

Deep Mixture Density Network for Probabilistic Object Detection

Yihui He* and Jianren Wang*

Abstract—Mistakes/uncertainties in object detection could lead to catastrophes when deploying robots in the real world. In this paper, we measure the uncertainties of object localization to minimize this kind of risk. Uncertainties emerge upon challenging cases like occlusion. The bounding box borders of an occluded object can have multiple plausible configurations. We propose a deep multivariate mixture of Gaussians model for probabilistic object detection. The covariances help to learn the relationship between the borders, and the mixture components potentially learn different configurations of an occluded part. Quantitatively, our model improves the AP of the baselines by 3.9% and 1.4% on CrowdHuman and MS-COCO respectively with almost no computational or memory overhead. Qualitatively, our model enjoys explainability since the resulting covariance matrices and the mixture components help measure uncertainties.

I. INTRODUCTION

Object detection provides crucial perception information for real-world applications like robotics grasping [1] and self-driving cars [2]. Mistakes/uncertainties in object detection could potentially risk the success of the robot’s operation and even human lives [3], [4], [5]. In this paper, we minimize this kind of risk by measuring uncertainties of bounding box localization. The uncertainties are quantified by a simple Gaussian mixture model.

Currently, there are two styles of bounding box annotation among the large-scale object detection datasets: (1) visible box that only contains visible parts (*e.g.*, MS-COCO [6] and PASCAL VOC [7]) (2) full box that contains both visible and occluded parts (*e.g.*, CrowdHuman [8] and VehicleOcclusion [9]). For full box annotation, regressing a single set of bounding box coordinates works well for fully visible objects, since it is a unimodal problem. However, when an object is occluded, we observe that its occluded parts can have several plausible configurations (*e.g.*, Figure 1 (b)), which is a multimodal problem. Even for visible box annotation, an object sometimes still exhibits multiple modes due to inaccurate labeling (*e.g.*, Figure 1 (c) vs. (d)). We argue that an object detector robust to occlusion should learn a multimodal distribution with the capability of proposing more than one plausible hypothesis for the configuration of an occluded part.

Besides, we also observe that the bounding box coordinates have correlations by nature. Take Figure 1 (c) as an example, by knowing the position of the car’s roof, we can easily infer the location of the left border even without looking at it. Therefore, an object detector robust to occlusion

also needs to be capable of inferring the correlations between the occluded bounding box borders and the visible ones.

Motivated by these two observations, we propose a deep multivariate mixture of Gaussians model for object detection under occlusion. Concretely, instead of regressing a single set of bounding box coordinates, our model regresses several sets of coordinates, which are the means of the Gaussians. Moreover, we learn a covariance matrix for the coordinates of each Gaussian mixture component. These components are summed together as the prediction for the distribution of plausible bounding box configurations. At inference time, we choose the expectation of our model’s distribution as the final predicted bounding box.

To demonstrate the generalizability of our proposed model, we conduct experiments on four datasets: CrowdHuman, MS-COCO, VehicleOcclusion, and PASCAL VOC. Quantitatively, our model improves the AP (Average Precision) of the baselines by 3.9% and 1.4% on CrowdHuman and MS-COCO respectively (Table I and Table II). Qualitatively, our model enjoys explainability since the resulting bounding boxes can be interpreted using the covariance matrices and the Gaussian mixture components. More importantly, our model is almost computation and memory free, since predicting the mixture components only requires a fully-connected layer, and we can discard the covariance matrices at inference time (Table VI).

