

Distributional Instance Segmentation: Modeling Uncertainty and High Confidence Predictions with Latent-MaskRCNN

YuXuan Liu^{1,2}, Nikhil Mishra^{1,2}, Pieter Abbeel^{1,2}, Xi Chen¹

Abstract—Object recognition and instance segmentation are fundamental skills in any robotic or autonomous system. Existing state-of-the-art methods are often unable to capture meaningful uncertainty in challenging or ambiguous scenes, and as such can cause critical errors in high-performance applications. In this paper, we explore a class of distributional instance segmentation models using latent codes that can model uncertainty over plausible hypotheses of object masks. For robotic picking applications, we propose a confidence mask method to achieve the high precision necessary in industrial use cases. We show that our method can significantly reduce critical errors in robotic systems, including our newly released dataset of ambiguous scenes in a robotic application. On a real-world apparel-picking robot, our method significantly reduces double pick errors while maintaining high performance.

I. INTRODUCTION

Instance segmentation is a fundamental problem in many real-world robotic systems. The goal of instance segmentation is to enumerate the objects (or *instances*) that appear in an image, specifying which pixels in the image belong to each object.

In the past few years, recent work has mostly focused on developing specialized architectures that make the instance segmentation task more amenable to deep learning. For example, *detect-then-segment* methods [1], [2], [3], [4], [5], [6], rely on a cascade of classification, regression, and filtering to first identify a *bounding box* for each instance (a related problem known as *object detection*), followed by an additional step to predict each instance’s mask given its bounding box. Another example is *pixel-embedding* methods [7], [8], which optimize pixel-level auxiliary tasks, and then use a specialized clustering procedure to extract instance predictions from the dense pixel representation.

We observe that existing methods are not well equipped to deal with the inherent ambiguity that exists in the real world. We posit that this stems from a phenomenon we describe as limited *distributional expressiveness*, namely, that most instance segmentation models are designed to predict only *one* possible segmentation hypothesis (a single set of objects). Making only a single prediction is limiting in terms of the accuracy attainable by high-performance autonomous systems: a robot picking application may only tolerate $< 1\%$ of errors caused by incorrect segmentation.

To overcome these limitations, we propose *distributional instance segmentation* which models a distribution over plausible hypotheses of objects. The key contributions of this work are:

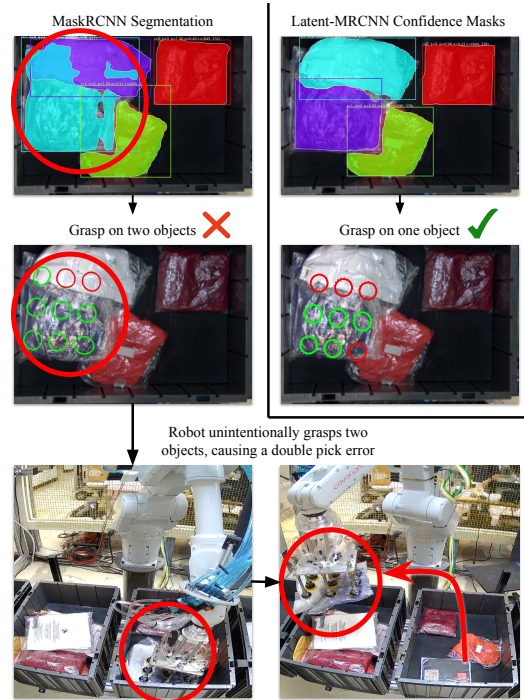


Fig. 1. Traditional instance segmentation models such as MaskRCNN cannot model uncertainty over object masks. For robotics, this can result in critical errors such as unintentionally picking two objects. Our Latent-MRCNN can predict multiple hypotheses of object masks and use these to make high-confidence predictions, reducing the rate of double pick errors.

1. We introduce a distributional instance segmentation model using latent codes, Latent-MaskRCNN, which can predict multiple hypotheses of object masks.
2. We propose new methods for using the output of a distributional instance segmentation model. For robotic applications, we propose high-precision predictions with Confidence Masks, and we achieve high recall with Union-NMS.
3. We are releasing a dataset of over 5000 annotated images from a real-world robotics application that highlights the ambiguity in instance segmentation. We show our method achieves high performance on this dataset as well as popular driving and instance segmentation datasets.
4. On a real-world apparel picking robot, our method can significantly reduce critical errors while achieving a high level of performance (Fig. 1).

