evidential deep learning for open set action recognition. In *ICCV*, pages 13329–13338, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [5] Solon Barocas and Andrew D. Selbst. Big data’s disparate impact. *California Law Review* 104, 2016. 4
- [6] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *CVPR*, pages 1563–1572, 2016. 1, 2
- [7] Jinwoo Choi, Chen Gao, Joseph C. E. Messou, and Jia-Bin Huang. Why can’t I dance in the mall? learning to mitigate scene bias in action recognition. In *NeurIPS*, pages 851–863, 2019. 2
- [8] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *ICLR*, 2016. 4
- [9] Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. In *ICML*, pages 233–240, 2006. 6
- [10] Dawei Du, Ameeya Shringi, Anthony Hoogs, and Christopher Funk. Reconstructing humpty dumpty: Multi-feature graph autoencoder for open set action recognition. In *WACV*, pages 3360–3369, 2023. 1
- [11] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *CVPR*, pages 203–213, 2020. 2, 8
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6201–6210, 2019. 2, 6, 7, 8
- [13] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, pages 1050–1059, 2016. 1, 3
- [14] Chuanxing Geng, Sheng-Jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *TPAMI*, 43(10):3614–3631, 2021. 1
- [15] Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *CVPR*, pages 12046–12055, 2019. 8
- [16] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, pages 6047–6056, 2018. 1, 2, 6, 7, 8
- [17] Audun Jøsang. Subjective logic: A formalism for reasoning under uncertainty. *Springer Verlag*, 2016. 1, 3
- [18] Audun Jøsang. Belief calculus. *arXiv preprint arXiv:0606029*, 2006. 5
- [19] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017. 6
- [20] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *NeurIPS*, pages 5574–5584, 2017. 2
- [21] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *ECCV*, pages 513–528, 2018. 2, 4
- [22] Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018. 1, 3, 6
- [23] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 6
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014. 6
- [25] Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *NeurIPS*, 2020. 1, 3, 6
- [26] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheroface: Deep hypersphere embedding for face recognition. In *CVPR*, pages 212–220, 2017. 1
- [27] Chris Manning and Hinrich Schütze. Foundations of statistical natural language processing. *MIT Press*, 1999. 6
- [28] Angelia Nedic and Asuman E. Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM J. Optim.*, 19(4):1757–1780, 2009. 5
- [29] Junting Pan, Siyu Chen, Mike Zheng Shou, Yu Liu, Jing Shao, and Hongsheng Li. Actor-context-actor relation network for spatio-temporal action localization. In *CVPR*, pages 464–474, 2021. 3, 6, 7
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019. 6
- [31] Alessandro Prest, Vittorio Ferrari, and Cordelia Schmid. Explicit modeling of human-object interactions in realistic videos. *TPAMI*, 35(4):835–848, 2013. 2
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, volume 28, 2015. 3, 6
- [33] Alina Roitberg, Ziad Al-Halah, and Rainer Stiefelhagen. Informed democracy: Voting-based novelty detection for action recognition. In *BMVC*, page 52, 2018. 2

- [34] Alina Roitberg, Chaoxiang Ma, Monica Haurilet, and Rainer Stiefelwagen. Open set driver activity recognition. In *IV*, pages 1048–1053, 2020. [1](#)
- [35] Ethan M. Rudd, Lalit P. Jain, Walter J. Scheirer, and Terrence E. Boulton. The extreme value machine. *TPAMI*, 40(3):762–768, 2018. [2](#)
- [36] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. In *NeurIPS*, volume 31, 2018. [1](#), [2](#), [3](#), [4](#), [6](#)
- [37] Gaurav Sharma, Frédéric Jurie, and Cordelia Schmid. Expanded parts model for semantic description of humans in still images. *TPAMI*, 39(1):87–101, 2017. [2](#)
- [38] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, volume 9905, pages 510–526, 2016. [1](#), [2](#), [6](#), [8](#)
- [39] Konstantin Sozykin, Stanislav Protasov, Adil Khan, Rasheed Hussain, and Jooyoung Lee. Multi-label class-imbalanced action recognition in hockey videos via 3d convolutional neural networks. In *SNPD*, pages 146–151, 2018. [1](#)
- [40] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *ECCV*, pages 318–334, 2018. [3](#)
- [41] Jiajun Tang, Jin Xia, Xinzhi Mu, Bo Pang, and Cewu Lu. Asynchronous interaction aggregation for action detection. In *ECCV*, pages 71–87, 2020. [7](#), [8](#)
- [42] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *ICCV*, pages 5552–5561, 2019. [8](#)
- [43] Haoran Wang, Weitang Liu, Alex Bocchieri, and Yixuan Li. Can multi-label classification networks know what they don't know? In *NeurIPS*, pages 29074–29087, 2021. [1](#), [6](#)
- [44] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, pages 284–293, 2019. [6](#)
- [45] Ronald R. Yager and Liping Liu, editors. *Classic Works of the Dempster-Shafer Theory of Belief Functions*, volume 219 of *Studies in Fuzziness and Soft Computing*. Springer, 2008. [1](#), [3](#)
- [46] Junzhe Zhang and Elias Bareinboim. Fairness in decision-making - the causal explanation formula. In *AAAI*, pages 2037–2045, 2018. [4](#)
- [47] Yanyi Zhang, Xinyu Li, and Ivan Marsic. Multi-label activity recognition using activity-specific features and activity correlations. In *CVPR*, pages 14625–14635, 2021. [2](#), [6](#), [7](#), [8](#)
- [48] Chen Zhao and Feng Chen. Rank-based multi-task learning for fair regression. *IEEE International Conference on Data Mining (ICDM)*, 2019. [4](#)
- [49] Chen Zhao, Feng Chen, and Bhavani Thuraisingham. Fairness-aware online meta-learning. In *KDD*, pages 2294–2304, 2021. [4](#)
- [50] Chen Zhao, Feng Mi, Xintao Wu, Kai Jiang, Latifur Khan, and Feng Chen. Adaptive fairness-aware online meta-learning for changing environments. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 2565–2575, 2022. [4](#)