II. APPROACH

A. Formulation

We observe that when an object is partially occluded, the occluded bounding box border can usually be inferred to some extent by other visible parts of the object (*e.g.*, it is easy to infer the left border of the car given the car roof position in Figure 1 (c)). Besides, the occluded bounding box exhibits multiple modes. For example, the left arm of the teddy bear could have several possible configurations in Figure 1 (b). Motivated by these two observations, we propose to estimate the bounding box coordinates as a probability distribution during bounding box regression instead of a set of deterministic coordinates. Specifically, we propose to estimate a multivariate mixture of Gaussians distribution with a deep network [10]. Multivariate Gaussian helps the case where bounding box borders have correlations, and a mixture of Gaussians helps the case where an occluded bounding box border exhibits multiple modes. Formally, we predict the distribution $p_{\theta}(\mathbf{x}|\mathbf{I})$ given the feature maps \mathbf{I} of a region of interest (RoI). The distribution is parameterized by θ , which is a neural network (*e.g.*, Faster R-CNN head, Figure 2). The distribution has K components $\mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$.

see arxiv.org/abs/1911.10614 for up-to-date version.

*indicates equal contribution. Yihui He and Jianren Wang are Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, PA 15213, USA. {he2@alumni, jianrenw@andrew}.cmu.edu




Fig. 1: We observe that an occluded bounding box usually exhibits multiple modes in most detection datasets, no matter whether the ground truth annotation is visible box or full box: (a) visible bounding box annotation (b) full object bounding box labeled by different annotators (c) visible bounding box annotated accurately (d) visible bounding box annotated inaccurately




Fig. 2: Faster R-CNN head architecture for our approach: We extended the existing Faster R-CNN head to predict the parameters of multivariate mixture of Gaussian μ , ϕ and Σ

Each component i has mean $\mu_{i=1\dots K} = [x_1, y_1, x_2, y_2]^T$, which is the most probable bounding box coordinates relative to the RoI, estimated by the component:

$$p_{\theta}(\mathbf{x}|\mathbf{I}) = \sum_{i=1}^K \phi_i \mathcal{N}(\mu_i, \Sigma_i)$$

$$\text{where } \sum_{i=1}^K \phi_i = 1 \text{ and } 0 \leq \phi_i \leq 1$$

$$\mathcal{N}(\mu_i, \Sigma_i) = \frac{\exp\left(-\frac{1}{2}(\mathbf{x} - \mu_i)^T \Sigma_i^{-1} (\mathbf{x} - \mu_i)\right)}{\sqrt{(2\pi)^4 |\Sigma_i|}}$$
(1)

where each $\mathcal{N}(\mu_i, \Sigma_i)$ is a multivariate Gaussian distribution. ϕ_i is a mixture weight scalar for $\mathcal{N}(\mu_i, \Sigma_i)$. $|\Sigma_i|$ is the determinant of Σ_i . Σ is the covariance matrix, which is a symmetric semi-positive definite matrix in general. To be able to compute the inverse Σ^{-1} , we constrain the covariance

matrix to be a symmetric positive definite matrix. In this case, the precision matrix Σ^{-1} is also a symmetric positive definite matrix. During training, the model estimates the precision matrix Σ^{-1} instead of the covariance matrix Σ , so that we do not need to compute the inverse every time during training which we also find more stable in our experiments. To ensure the properties of the precision matrix Σ^{-1} , we parameterize it using the Cholesky decomposition:

$$\Sigma^{-1} = \mathbf{U}^T \mathbf{U} \quad (2)$$

where \mathbf{U} is an upper triangular matrix with strictly positive diagonal entries, such that Cholesky decomposition is guaranteed to be unique:

$$\mathbf{U} = \begin{bmatrix} \exp(u_{11}) & u_{12} & u_{13} & u_{14} \\ & \exp(u_{22}) & u_{23} & u_{24} \\ & & \exp(u_{33}) & u_{34} \\ & & & \exp(u_{44}) \end{bmatrix}$$

$$|\Sigma| = \frac{1}{|\Sigma^{-1}|} = \frac{1}{|\mathbf{U}^T \mathbf{U}|} = \frac{1}{|\mathbf{U}^T| |\mathbf{U}|} = \frac{1}{|\mathbf{U}|^2}$$

$$= \frac{1}{\exp(\sum_{i=1}^4 u_{ii})^2}$$
(3)