¹Covariant.ai, ²University of California, Berkeley

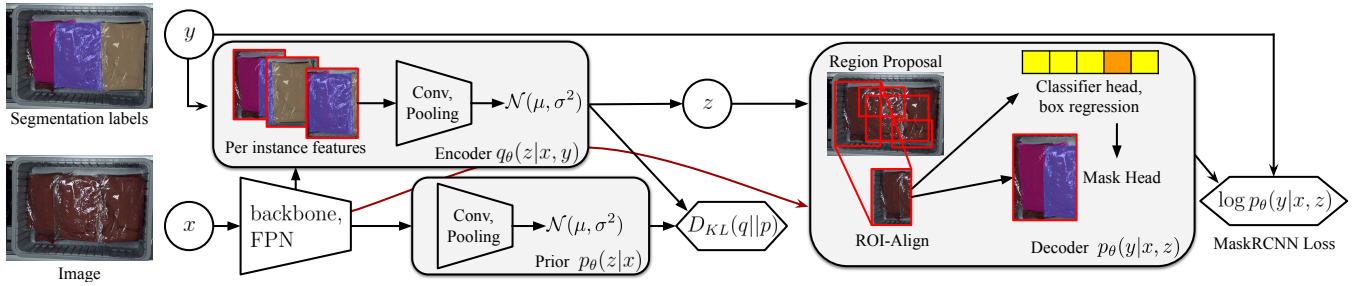


Fig. 2. Overview of Latent-MaskRCNN: At training time, the encoder q_θ uses features extracted from the image x and labels y to sample a latent code z which is passed into the decoder. The decoder conditions on z and uses a typical MaskRCNN architecture to predict masks including region proposal, classifier, box, and mask heads. At inference time, z is sampled through the prior $p_\theta(z|x)$ which only takes the image x as input. $D_{KL}(q||p)$ ensures the prior has good coverage over the latent space.

II. RELATED WORK

Detect-then-segment methods are the most popular instance segmentation methods, and MaskRCNN belongs to this category. While they all first perform object detection and then segment each instance given its bounding box, there are some variations. For example, YOLACT [9] follows the same structure as MaskRCNN, but uses YOLO [10] as the object detector instead of FasterRCNN [2]. YOLO is very similar to FasterRCNN, making architectural changes that sacrifice some accuracy in exchange for real-time inference speed. Thus, we expect YOLACT to have the same distributional limitations as MaskRCNN. Other methods [11], [12], [13] explore how to express uncertainty during the detection step, but they consider distributions of individual boxes rather than over sets of object masks.

Mask-proposal methods [14], [15], [16] aim to circumvent bounding boxes as an intermediate representation. They are structured like FasterRCNN [2], but propose masks directly. Empirically, they do not behave much differently than MaskRCNN. Distributionally, they suffer from many of the same limitations as MaskRCNN: each proposal still models each pixel independently of the others, and they still rely on NMS to filter proposals.

Pixel-embedding [7], [8], [17], [18] methods work in a substantially different way than either of the above two families. They generally optimize some auxiliary task that encourages pixels in the same instance to have similar representations. Then they rely on a clustering-based inference procedure to extract instance predictions from their pixelwise representations. However, their performance has lagged quite far behind that of detect-then-segment methods, which has made them relatively unpopular. They can model per-pixel uncertainty in a manner similar to a naive semantic segmentation method, but this is likely insufficient for distributional expressiveness.

A number of methods explore how to express uncertainty in other structured prediction tasks. However, many of these do so by training multiple replicas of the entire model or some subset of the parameters, and modifying the training objective in a way that encourages diversity amongst the replicas [19], [20], [21], [22]. This incurs a multiplicative increase in the computational cost and memory footprint

required at training time, which can be prohibitively expensive for large models. Other latent-variable formulations [23], [24], [25] offer improvements on medical *semantic* segmentation and video segmentation tasks. We find, however, that *instance* segmentation poses a richer set of challenges and has different application-specific uses.

III. DISTRIBUTIONAL INSTANCE SEGMENTATION WITH LATENT VARIABLES

A. Latent Variable Formulation

How can we turn instance segmentation into distributionally expressive models, while retaining the inductive biases of existing model architectures? Drawing on prior work in variational inference [23], [26] we consider a latent-variable formulation where we incorporate latent codes in the style of a variational autoencoder. If we adopt this framework, then an instance segmentation model becomes a conditional VAE that is trained to maximize the evidence lower-bound:

$$\log p(y|x) \geq \mathbb{E}_{z \sim q} [\log p(y|x, z)] - D_{KL}(q||p(z|x)) \quad (1)$$

Typically $q(z|y, x)$ is known as the encoder, $p(y|x, z)$ as the decoder, and $p(z|x)$ as the prior, and these components are all learned to maximize the lower-bound. The decoder is essentially an instance segmentation model in the traditional sense, except that it is augmented to additionally consume a latent code z . This general technique allows us to reuse any existing instance segmentation model to implement our decoder (and train it in the same way), with only a slight modification to incorporate z as an input.

During inference, we can sample from $p(y|x)$ by sampling different latent codes $z^{(k)} \sim p(z|x)$, and decoding them into different instance predictions $y^{(k)} \sim p(y|x, z)$. This can be quite powerful since we can now sample multiple structured and expressive hypotheses for a given image.

B. Latent-MaskRCNN

In principle, the latent-variable method can be applied to any existing instance segmentation model. In this section, we explore how it might be applied to MaskRCNN. We call the resulting model *Latent-MaskRCNN* (Fig. 2). We chose MaskRCNN since it is one of the most popular instance segmentation models and has served as the basis for most state-of-the-art methods in recent years.

The decoder of Latent-MaskRCNN uses the same architecture and training objective as MaskRCNN, with the main change being that it needs to incorporate latent codes. To allow them to influence as much of the prediction as possible, we want to do this relatively early in the model. We chose to inject the latent codes directly before region proposal, so that they can influence region proposal network, object detection head, and mask head. We tile the latent codes across the spatial dimensions of the image and concatenate them with the feature maps from the Feature Pyramid Network (FPN) [27]. Then we use a few convolutional layers to project the combined feature maps back down to their original channel dimensionality.

The encoder of Latent-MaskRCNN takes in an image x along with a set of ground-truth instances y , and produces a distribution over latent codes $q_\theta(z|y, x) = \mathcal{N}(\mu_\theta(y, x), \sigma_\theta^2(y, x))$. The architecture for $\mu_\theta(y, x)$ and $\sigma_\theta(y, x)$ takes inspiration from the mask head of MaskRCNN: it acts like a "reverse mask head" that operates on each ground-truth instance, and then pools features from across all instances. For each ground truth instance y_i , we extract ROI-aligned features from the FPN feature maps. Then we use a small CNN to embed each one into a single feature vector. At this point, we employ a graph neural network [28] to accumulate information from per-instance features since we need to a single latent code for the entire image. After several graph network layers, we mean-pool across the node features and use a fully-connected layer to produce a mean and log-variance for our latent distribution. The encoder is only used at training time since it has access to the ground truth mask labels.

At inference time, we must sample latent codes from the prior to produce mask samples. This prior takes in an image x and produces a distribution over latent codes $p_\theta(z|x) = \mathcal{N}(\mu_\theta(x), \sigma_\theta^2(x))$. We apply a few convolutional layers to the FPN feature maps, mean-pool across the spatial dimensions, and then predict a mean and log-variance using a small MLP. For all latent distributions, we use a fixed 64-dimension Gaussian with diagonal covariance.

For training, we use the encoder q_θ to sample latent codes z , which are passed to the Mask-RCNN decoder. We maximize the evidence lower bound (Equation 1) objective, where $-\log p(y|x, z) = \mathcal{L}_M(x, y, z)$ is the usual Mask-RCNN loss:

$$\mathcal{L}_M(x, y, z) = \mathcal{L}_{RPN} + \mathcal{L}_{cls} + \mathcal{L}_{box} + \mathcal{L}_{mask} \quad (2)$$

$D_{KL}(q||p(z|x))$ ensures the prior has good coverage over the encoder distribution. During inference, we sample latent codes from the prior (instead of from the encoder), but the decoder consumes them in the same way as during training.