We parameterize the mixture weights ϕ_i using Softmax, so that they range from 0 to 1 and sum to 1:

$$\phi_i = \frac{\exp(z_i)}{\sum_{k=1}^K \exp(z_k)} \quad (4)$$

z_i , u_{ij} and μ_i are outputs produced by a fully-connected layer on top of the final fully-connected layer fc7 on the Faster R-CNN head. Take Faster R-CNN with RPN as an example, Figure 2 shows the architecture of our model. Since we only modify a small part of the architecture, our approach might also be applied to other object detectors than Faster R-CNN, like one-stage object detectors YOLO and RetinaNet.

a) *Learning*: Our model parameterizes the distribution over bounding boxes using a neural network which depends on RoI features. During training, we estimate the parameters θ with maximum likelihood estimation on a given dataset $\{\mathbf{I}_{\ell}, \mu_{\ell}^* | \ell = 1, 2, \dots, N\}$, where μ_{ℓ}^* represents the ground truth

coordinates for RoI feature maps \mathbf{I}_ℓ and N is the number of observations:

$$\hat{\theta} = \arg \max_{\theta} \frac{1}{N} \sum_{\ell} \ln p_{\theta}(\boldsymbol{\mu}_{\ell}^* | \mathbf{I}_{\ell}) \quad (5)$$

In practice, N is the number of samples in a mini-batch. We use momentum stochastic gradient descent (SGD) to minimize the localization loss L_{loc} and the classification loss L_{cls} :

$$L = L_{cls} + \lambda L_{loc} \quad (6)$$

where $L_{loc} = -\frac{1}{N} \sum_{\ell} \ln p_{\theta}(\boldsymbol{\mu}_{\ell}^* | \mathbf{I}_{\ell})$

Note that we use different parameters θ for different classes in practice. For simplicity, the formulation above only considers the regression problem for a single class.

b) Inference: During testing, we perform inference in two ways: The most obvious one is **average**. we use the expectation of our mixture module as prediction:

$$\mathbb{E}[p_{\theta}(\mathbf{x} | \mathbf{I})] = \frac{1}{K} \sum_{i=1}^K \phi_i \boldsymbol{\mu}_i \quad (7)$$

Notice that the covariance matrix $\boldsymbol{\Sigma}_i$ is not involved in inference. In practice, we discard the neurons that produce the covariance matrix to speed up inference. In our experiments (Table VI), our model has almost the same inference latency and memory consumption as the baseline network.

The second way is most **probable** inference, where we use the prediction from the mixture component with highest confidence:

$$\boldsymbol{\mu}_i \text{ where } i = \arg \max_i \phi_i \quad (8)$$

This is not robust in practice. We propose to use the inference when the ϕ_i from the most confident component (i) is highest than $\phi_j + t$ from the second most confident component (j). $t > 0$ is a threshold, (e.g., 0.4). If this is not satisfied, we use the average inference.

B. Degradation Cases

a) Multivariate Gaussian: When the number of mixture components $K = 1$, our model degrades into a multivariate Gaussian model. And the localization loss can be rewritten as follow (for simplicity, we only illustrate the loss for a single sample ℓ):

$$L_{loc}^{\ell} = \frac{(\boldsymbol{\mu}^* - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}^* - \boldsymbol{\mu})}{2} + \frac{\log |\boldsymbol{\Sigma}|}{2} + 2 \ln 2\pi \quad (9)$$

$$= \frac{(\boldsymbol{\mu}^* - \boldsymbol{\mu})^T \mathbf{U}^T \mathbf{U} (\boldsymbol{\mu}^* - \boldsymbol{\mu})}{2} - \sum_{i=1}^4 u_{ii} + 2 \ln 2\pi$$

where $2 \ln 2\pi$ is a constant which can be ignored during training. Multivariate Gaussian model is helpful under occlusion since the borders of a bounding box have correlations with each other inherently. For example, by looking at the location of a car's door, we can guess the location of its roof even if it is occluded.

TABLE I: CrowdHuman on ResNet-50 FPN Faster R-CNN

method	mAP
RetinaNet [12]	80.8
FPN [13]	85.0
FPN+Gaussian	86.2
FPN+mixture of 8 Gaussian (average)	87.7
FPN+mixture of 8 Gaussian (probable)	87.8
FPN+multivariate Gaussian	88.3
FPN+multivariate mixture of 8 Gaussian (average)	88.8
FPN+multivariate mixture of 8 Gaussian (probable)	88.9

b) Mixture of Gaussians: When the covariance matrix is constrained to be a diagonal matrix, our model becomes a mixture of Gaussians model with independent variables:

$$L_{loc}^{\ell} = -\ln \sum_{i=1}^K \phi_i \sum_{j=1}^4 \frac{\exp\left(-\frac{(\boldsymbol{\mu}_{ij}^* - \boldsymbol{\mu}_{ij})^2 / 2 (\mathbf{U}_i)_{jj}^2}{\sqrt{2\pi} (\mathbf{U}_i)_{jj}}\right)}{\sqrt{2\pi} (\mathbf{U}_i)_{jj}} \quad (10)$$

where $(\mathbf{U}_i)_{jj}$ is the j th diagonal element of the matrix \mathbf{U}_i . Multimodality is helpful under occlusion because an occluded object usually has multiple modes.

c) Gaussian: When the number of mixture components $K = 1$ and the covariance is constrained to be a diagonal matrix, it becomes a simple Gaussian model where different variables are independent:

$$L_{loc}^{\ell} = \sum_{j=1}^4 (\mathbf{U})_{jj}^2 \frac{(\boldsymbol{\mu}_j^* - \boldsymbol{\mu}_j)^2}{2} - \ln (\mathbf{U})_{jj} + \frac{\ln 2\pi}{2} \quad (11)$$

We argue that this simple model helps detection in most cases. Here $(\mathbf{U})_{jj}$ behaves like a balancing term. When the bounding box regression is inaccurate (large $(\boldsymbol{\mu}_j^* - \boldsymbol{\mu}_j)^2 / 2$), the variance $1 / (\mathbf{U})_{jj}^2$ tends to be larger. Therefore smaller gradient will be provided to bounding box regression $(\mathbf{U})_{jj}^2 (\boldsymbol{\mu}_j^* - \boldsymbol{\mu}_j)^2 / 2$ in this case, which might help training the network (Table I and Table II). If bounding box regression is perfect, \mathbf{U} tend to infinity (i.e., the variance should be close 0). However, regression is not that accurate in practice, \mathbf{U} will be punished for being too large.

d) Euclidean Loss: When all the diagonal elements $(\mathbf{U})_{jj}$ are one ($u_{jj} = 0$), our model degenerates to the standard euclidean loss:

$$L_{loc}^{\ell} = \sum_{j=1}^4 \frac{(\boldsymbol{\mu}_j^* - \boldsymbol{\mu}_j)^2}{2} + \frac{\ln 2\pi}{2} \quad (12)$$

III. EXPERIMENTS

We initialize the weights of $\boldsymbol{\mu}_i$, z_i and u_{ii} layers (Figure 2) using random Gaussian initialization with standard deviations 0.0001 and biases 0, -1 and 0 respectively. So that at the start of training, bounding box coordinate $\boldsymbol{\mu}_i$ is at an unbiased position, \mathbf{U}_i is an identity matrix and ϕ_i treats each mixture component equally. Our model can be trained end-to-end. Unless specified, we follow settings in Detectron [11] and those original papers.