We found it helpful to use a KL warm-up, as is a common practice for training VAEs [29]. The total training loss for Latent-MaskRCNN then becomes:

$$\mathcal{L}(x, y) = \mathbb{E}_{z \sim q}[\mathcal{L}_M(x, y, z)] + \beta D_{KL}(q||p(z|x)) \quad (3)$$

In the first part of training, we use $\beta = 0$ and increase β towards the end of training. This allows the latent code to

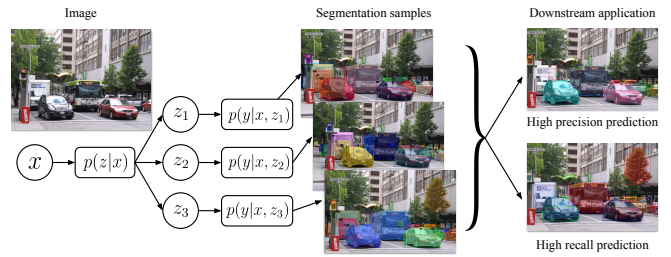


Fig. 3. At inference time, the encoder q_θ is discarded, and latent variables z_i are sampled from the image x conditioned prior $p_\theta(z|x)$. Each latent is decoded using $p_\theta(y|x, z_i)$ into a set of masks, which can be used for our high precision or recall predictions depending on the application.

encode useful information early on as the rest of the model is still learning; towards the end of the training, higher β pushes the latent space to be covered by the prior for better samples. For more details on our models and code, please refer to our website segm.yuxuanliu.com.

IV. APPLYING DISTRIBUTIONAL INSTANCE SEGMENTATION

Given a distributionally-expressive segmentation model, a natural question might be, how a downstream application can consume its distributional output? Instance segmentation often occurs at the beginning of the perception pipeline, and it's not immediately clear how samples from a distributional segmentation model can be used downstream. Moreover, each application may have varying error asymmetries: failing to detect an object can be catastrophic in autonomous driving but acceptable in robotic picking, while grouping two objects as one is a critical error in robotic picking but more reasonable in driving applications. In this section, we show how a single Latent-MaskRCNN model can be used flexibly across a number of applications with different requirements (Fig. 3).

A. High-Precision with p -Confidence Masks

In some applications, it can be very costly to make *under-segmentation* errors, when an instance's mask is predicted to be larger than it actually is. For example, consider a robotic manipulation application, where the robot must pick objects one at a time and feed them into a sortation process. If the model undersegments an instance, it may inadvertently pick multiple objects, which can be an expensive error for the downstream application. How can we ensure that these errors don't occur?

Suppose we draw several samples from Latent-MaskRCNN. If two pixels belong to the same instance mask in many samples, then we can be reasonably confident that they actually do belong to the same ground-truth instance. Drawing on this intuition, we can then compute a *p -confidence mask*, consisting of pixels that are all likely to be contained in a single ground-truth instance.

For a given confidence requirement p , we define a confidence mask c_p as a mask that is fully contained in a ground truth mask m with probability at least p : $\mathbb{P}(c_p \subseteq$




Fig. 4. Top: Confidence Mask predictions with different p . Notice that as the confidence requirement p increases, single objects can be split into two, ambiguous object extents are reduced, and uncertain objects are eliminated entirely. Bottom: Constructing an empirical confidence mask \hat{c}_p by taking intersection of samples m_1, m_2, m_3 . When an object’s extent is uncertain, a high confidence mask prediction will only consist of pixels that are highly likely to be contained within an object as determined by the samples.

$m) \geq p$. Using Latent-MaskRCNN, we can approximate this probability as:

$$\mathbb{P}(c_p \subseteq m) = \mathbb{E}_{p(m|x)}[\mathbb{1}\{c_p \subseteq m\}] \approx \frac{1}{k} \sum_{i=1}^k \mathbb{1}\{c_p \subseteq m_i\}$$

In the finite sample regime, \hat{c}_p is an empirical confidence mask if it is contained within a sampled mask for p fraction of the samples.

Now consider any subset of masks I consisting of one mask m_j from at least kp different samples. If we take the intersection of all of the masks in I , $\hat{c}_p = \bigcap_{m_j \in I} m_j$, then this intersection mask must be contained in each of the masks used in the intersection $\hat{c}_p \subseteq m_j$. Therefore we have:

$$\frac{1}{k} \sum_{m_j \in I} \mathbb{1}\{\hat{c}_p \subseteq m_j\} = \frac{|I|}{k} \geq p$$

and \hat{c}_p is an empirical confidence mask by construction.

Figure 4 illustrates confidence mask predictions for different p . Notice as the confidence requirement increases, the unconfident extents of the masks shrink, and some uncertain masks are eliminated. We can also see how constructing confidence masks via intersection leads to high confidence region predictions.