A. Datasets

To demonstrate the generalizability of our method, we conduct experiments on four datasets:

TABLE II: Performance of our models on MS-COCO on ResNeXt-101 FPN Faster R-CNN

method	AP	AP ⁵⁰	AP ⁷⁵	AP ^S	AP ^M	AP ^L
baseline	40.8	62.5	44.4	23.4	44.4	53.8
Gaussian	41.2	61.2	44.3	23.0	44.7	54.4
mixture of 8 Gaussian (average)	41.4	61.2	44.7	23.6	44.6	55.6
mixture of 8 Gaussian (probable)	41.6	61.2	44.9	23.6	44.7	55.5
multivariate Gaussian	41.5	61.5	44.7	23.6	45.0	55.2
multivariate mixture of 8 Gaussian (average)	42.0	61.6	45.2	23.7	45.1	56.0
multivariate mixture of 8 Gaussian (probable)	42.3	61.6	45.4	23.8	45.2	56.2

TABLE III: Probabilistic object detectors comparison. Legend: moLRP = Localization Recall Precision [14], PDQ = Probability-based Detection Quality measure [5].

Approach (τ)	mAP (%)	moLRP (%)	PDQ (%)
MC-Dropout SSD (0.5) [15]	15.8	84.4	12.8
MC-Dropout SSD (0.05) [15]	19.5	83.4	1.3
SSD-300 (0.5) [16]	15.0	85.7	3.9
SSD-300 (0.05) [16]	19.3	84.0	0.6
YOLOv3 (0.5) [17]	29.7	69.2	5.7
YOLOv3 (0.05) [17]	30.1	72.3	3.3
FRCNN R (0.5) [18]	32.8	70.9	6.7
FRCNN R (0.05) [18]	34.3	70.9	3.0
FRCNN R+FPN (0.5) [19]	34.6	68.8	11.8
FRCNN R+FPN (0.05) [19]	37.0	69.6	4.2
FRCNN X+FPN (0.5) [19]	37.4	67.3	11.9
FRCNN X+FPN (0.05) [19]	39.0	67.9	4.4
probFRCNN (0.5) [5]	35.5	67.8	29.4
FRCNN+multivariate Gaussian (probable)	41.6	65.3	28.2

a) *CrowdHuman*: [8] is a large, rich-annotated and highly diverse dataset for better evaluation of detectors in crowd scenarios. Its training and validation sets contain a total of 470k human instances, and around 22.6 persons every image under various kinds of occlusions. The annotations for occluded bounding boxes are full boxes (Figure 1 (b)) instead of visible boxes (Figure 1 (a)). The experiments are in Table I.

b) *VehicleOcclusion*: is a synthetic dataset designed for object detection under occlusion [9]. Same as above, the annotations are full boxes. The occlusion annotations are more accurate since the occluders (occluding objects) are randomly placed on the annotated visible object. It contains six types of vehicles and occluded instances at various difficulty levels. Specifically, it consists of four occlusion levels: No occlusion (0%), L1 (20% ~ 40%), L2 (40% ~ 60%), L3 (60% ~ 80%). The percentages are computed by pixels. At level L1, L2 and L3, there are two, three, and four occluders placed on the object, respectively (Table V).

c) *MS-COCO*: [6] is a large-scale object detection dataset containing 80 object categories, 330k images (> 200k labeled) and 1.5 million object instances. Compared with the two datasets above, MS-COCO has fewer occlusion cases. For example, the IoU (intersection over union) between overlapped human bounding boxes in MS-COCO are less than 0.7 [8]. We use train2017 for training and val2017 for testing (Table II). Different from above, the annotations are visible boxes.

d) *PASCAL VOC*: is a classic dataset for object detection [20]. Similar with MS-COCO, this dataset has less

TABLE IV: comparison on PASCAL VOC with YOLO-v3

method	mAP
baseline	79.3
Gaussian	80.8
mixture of 8 Gaussian (average)	80.9
mixture of 8 Gaussian (probable)	81.1
multivariate Gaussian	81.2
multivariate mixture of 8 Gaussian (average)	81.3
multivariate mixture of 8 Gaussian (probable)	81.5




Fig. 3: AP/AP⁵⁰ when varying # mixture components

occlusion cases than the first two datasets. We use voc_2007 and voc_2012 for training and voc_2007_test for testing (Table IV). The annotations are visible boxes.