B. Scoring Confidence Masks

Since each confidence mask is an intersection of masks $c_p = \bigcap_I m_j$, how should we assign the score of a confidence mask prediction? One intuitive approach might be to take the average of the score s_i of each mask in the intersection: $\frac{1}{|I|} \sum s_j$. However, a confidence mask is not an average of masks but rather an intersection.

To formulate a better score for our confidence mask prediction, consider two scenarios. In the first scenario, the model is very confident about an object’s mask so it predicts roughly the same mask in every sample. The resulting confidence mask c_p has high IoU with each of the masks used in the sample m_j . On the other hand, consider an unconfident prediction where the object’s mask varies significantly across samples. Here, the confidence mask c_p represents a small

but confident region of the object whose extent is highly uncertain. The resulting IoU between c_p and each m_j will be smaller than when the model is confident and masks are not varying across samples.

We score each confidence mask as the mean score-weighted IoU between the predicted mask c_p and every mask used in the intersection:

$$s_{c_p} = \frac{1}{|I|} \sum_{m_j \in I} s_j \frac{|c_p \cap m_j|}{|c_p \cup m_j|}$$

When s_{c_p} is large, this indicates that c_p is a confident intersection of masks with very similar IoU. On the other hand, a small s_{c_p} indicates that c_p has low score or IoU with its samples and likely does not capture the full extent of the object well.

To predict a set of confidence masks for an image, we iteratively select the highest scoring confidence mask, excluding the pixels of all of the confidence masks predicted so far. This algorithm greedily approximates the maximum scoring confidence mask selection optimization.

C. High-Recall with Union-NMS

Other applications might be concerned about *over-segmentation*, the complement of under-segmentation. For example, in autonomous driving, failing to identify a pedestrian, or predicting them to be smaller than they actually are, can lead to a catastrophic error.

To make high-recall predictions with Latent-MaskRCNN, we use a procedure called *Union-NMS*. We first sample multiple segmentations from the model, and run NMS on the predicted masks. It checks if any two masks m_i, m_j have IoU greater than some threshold, and then discards the lower-scoring one. Suppose that some mask m_i remains after we perform NMS. Then Union-NMS returns the union of m_i with every mask that it suppressed, achieving higher recall by incorporating masks that would have otherwise been ignored.

D. Vanilla Prediction with the Prior Mean

Some applications may not have any specific performance requirements or may have strict inference time requirements. In these cases, a point estimate can be sufficient. With Latent-MaskRCNN, we can achieve this by always decoding the mean of the prior: $z = \mu_\theta(x)$ where $\mu_\theta(x)$ is the mean of the prior $p_\theta(z|x)$ (for Gaussian $p_\theta(z|x)$, it is also the mode of the distribution). We found that this scheme typically matches or yields a small improvement over MaskRCNN predictions, suggesting that Latent-MaskRCNN strictly increases the expressiveness of MaskRCNN and no performance is lost by using a more expressive distribution.

V. EXPERIMENTS

We conducted experiments seeking to answer the following questions:

1. Can Latent-MaskRCNN with Confidence Masks make high-precision predictions across a variety of datasets?
2. Can Union-NMS make high-recall predictions?
3. Can Latent-MaskRCNN reduce critical double pick errors in robotic picking applications?

TABLE I
EVALUATION OF MASKRCNN AND LATENT-MASKRCNN ACROSS VARIOUS DATASETS AND METRICS.

Method	COCO			Cityscapes			Apparel-5k		
	MR@HP	AR	mAP	MR@HP	AR	mAP	MR@HP	AR	mAP
MaskRCNN	20.0	66.0	35.0	25.3	55.1	35.8	23.6	39.9	26.9
Latent Union NMS	7.4	72.3	26.5	17.7	57.5	33.6	13.6	61.4	34.1
Latent Confidence Mask	22.0	48.1	30.5	28.0	48.6	34.1	42.4	41.8	35.1
Latent Prior Mean	19.5	65.8	35.3	25.7	53.8	35.0	26.9	49.4	34.3

A. Datasets

To help us answer these questions, we compared MaskRCNN and Latent-MaskRCNN across several datasets, each with its own set of challenges.

COCO [30]: This large dataset is the standard benchmark for instance segmentation. There are many object categories and a huge variety in image composition.

Cityscapes [31]: A real-world dataset from an autonomous driving application. Although it is smaller and more specialized than COCO, it is still a popular benchmark for instance segmentation. One notable challenge is that there are many background instances that are still important to segment (e.g. pedestrians), but the limited image resolution can introduce some uncertainty.

Apparel-5k: We collected this dataset of roughly 5000 images from a robot picking application. We use 4198 images in the training set and 463 in the validation set. There is only one object category, but the images exhibit a lot of inherent ambiguity due to complex occlusions, lighting, transparency, etc. We are releasing this dataset on our website segm.yuxuanliu.com for the broader community to build upon our work.

For each dataset, we trained both MaskRCNN and Latent-MaskRCNN on 8 GPUs using MaskRCNN’s released hyperparameters and training schedules. We use the same publicly available train/val splits for all experiments and datasets. We used a Resnet-50 backbone [32], initialized from pretrained-Imagenet [33] weights (for COCO) or pretrained COCO weights (for other datasets). Inference with MaskRCNN can take 80-100ms depending on the number of objects, and inference with Latent-MaskRCNN can take 500-1000ms depending on the number of samples and objects.

B. Evaluating p -Confidence Masks

In Section IV-A, we introduced high precision predictions with p -Confidence masks to address the problem of under-segmentation. In those cases, we care that predictions have high Intersection-over-Prediction: $\text{IoP}(m_i, g) = \frac{|m_i \cap g|}{|m_i|}$. When IoP is high, errors due to under-segmentation are less likely to occur.

When evaluating models in this regime, we need to trade off precision (in terms of IoP) with recall (to avoid degenerate solutions). To do this, we consider the *max recall at high precision* (MR@HP):

$$\text{MR@HP} = \frac{1}{|p| \cdot |\tau|} \sum_{p_i \in p, \tau_j \in \tau} \max_{t: \text{Precision}(\tau_j) \geq p_i} \text{Recall}(t, \tau_j)$$

For a given precision threshold p_i and IoP threshold τ_j , we can compute the max recall that each model achieves (or zero, if it never achieves precision p_i). The MR@HP metric is the average of these recalls, over a range of precision threshold p and IoP thresholds τ . For high precision use-cases, we care about performance at high values of these thresholds, therefore we use $p = \tau = [0.75, 0.8, 0.85, 0.9, 0.95]$.

In Table I, we evaluated Latent-MaskRCNN using both the prior-mean scheme from Section IV-D as well as confidence masks with a confidence level $p = 0.9$. Across all three datasets, we find that latent confidence masks yield the best performance in terms of MR@HP. As for mAP, we find that Latent Prior Mean can match if not exceed the performance of MaskRCNN on all three datasets. On the challenging Apparel-5k dataset, we find that Latent Confidence Mask and Latent Prior Mean significantly outperform MaskRCNN in terms of MR@HP and mAP. Overall, we find that Latent-MaskRCNN is a strict improvement over MaskRCNN by matching overall detection performance in terms of mAP and offering the best high-precision performance in terms of MR@HP.

C. Evaluating Union NMS

For the over-segmentation problem, we introduced the Union NMS method in Section IV-C. In such cases, we care that we have high recall (that we detect every instance that exists), and that each mask prediction has high IoG (intersection-over-ground-truth): $\text{IoG}(m_i, g) = \frac{|m_i \cap g|}{|g|}$. To capture both of these considerations, we consider the average-recall (AR) [30] using IoG. This measures both recall while also penalizing over-segmentation (are there any predicted masks that are too small).

We evaluated Latent-MaskRCNN using both its prior-mean (Section IV-D) and Union-NMS. In Table I, we show that the prior-mean predictions are similar to MaskRCNN on all three datasets, while Union-NMS achieves substantially higher AR (using IoG). This suggests that Latent-MaskRCNN with Union NMS can more effectively cover different modes of uncertainty, for high-recall applications.

D. Can confidence masks reduce double pick errors on Apparel-5K?

In a robotic picking application, it is costly for the robot to pick up two items accidentally, thinking it had only picked one item since it affects inventory counts and downstream orders. For the Apparel-5k dataset, we can estimate the double pick rate of a model’s segmentation prediction by




Fig. 5. Latent confidence masks achieve lower double pick rates and generally more pickable area compared to MaskRCNN.

approximating the robot’s gripper as a circle with a fixed radius in pixel space. Then, we randomly sample circles on the image and count the number of circles D that land within one predicted mask but more than one ground truth mask. We divide this by the number of circles N that land within one predicted mask, to arrive at the estimated double pick rate $R = \frac{D}{N}$. Empirically we find that this simulated double pick rate is correlated with double pick rates on a real robot.