B. Ablation Study

a) *Number of Mixture Components*: Shown in Figure 3, we test our mixture of Gaussians model by varying the number of mixture components. The baseline is ResNet-50 FPN Faster R-CNN [21], [13] on CrowdHuman. As the number of components increases from 1, 4 to 8, we observe consistent performance improvement. The mixture of eight Gaussians model (Equation 10) outperforms Gaussian model (Equation 11) by 1% AP. However, the performance goes down when there are more than 16 components. This might be because the objects in the dataset might not have as many as 16 modes when occluded. Besides, the more components we have, the higher the chance of over-fitting. Unless specified, we use eight components for the mixture of Gaussians model.

b) *Average v.s. Probable*: Shown in Figure 4, we test ResNet-50 FPN Faster R-CNN mixture of eight Gaussians




Fig. 4: AP by varying the threshold t of probable inference

model with probable inference by varying its threshold t . The threshold indeed has a sweet spot. We used $t = 0.4$ by default for probable inference for the rest of our paper. Not shown on the figure, when $t = 0$, probable inference is very unstable, which only achieves 33.45%. Also shown in Table I and Table II, probable inference consistently performs better than average inference.

c) Mixture of Gaussian vs. Multivariate Gaussian:

Shown in Table I and II, we compare the degradation cases of our complete model (Equation 1): Gaussian (Equation 11), mixture of Gaussians (Equation 10) and multivariate Gaussian (Equation 9) on CrowdHuman and MS-COCO. For CrowdHuman, we use ResNet-50 FPN Faster R-CNN as the baseline. For MS-COCO, we use ResNeXt-101 (64x4d) FPN Faster R-CNN.

On CrowdHuman which has a lot of crowded scenes, our model greatly improves the baseline. Gaussian improves the baseline by **1.2%** mAP. A mixture of eight Gaussians improves **2.7%** mAP, and multivariate Gaussians improves **3.3%** mAP. The complete model improves the performance by **3.9%** mAP. The improvements indicate all these assumptions are helpful under heavy occlusion. Gaussian helps training the regression network by learning to decrease the gradients for high variance cases. Multivariate Gaussian helps to learn the correlations between an occluded border and the visible borders. Mixture of Gaussians helps to learn a multimodal model for the occluded cases which have multiple modes.

On MS-COCO, the bounding box annotations are visible boxes instead of full boxes used in CrowdHuman. Gaussian still works here which improves the baseline by **0.4%** AP, since there are variances in the dataset caused by inaccurate annotation (*e.g.*, Figure 1 (d)). Gaussian helps to reduce the gradients for these ambiguous cases. A mixture of eight Gaussians improves **0.8%** AP, and multivariate Gaussians improves **0.7%** AP. The complete model improves the performance by **1.4%** AP. The improvements are noticeable, however less significant than on CrowdHuman. On the one

TABLE V: Comparison with a state-of-the-art occlusion-aware detector on VehicleOcclusion. The metric is mAP. occ.: occlusion

	no occ.	L1	L2	L3
baseline	73.6	48.3	35.0	23.0
DeepVoting	72.0	53.7	42.6	31.6
DeepVoting+	74.0	58.0	46.9	35.2
Ours (average)	74.4	62.1	50.9	38.4
Ours (probable)	74.4	62.2	51.0	38.6

TABLE VI: Model size and FPS comparison with ResNet-50 FPN Mask R-CNN on a single GPU

method	# params	FPS
baseline	91M	11.1
Gaussian	91M	11.1
multivariate Gaussian	91M	11.1
mixture of 8 Gaussian	93M	10.2
multivariate mixture of 8	93M	10.2

hand, there are fewer occluded instances in MS-COCO, multimodality and covariances might not be as helpful as in CrowdHuman. On the other hand, predicting full boxes require guessing the invisible parts where multimodality and covariances are more useful.