Another metric that we are concerned with in industrial robot picking is pickable area, the amount of visible surfaces that the robot can pick from. A model that predicts higher, more accurate pickable areas enables the robot to have more flexibility in its grasping strategy. To this end, we compute the area of all the predicted masks over the area of all the ground truth masks, as the fraction of pickable area available.

We compare MaskRCNN and Latent-MaskRCNN with varying p -confidence masks in Figure 5. We find that Latent-MaskRCNN outperforms MaskRCNN in fraction of pickable area and double pick rate in all cases. Moreover, the tunable parameter p in Latent confidence masks allows for application-specific tradeoffs between double pick rate and pickable area. Higher values of p tend to correspond to lower double pick rate and less pickable area, as the confidence requirement for each prediction is increased. With traditional MaskRCNN, only one double pick rate and pickable area fraction is realizable since no tunable knob exists.

E. Can confidence masks reduce double pick errors on a real-world apparel-picking robot?

To evaluate whether our dataset evaluation translates to real-robot performance, we compare MaskRCNN and Latent-MaskRCNN on an apparel-picking robot. We use an ABB1300 with a 9-cup suction gripper to pick apparel items in polybags between two totes (Fig. 1). The robot uses two overhead camera systems to perform instance segmentation and then grasp point generation. The grasp points are optimized to land as many suction cups as possible on a single object detected by the segmentation model.

For our evaluation, we only change which segmentation model is used while holding other parts of the system

TABLE II
APPAREL-PICKING ROBOT EVALUATION. * INDICATES A STATISTICALLY SIGNIFICANT DIFFERENCE

Method	Double Pick Rate	Average Scaled Cups
MaskRCNN	4.40%*	4.76
Latent-MRCNN	0.82%*	4.91

constant, including hardware, object set, and grasp point generation. Each segmentation model is trained on the same Apparel-5k dataset. We run each model with several hundred grasps and record the number of double picks, grasps that unintentionally pick two objects. In an industrial warehouse application, these double picks are very costly errors since they result in incorrect inventory counts and cause errors in downstream sortation and order fulfillment systems. A typical high automation warehouse can tolerate at most 1% double pick rate before the robot is causing more problems than it solves.

We also measure the average number of sealed cups on a grasped item, since the suction holding force is proportional to the number of sealed cups. Grasps that use less sealed cups tend to result in more dropped objects, which leads to jams, lost inventory, and costly human intervention. A robotic system can reduce double pick rate by shrinking object mask sizes, chopping up bigger masks into smaller ones, or only using a single small suction cup. However, all of these approaches indiscriminately reduce the suction holding force on all items, whereas our approach will be conservative only when ambiguity is present.

Table II reports the results of our apparel-picking experiments. We find that Latent-MaskRCNN with 0.9-Confidence mask significantly reduces the double pick rate. This validates our simulated findings on Apparel-5K in Section V-D. Moreover, Latent-MaskRCNN achieves slightly better average number of sealed cups, suggesting that suction stability was not sacrificed. This suggests that our method can make high-confidence predictions and make the appropriate tradeoffs in the face of uncertainty.

VI. DISCUSSION

We proposed a new family of models that builds on top of existing instance segmentation models by using latent variables to achieve more distributional expressiveness. Latent-MaskRCNN can express a wide range of uncertainty where existing instance segmentation models often fall short. We can leverage uncertainty expressed by the model using Confidence Masks and Union-NMS to achieve high precision and high recall respectively. These methods demonstrate strong performance across robotics, autonomous driving, and general object datasets. On a real apparel-picking robot, we find that our model can significantly reduce the rate of critical errors while maintaining high performance. Finally, we have highlighted the importance of distributional expressiveness and hope that future work in instance segmentation can continue to build on top of our work and datasets shared in this paper.