We further conduct experiments on PASCAL VOC, shown in Table IV. YOLO-v3 [17] is the baseline. Similar to MS-COCO, the bounding box annotations are visible boxes instead of full boxes used in CrowdHuman. We observe that Gaussian improve the mAP (mean Average Precision) by **1.5%**. The complete model improves the mAP by **2.2%**. Multimodality and multivariate Gaussian do not substantially improve the performance. These observations coincide with the observations on MS-COCO.

d) Comparison with State-of-the-art: Shown in Table V, we compare multivariate mixture of eight Gaussians model to DeepVoting [22] on VehicleOcclusion. Similar to CrowdHuman, the bounding box annotations are full boxes. The baseline is VGG-16 Faster R-CNN.

Our multivariate mixture of eight Gaussians model outperforms DeepVoting by a large margin at different occlusion levels. Without occlusion, our model also helps to learn a better detector, coinciding the experiments above. We argue that our model considers multiple modes of an object and the correlations between each border of a bounding box, which helps detection under occlusion.

e) Model Size and Inference Speed: We measure the inference speed of our models using ResNet-50 FPN Mask R-CNN with a TITAN Xp, CUDA 10.1 and cuDNN 7.5.0 on MS-COCO val2017 [23], [24], [25], [26]. Shown in Table VI, Gaussian (Equation 11) and multivariate Gaussian (Equation 9) neither slow down the inference nor increase the number of parameters, since we can discard the covariance Σ at inference time (Section II-A). The complete model, multivariate mixture of eight Gaussians (Equation 1), only increases **2M** parameters and sacrifices **0.9 FPS** on GPU. Our models outperform the baselines by large margins (Table I, II and V), while requires almost no additional computation and memory.

Note that we measure the inference latency on MS-COCO where there are 80 classes, such that the number of parameters for μ is $1024 \times 80 \times K$ (1024 is the number of output channels of fc7, Figure 2). On CrowdHuman where there is only one class (human), the number of parameters for μ is only $1024 \times K$, which will consume even fewer computation and memory resources.

f) *probabilistic detection challenge*: [5] proposed a new metric (PDQ) for measuring probabilistic detection quality on MS-COCO val2017. We follow [5] and trained a ResNeXt FPN Faster RCNN with our multivariate Gaussian model (probable, Table II). Shown in Table III, our approach outperformed [5] by a large margin in every metric.

IV. CONCLUSION

We propose a multivariate mixture of Gaussians model for object detection under occlusion. Quantitatively, it demonstrates consistent improvements over the baselines among MS-COCO, PASCAL VOC, CrowdHuman, and VehicleOcclusion. Qualitatively, our model enjoys explainability as the detection results can be diagnosed via the covariance matrices and the mixture components.