REFERENCES

- [1] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comp. Vis.*, 2015, pp. 1440–1448.
- [2] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Proc. Advances in Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [3] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comp. Vis.*, 2017, pp. 2961–2969.
- [4] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014, pp. 580–587.
- [5] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *ECCV*, 2020.
- [6] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. M. Ni, and H.-Y. Shum, "Dino: Detr with improved denoising anchor boxes for end-to-end object detection," 2022.
- [7] A. Arnab and P. H. S. Torr, "Pixelwise instance segmentation with a dynamically instantiated network," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [8] M. Bai and R. Urtasun, "Deep watershed transform for instance segmentation," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2017, pp. 5221–5229.
- [9] D. Bolya, C. Zhou, F. Xiao, and Y. J. Lee, "Yolact: Real-time instance segmentation," *Proc. IEEE Int. Conf. Comp. Vis.*, 2019.
- [10] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016, pp. 779–788.
- [11] D. Hall, F. Dayoub, J. Skinner, H. Zhang, D. Miller, P. Corke, G. Carneiro, A. Angelova, and N. Sunderhauf, "Probabilistic object detection: Definition and evaluation," 03 2020, pp. 1020–1029.
- [12] A. Harakeh, M. Smart, and S. L. Waslander, "Bayesod: A bayesian approach for uncertainty estimation in deep object detectors," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 87–93, 2020.
- [13] D. Miller, N. Sunderhauf, H. Zhang, D. Hall, and F. Dayoub, "Benchmarking sampling-based probabilistic object detectors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [14] X. Chen, R. Girshick, K. He, and P. Dollár, "Tensormask: A foundation for dense object segmentation," *Proc. IEEE Int. Conf. Comp. Vis.*, 2019.
- [15] E. Xie, P. Sun, X. Song, W. Wang, X. Liu, D. Liang, C. Shen, and P. Luo, "Polarmask: Single shot instance segmentation with polar representation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [16] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," 2022.
- [17] D. Neven, B. D. Brabandere, M. Proesmans, and L. V. Gool, "Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [18] L. Rumberger, L. Mais, and D. Kainmueller, *Probabilistic Deep Learning for Instance Segmentation*, 01 2020, pp. 445–457.
- [19] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *arXiv preprint arXiv:1612.01474*, 2016.
- [20] A. Guzman-Rivera, D. Batra, and P. Kohli, "Multiple choice learning: Learning to produce multiple structured outputs," in *NIPS*, vol. 1, no. 2. Citeseer, 2012, p. 3.
- [21] C. Rupprecht, I. Laina, R. DiPietro, M. Baust, F. Tombari, N. Navab, and G. D. Hager, "Learning in an uncertain world: Representing ambiguity through multiple hypotheses," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3591–3600.
- [22] B.-B. Gao, C. Xing, C.-W. Xie, J. Wu, and X. Geng, "Deep label distribution learning with label ambiguity," *IEEE Transactions on Image Processing*, vol. 26, no. 6, pp. 2825–2838, 2017.
- [23] S. A. A. Kohl, B. Romera-Paredes, C. Meyer, J. D. Fauw, J. R. Ledsam, K. H. Maier-Hein, S. M. A. Eslami, D. J. Rezende, and O. Ronneberger, "A probabilistic u-net for segmentation of ambiguous images," *CoRR*, vol. abs/1806.05034, 2018. [Online]. Available: <http://arxiv.org/abs/1806.05034>
- [24] W. Jang, D. Wei, X. Zhang, B. D. Leahy, H. Y. Yang, J. Tompkin, D. Ben-Yosef, D. Needleman, and H. Pfister, "Learning vector quantized shape code for amodal blastomere instance segmentation," *CoRR*, vol. abs/2012.00985, 2020. [Online]. Available: <https://arxiv.org/abs/2012.00985>
- [25] C.-C. Lin, Y. Hung, R. Feris, and L. He, "Video instance segmentation tracking with a modified vae architecture," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 13 144–13 154.
- [26] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [27] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, July 2017.
- [28] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," 2019, cite arxiv:1901.00596Comment: updated tables and references. [Online]. Available: <http://arxiv.org/abs/1901.00596>
- [29] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-vae: Learning basic visual concepts with a constrained variational framework," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [Online]. Available: <https://openreview.net/forum?id=Sy2fzU9gl>
- [30] 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 *Proc. Eur. Conf. Comp. Vis.* Springer, 2014, pp. 740–755.
- [31] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016, pp. 3213–3223.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2016.
- [33] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* Ieee, 2009, pp. 248–255.