REFERENCES

- [1] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine, "More than a feeling: Learning to grasp and regrasp using vision and touch," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3300–3307, 2018.
- [2] B. Huval, T. Wang, S. Tandon, J. Kiske, W. Song, J. Pazhayampallil, M. Andriluka, P. Rajpurkar, T. Migimatsu, R. Cheng-Yue, et al., "An empirical evaluation of deep learning on highway driving," *arXiv preprint arXiv:1504.01716*, 2015.
- [3] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in ai safety," *arXiv preprint arXiv:1606.06565*, 2016.
- [4] K. R. Varshney and H. Alemzadeh, "On the safety of machine learning: Cyber-physical systems, decision sciences, and data products," *Big data*, vol. 5, no. 3, pp. 246–255, 2017.
- [5] D. Hall, F. Dayoub, J. Skinner, H. Zhang, D. Miller, P. Corke, G. Carneiro, A. Angelova, and N. Sünderhauf, "Probabilistic object detection: Definition and evaluation," in *The IEEE Winter Conference on Applications of Computer Vision*, 2020, pp. 1031–1040.
- [6] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, Cham, 2014, pp. 740–755.
- [7] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [8] S. Shao, Z. Zhao, B. Li, T. Xiao, G. Yu, X. Zhang, and J. Sun, "Crowdhuman: A benchmark for detecting human in a crowd," *arXiv preprint arXiv:1805.00123*, 2018.
- [9] J. Wang, C. Xie, Z. Zhang, J. Zhu, L. Xie, and A. Yuille, "Detecting semantic parts on partially occluded objects," *arXiv preprint arXiv:1707.07819*, 2017.
- [10] C. M. Bishop, "Mixture density networks," 1994.
- [11] R. Girshick, I. Radosavovic, G. Gkioxari, P. Dollár, and K. He, "Detectron," <https://github.com/facebookresearch/detectron>, 2018.
- [12] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [13] T.-Y. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie, "Feature pyramid networks for object detection," in *CVPR*, 2017, p. 3.
- [14] K. Oksuz, B. Can Cam, E. Akbas, and S. Kalkan, "Localization recall precision (lrp): A new performance metric for object detection," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 504–519.
- [15] D. Miller, F. Dayoub, M. Milford, and N. Sünderhauf, "Evaluating Merging Strategies for Sampling-based Uncertainty Techniques in Object Detection," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [16] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [17] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [18] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [19] F. Massa and R. Girshick, "maskrcnn-benchmark: Fast, modular reference implementation of Instance Segmentation and Object Detection algorithms in PyTorch," <https://github.com/facebookresearch/maskrcnn-benchmark>, 2018, accessed: [Insert date here].
- [20] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results," <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [22] Z. Zhang, C. Xie, J. Wang, L. Xie, and A. L. Yuille, "Deepvoting: A robust and explainable deep network for semantic part detection under partial occlusion," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1372–1380.
- [23] Y. He, X. Zhang, and J. Sun, "Channel pruning for accelerating very deep neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1389–1397.
- [24] Y. He, J. Lin, Z. Liu, H. Wang, L.-J. Li, and S. Han, "Amc: Auttml for model compression and acceleration on mobile devices," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 784–800.
- [25] H. Phan, Y. He, M. Savvides, Z. Shen, et al., "Mobinet: A mobile binary network for image classification," in *The IEEE Winter Conference on Applications of Computer Vision*, 2020, pp. 3453–3462.
- [26] Y. Liang, Z. Yang, K. Zhang, Y. He, J. Wang, and N. Zheng, "Single image super-resolution via a lightweight residual convolutional neural network," *arXiv preprint arXiv:1703.08173*, 2017.
- [27] J. Choi, D. Chun, H. Kim, and H.-J. Lee, "Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 502–511.
- [28] C. Zhu, Y. He, and M. Savvides, "Feature selective anchor-free module for single-shot object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 840–849.
- [29] J. Wang and Y. He, "Physics-aware 3d mesh synthesis," in *2019 International Conference on 3D Vision (3DV)*, 2019, pp. 502–512.
- [30] J. Wang, Z. Fang, and H. Zhao, "Alignnet: A unifying approach to audio-visual alignment," in *The IEEE Winter Conference on Applications of Computer Vision (WACV)*, March 2020.
- [31] Y. He, R. Yan, K. Fragkiadaki, and S.-I. Yu, "Epipolar transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7779–7788.
- [32] Y. He, X. Liu, H. Zhong, and Y. Ma, "Addressnet: Shift-based primitives for efficient convolutional neural networks," in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 1213–1222.
- [33] S. Choi, K. Lee, S. Lim, and S. Oh, "Uncertainty-aware learning from demonstration using mixture density networks with sampling-free variance modeling," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6915–6922.